GENI System Requirements Document (SRD)

GDD-07-46

GENI: Global Environment for Network Innovations

Version 0.3 of April 19, 2007

Status: Draft (Version 0.3)

This version released in anticipation of the document review scheduled for late spring 2007.

This document was prepared based on input from the Research Coordination Working Group and the GENI Planning Committee.

Author:

Aaron Falk, *USC ISI*

Contributors:

John Wroclawski, *USC ISI*

Change Log

Version	Author	Date	Comments
0.1	Aaron Falk	Jan 26, 2007	Initial draft
0.2	Aaron Falk	March 4, 2007	Includes initial list of A-spec and B-spec requirements, coarse cut at mapping table.
0.3	Aaron Falk	April 19, 2007	Added research goals & requirements

This work is supported in part by NSF grant CNS-0631422.

1 Introduction

This document contains the GENI system requirements. GENI is intended to meet the needs of the networking and distributed systems research communities – its users. These research communities should be considered the *customers* of the GENI system, and the system needs to perform according to their expectations.

At this conceptual design stage, the GENI Research Coordination Working Group has produced a peer-reviewed Research Plan describing what types of activities the community of users desires GENI to support. This plan leads directly to two levels of *research* requirements, which are captured in Section 2 of this document. The first level of research requirement presented captures the high level objectives of the research enabled by GENI, in terms of broader impact on technical development and society as a whole. The second level outlines concrete research directions, areas, and capabilities identified as necessary to achieve these broader objectives, and thus as research areas to be supported by GENI.

The GENI Research Plan also summarizes the requirements placed on the GENI *facility* by the previously mentioned research objectives. The rationale for these facility requirements is directly linked to the research objectives outlined in earlier sections of the Research plan. Section 3 of this document translates the Research Plan's summary of needed *facility* capabilities into individual, verifiable, statements of requirement. These statements of top-level customer requirements are traditionally referred to as the *"A-spec"* in the aerospace community, and we use that terminology here.

For GENI, the A-spec should fully represent the needs of the research community, as well as other critical communities such as the community of GENI site hosts and operators. In other words, if a system that completely fulfilled the list of A-spec requirements was delivered, it should be satisfactory to all concerned.

During the design process, A-spec requirements drive the creation of a high-level, functional architecture, and the combination of the architecture and the A-spec is used by the design team to derive a set of more detailed system requirements, often known as the *B-spec*. Every requirement in the B-spec should map to one or more A-spec requirements. (If it is known that a B-spec requirement exists and there is no A-spec requirement to motivate it, it is an indication of either a missing A-spec requirement or an over-constrained design.) The A-spec catalogs what the customers are asking for; the B-spec catalogs what the designers are building. The mapping between the two sets of requirements should be clear. In this document, GENI's A-spec, or customer facility requirements, are presented in Section 3; a brief summary of the high-level architecture is given in Section 4; the B-spec, or derived requirements, in Section 5; and tables showing the mapping between A-spec and B-spec requirements in Section 6.

1.1 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 requirement has been met.

Here are some principles of requirements that guide the discussion to follow[1]:

- 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").

It is important to keep the list of requirements requested by the customer community separate from the derived requirements. It is a serious matter when a top-level requirement cannot be achieved, typically requiring specific acknowledgement (or negotiation) with the customer community. However, if a derived requirement cannot be met, it may be possible to still meet all the customer requirements, possibly by adjusting the architecture.

---

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

1.2 On Quantitative or Numerical Requirements

In many systems, the precise functions to be provided and the conditions under which they will be performed are known in advance. In such systems, it is often fairly easy to determine and specify a range of quantitative requirements that define performance, scale, and similar metrics. Once defined as top-level requirements, a process called "requirements allocation" establishes similar requirements for each subsystem, so that the overall system meets the high level requirements.

In the case of GENI, identification and determination of quantitative requirements on scale, sizing of specific subsystems, and performance is made more difficult because of the general-purpose nature of the system. Rather than being designed for a specific purpose, GENI is intended to be a flexible and programmable facility designed to meet the needs of a wide variety of users and user experiments, including experiments not yet defined at the time GENI is first designed. This problem is typical of many general-purpose facilities. And yet, decisions about quantitative requirements remain critical to the design of the system.

For this reason, we note several approaches that may be useful for determining appropriate quantitative requirements in different circumstances.

- The quantitative requirement can be specified as an absolutely defined customer/user requirement, such as "each tail circuit shall have 10Gb of bandwidth".
- The quantitative requirement can be defined as a relative value with respect to an existing external technology or scenario, such as "The number of available CPU cores in a computer cluster should be at least 10% of that available to commercial search engine providers". In this circumstance it is also appropriate to state in the requirement whether the value can be determined and set once, or whether the system needs to *track* the existing technology or scenario.
- The quantitative requirement can be expressed as a constraint or relationship on some other quantitative requirement or set of requirements also defined for the system. An example would be "each computational cluster shall have a connection to the GENI core with at least 1Gb of available bandwidth per 1000 available physical CPU cores".

2 GENI Research Requirements

2.1 Scope

This section presents *research* goals and requirements related to GENI. These are descriptions of the research that GENI is intended to facilitate. (Later sections of the document present GENI *facility* requirements that derive from these research requirements.)

The requirements in this section are all captured directly from sections 2 – 4 of the GENI Research Plan.

There are two general categories of research requirements. Section 2.3, High Level Objectives, captures the major research challenges that motivate GENI. These can be thought of as the "broader impacts" of the GENI-supported research program. Section 2.4, Concrete Research Requirements, contains requirements about specific features and characteristics GENI is expected to have and research it is expected to support.

All requirements are annotated with the section of the GRP from which they are derived.

2.2 Related Documents

GRP, "GENI Research Plan" version 4.4a, March 6, 2007

2.3 High Level Objectives

2.3.1 Overall

[R1000] To realize its potential, a GENI shall enable and foster:

- A world where mobility and universal connectivity is the norm, in which any piece of information is available anytime, anywhere.
- A world where more and more of the world's information is available online-a world that meets commercial concerns, provides utility to users, and makes new activities possible. A world where we can all search, store, retrieve, explore, enlighten and entertain ourselves.
- A world that is made smarter-safer, more efficient, healthier, more satisfactory-by the effective use of sensors and controllers.
- A world where we have a balanced realization of important social concerns such as privacy, accountability, freedom of action and a predictable shared civil space.
- A world where "computing" and "networking" is no longer something we "do", but a natural part of our everyday world. We no longer use the Internet to go to cyber-space. It has come to us: a world where these tools are so integrated into our world that they become invisible. (2.0)

2.3.2 Security & Robustness

[R1010] GENI shall support architectural research to respond to the following design challenges posed by security & robustness needs:

- Any set of "well-behaved" hosts should be able to communicate among themselves as they desire, with high reliability and predictability, and malicious or corrupted nodes should not be able to disrupt this communication. Users should expect a level of availability that matches or exceeds the telephone system of today.
- Security and robustness should be extended across layers, because security and reliability to an end user depends on the robustness of both the network layer and the distributed applications.
- There should be a reasoned balance between identity for accountability and deterrence and privacy and freedom from unjustified observation and tracking. (2.1.1)

2.3.3 New Underlying Technologies

2.3.3.1 Wireless Technologies

[R1020] GENI shall support architectural research to respond to the following design challenges posed by new wireless technologies:

- A Future Internet must support node mobility as a first-level objective. Nodes must be able to change their attachment point to the Internet.
- A Future Internet must provide adequate means for an application to discover characteristics of varying wireless links and adapt to them.
- A Future Internet (or a service running on that Internet) must facilitate the process by which nodes that are in physical proximity discover each other.
- Wireless technologies must be developed to work well in an Internet context, with robust security, resource control, and interaction with the wired world. (2.1.2)

2.3.3.2 Optical Technologies

[R1030] GENI shall support architectural research to respond to the following design challenges posed by new optical technologies:

- A Future Internet must be designed to enable users to leverage these new capabilities of the underlying optical transport, including better reliability through cross-layer diagnostics, better predictability at lower cost through cross-layer traffic engineering, and much higher performance to the desktop.
- A Future Internet must allow for dynamically reconfigurable optical nodes that enable the electronics layer to dynamically access the full fiber bandwidth.

- A Future Internet must include control and management software that allow a network of dynamically reconfigurable nodes to operate as a stable networking layer. (2.1.2)

2.3.3.3 Computing Technologies

[R1040] GENI shall support architectural research to respond to the following design challenges posed by new computing technology:

- A Future Internet must take account of the specialized networks that will support future computing devices, which will imply such architectural requirements as intermittent connectivity, data-driven communication, support of location-aware applications, and application-tuned performance.
- It should be possible to extend a given sensor application across the core of the Internet, to bridge two parts of a sensor net that are part of a common sensing application but partitioned at the level of the sensor net. In the limit, a future Internet should support sensing at a global level. (2.1.3)

2.3.4 New Applications and Systems

[R1050] GENI shall permit architectural research to respond to the following design challenges posed by new applications and systems:

- A Future Internet needs to develop and validate a new set of abstractions for managing the complexity of distributed services that can scale across the planet and down to the smallest device, in a robust, secure, and flexible fashion. This must include an architecture or framework that captures and expresses an "information-centric" view of what users do.
- A Future Internet must identify specific monitoring and control information that should be revealed to the application designer, and include the specification and interfaces to these features. For example, the Future Internet might reveal some suitable measure of expected throughput and latency between specified points.
- A Future Internet should include a coherent design for the various name-spaces in which people are named. This design should be derived from a socio-technical analysis of different design options and their implications. There must be a justification of what sort of identification is needed at different levels, from the packet to the application. (2.1.4)

2.3.5 Network Management

[R1060] GENI shall permit architectural research to respond to the following design challenges posed by network management needs:

- An operator of a network region should be able to describe and configure his region using high-level declarations of policy, and automatic tools should configure the individual devices to conform.

- A user detecting a problem should have a tool that diagnoses the problem, gives feedback to the user in meaningful terms, and reports this error to the responsible party, across the network as necessary.
- All devices on a Future Internet should have a way to report failures. (2.1.5)

[R1070] GENI shall be designed to help create real networks that better match desirable theoretical models. Examples applications are guaranteeing message delivery within a finite time bound and failure detection. (3.1.7)

2.3.6 Economics

[R1080] GENI shall permit architectural research to respond to the following design challenges posed by economic needs:

- Routing protocols must be redesigned to deal with the range of business policies that ISPs want to express. Issues to be considered include signaling the direction of value flow, provisioning and accounting for higher-level services, dynamic pricing, explicit distance-sensitive pricing, and alternatives to the simple interconnection models of peering and transit.
- A Future Internet must provide a means to link the long-term resource provisioning problems at one level to the short-term resource utilization decisions (e.g. routing) at higher levels. (2.1.6)

2.3.7 Service in Times of Disaster

[R1090] GENI shall permit architectural research to respond to the following design challenges posed by robustness needs during times of disaster:

- A Future Internet should be able to allocate its resources to critical tasks while it is under attack and some of its resources have failed. (For example, it should support some analog of priority telephone access that is provided today.)
- Users should be able to obtain information of known authority in a timely way during times of crisis. The network (and its associated applications) should limit opportunities for flooding, fraudulent and counterfeit mis-information, and denial of service.
- Users should be able to obtain critical information based on their location, and request assistance based on their location.
- Portions of the Internet should be usable by the attached computers even if they are disconnected from the rest of the network. (2.2.1)

[R1100] GENI should permit architectural research to respond to the following design challenges posed by management needs during times of disaster:

- It should be possible for parts of the network to operate while partitioned from the rest of the network. This raises issues of address management, routing, and name resolution.

- It should be possible for a device (or a region of the network) that has crashed and lost all of its dynamic state to rejoin the network in a way that is both secure and as automatic as possible. (2.2.1)

2.3.8 Expansion to the Developing World

[R1110] GENI shall permit architectural research to respond to the following design challenges posed by needs of the developing world:

- autonomic management of large-scale distributed infrastructures,
- autonomic management of large mobile ad-hoc networks,
- provisioning and update of large scale distributed storage systems,
- large scale content distribution systems, and
- sensing and control of routing and transport over large, heterogeneous networks with high latency and link loss. (2.2.5)

2.4 Concrete Research Requirements

This section presents concrete research requirements – the specific research areas, capabilities, and activities that GENI is expected to support, in furtherance of the broad objectives described in Section **Error! Reference source not found.**.

This mapping from broad high-level objectives to the concrete research functions necessary to meet them is established by the research communities that are the customers of GENI, rather than by the GENI designers. This mapping is not derivable through engineering practice but is heavily dependent on the judgment and insight of the customer research community.

2.4.1 General

2.4.1.1 Use

[R1120] GENI shall provide the substrate upon which new architectural proposals, and new features and protocols, can be experimentally tested and evaluated. (3.0)

[R1130] GENI shall be used to test the concept of virtualization. (3.2.1)

[R1140] GENI shall permit conceptualization and demonstration of one or more proposals for an Internet for tomorrow. (3.1)

[R1150] GENI shall use virtualization as the basis for permitting experimentation with new network architectures. (3.1.1)

2.4.1.2 Access

[R1160] Access to GENI shall be ubiquitous and affordable. (2.2.2)

[R1170] GENI shall be accessible by real application developers and users. (2.2.3)

[R1180] GENI shall have points of presence at multiple geographical locations and within multiple ISPs. (3.1.2)

[R1190] GENI shall have nodes collocated with the Internet backbone routers. (3.1.2)

[R1200] GENI shall enable experiments that capture the idea that different parts of the Internet are run by different organizations. (3.2.3)

[R1210] GENI shall span the world, different administrative domains and cultures. (3.1.2)

[R1220] GENI shall include real user traffic. (3.2.6)

[R1230] GENI shall have topological scale and scope at least of a reasonably-sized ISP. (3.2.6)

2.4.1.3 Measurement Functions

[R1240] GENI shall collect short-term numerical performance measures such as packet delay, packet loss rate and user throughput for specified applications and mobility patterns. (3.1.5)

[R1250] GENI shall include support for measurement as a first-class mechanism because of the importance of measurement in understanding and operating the network. (2.3.2)

[R1260] GENI shall provide rich capture and logging of data from experiments. (2.3.2)

[R1270] GENI shall collect long-term service quality measures such as the percentage of dropped connections and level of availability. (3.1.5)

[R1280] GENI shall permit deployment of a hierarchy of carefully-situated monitoring infrastructure nodes. (3.4.1)

[R1290] GENI shall permit large-scale evaluation and demonstration of bounded-delay interactions, including support for critical applications, such as low-latency and error-free delivery of high-quality imagery, and low-latency control of a remote robotic device. (3.1.6)

[R1300] GENI shall support experiments which have high requirements for reliability, resilience and availability, such as remote surgery. (3.1.6)

2.4.1.4 Other General Features

[R1310] GENI shall support experiments that require reliable delivery of high-bandwidth data with controlled latency. (2.2.2)

[R1320] GENI shall develop policies and mechanisms to balance rights of users (e.g., to privacy) against rights of experimenters. (2.3.2)

[R1330] GENI shall include:

- routers for which we can accurately set the size of the buffer, and measure the occupancy in real time,
- a network built from these routers (natively over links, not as an overlay, as buffer occupancy depends critically on link delays), and

- real user traffic that will allow the hypotheses to be tested with lots of users, and lots of applications, over long periods of time. (3.2.2)

[R1340] GENI shall have edge diversity, in terms of:

- participating computing devices,
- their upload/download performance, and
- their physical connectivity. (3.1.2)

[R1350] GENI shall permit experimentation with network failures and congestion. (3.1.5)

[R1360] GENI shall permit seamless operation across heterogeneous link layer technologies, including wireless, and permit mobility experiments with new link layer protocols. (3.1.5)

[R1370] GENI shall provide a virtualizable routing platform for mobility experiments with protocols. (3.1.5)

[R1380] GENI shall permit large-scale end-to-end experiments with transport protocols that support mobility. (3.1.5)

[R1390] GENI shall include allocation and scheduling mechanisms with schemes that provide for backup capacity over disjoint routes, rapid failover from one path to another (or perhaps even simultaneous transmission along these disjoint paths). (3.1.6)

[R1400] GENI shall have the ability to dedicate fixed capacity links to a particular experiments. (3.2.6)

2.4.2 Research in Naming, Routing, & Protocols

[R1410] GENI shall permit experimentation with alternatives to today's packet addressing and forwarding. (3.2.2)

[R1420] GENI shall permit experiments to investigate the following questions regarding naming and addressing: What is the appropriate division of functionality between naming and network addressing? What, if any, role should out-of-band signaling play in a future Internet? Should network addressing be eroded to the point where a naming/signaling system, rather than a global Internet address in every packet, is the unifying characteristic of the Internet? (3.2.2)

[R1430] GENI shall support experimentation with alternative naming approaches. (3.2.3)

[R1440] GENI shall enable experimentation with anycast services and protocols. (3.2.2)

[R1450] GENI shall enable experimentation with BGP alternatives that provide better convergence after equipment failures. (3.2.3)

[R1460] GENI shall enable experiments that measure routing-protocol convergence delay and the effects on end-to-end performance when topology changes occur. (3.2.3)

[R1470] GENI shall enable experiments that involve "injecting" link failures under today's Internet routing architecture and under the new design. (3.2.3)

[R1480] GENI shall accurately capture the overheads and delays on the routing software during intentional failures. (3.2.3)

[R1490] GENI shall support experiments with the use of tunneling as a first-class component of the architecture, leading to solutions for route selection, traffic engineering, mobility, diffusion routing, and fast failure recovery. (3.2.3)

[R1500] GENI shall use virtualization to provide a high-degree of route diversity. (3.2.3)

[R1510] GENI shall support experimentation with route computation services. In particular, some experiments may interact with operational ISPs using GENI-Internet interconnection. (3.2.3)

[R1520] GENI shall use virtual machines within routers to enable concurrent competing architectures that could provide the much needed independence of failures to build a robust routing infrastructure. (3.1.7)

[R1530] GENI shall support evaluation of proposals to make network routing more secure. (3.2.3)

[R1540] GENI shall support exploration and understanding of the fundamental trade-off that exist between distributed and centralized solutions for routing systems. (3.2.3)

[R1550] GENI shall permit research in relayed communication, such as Delay Tolerant Networking. Research challenges include the design of a naming protocol and a routing protocol. (3.1.4)

[R1560] GENI shall provide a real and controlled testbeds that allows researchers to study what the DTN network stack should look like all the way from the link layer to the application layer, and comparing this with the traditional IP-based network stack. (3.1.7)

2.4.3 Research in Security & Robustness

[R1570] GENI shall support experiments with security mechanisms permitting assessment of performance under load, usability, and resilience to known and new attacks. (3.2.5)

[R1580] GENI shall provide a platform where a new secure networking architecture, providing strong assurance of communication availability even under attack, can be demonstrated to work in practice on a national scale network connecting millions of users, at Internet speeds and reasonable cost. (3.1.7)

[R1590] GENI shall permit experimentation with architectures that provide trust-modulated transparency: trusting nodes should be able to communicate at will, as in the original conception of the Internet, but nodes should be protected from nodes they do not want to communicate with. (3.2.2)

[R1600] GENI shall permit 'white hat' attacks, for example allowing attackers, such as researchers from the security research community, to attempt to interfere with the operation of an experimental system, while it attempts to serve a developing user community? (3.2.5)

[R1610] GENI shall permit experimentation with cryptographically verifiable addresses, for example by supporting host network interfaces that have some functions under control of entities other than the host operating system. (3.2.5)

[R1620] GENI shall support network attacks to test security mechanisms in a manner that protects portions and users of the system that are not participating. (3.2.5)

[R1630] GENI systems experiments, data gathering and monitoring capabilities, and mechanisms for resource allocation shall be designed so that attacks against one experimental system do not interfere with other experiments. (3.2.5)

[R1640] GENI shall permit experimentation in spam-resistant email that allows interoperation with Internet users. (3.2.5)

[R1650] GENI systems shall be capable of being left open to attack in order to test robustness. (3.2.5)

[R1660] GENI shall support designing and testing worm detection and mitigation methods. (3.2.5)

[R1670] GENI shall be capable of implementing a trusted system using attestation as a lower layer underneath GENI hosts to provide compartmentalization between different network services, architectures, and applications implemented within the network, as well as to provide a secure, tamper- and masquerade-proof identification of these properties. (3.2.5)

[R1680] GENI shall permit deployment and validation practical approaches for Byzantine Fault Tolerance (BFT) through the use of a massive-scale distributed edge testbed comparable to large content distribution networks of today. (3.1.7)

[R1690] GENI shall permit large-scale experimentation and evaluation of the BFT approach to service robustness through the inclusion of large edge clusters that can host real Web services like content distribution or multi-tier Web services. (3.1.7)

[R1700] GENI's shall include support for delay-tolerant edge network testbeds to enable researchers to experiment with developing realistic failure models to replication approaches. (3.1.7)

[R1710] GENI shall provide a controlled network testbed to deploy and validate routing architectures as well as distributed systems based on the BAR model. (3.1.7)

[R1720] GENI shall contribute to understanding how to develop an infrastructure for secure identities that can be used for a broad range of network as well as end-system services. (3.1.7)

[R1730] GENI shall enable researchers to experiment with new techniques for building dependable distributed systems. (3.1.7)

2.4.4 Research in Network Management

[R1740] GENI shall enable network management experiments by:

- including real user traffic,
- exposing management functionality of network resources,

- supporting long-running experiments to gain operational experience, and
- permitting construction of multiple interacting domains. (3.2.6)

[R1750] GENI shall include a framework for testing network management ideas and integrating them with network experiments. (3.2.6)

[R1760] The GENI infrastructure shall be able to steer user traffic through certain virtual experimental networks, detect when performance through those networks is sub-par, and subsequently steer traffic away from those failing experiments. (3.2.6)

[R1770] GENI shall support experiments into a graceful maintenance architecture, which would make step-by-step coordinated changes in the configuration of network equipment at multiple protocol stack levels to prepare for removing equipment from the network. (3.2.3)

2.4.5 Integration of Wireless Technologies

[R1780] GENI shall support development of novel wireless applications and deployment them at scale to understand what services and systems components would be required in a future Internet. (3.3.1)

[R1790] GENI shall support design, prototyping, and evaluation of novel wireless architectural components and examination their performance, flexibility, and manageability. (3.3.1)

[R1800] GENI wireless networks shall include location determination services based on signal triangulation or other methods. (3.4.4)

[R1810] GENI shall support deployment of sensors at significant physical scale. (2.2.3)

[R1820] GENI shall permit multiple independent deployments of sensors to interact in a single experimental infrastructure. (2.2.3)

[R1830] GENI shall enable the deployment of new communication and in-network processing primitives that are unique to the sensor networking research area. (2.2.3)

[R1840] GENI shall provide the infrastructure for querying and fusing data across multiple (possibly overlapping) sensor networks in different scientific and administrative domains. (3.1.3)

[R1850] GENI shall include infrastructure for global sensing research, which involves integration and manipulation of data across the world, not in a locale. (3.1.3)

[R1860] GENI shall include wireless mesh networks built out of nodes with radios (e.g., 802.11) that communicate with each other to form an ad hoc "self-configuring" network without much manual involvement. (3.3.1.1)

[R1870] GENI shall include a cognitive radio network to explore a number of architectural issues towards understanding how this technology can be integrated into a future Internet architecture- these include control and management protocols, support for cooperative communication, dynamic spectrum coordination, flexible MAC layer protocols, ad hoc group formation and cross-layer adaptation. (3.3.1.2)

[R1880] GENI wireless networks shall include sufficient storage and programmability to permit experimentation in cache-and-forward architectures. (3.3.1.3)

[R1890] GENI shall include vehicular networks including advanced instrumentation of the car, so that it can observe its environment, and highly reliable communication among cars so they can plan common actions and maintain safe trajectories. (2.2.4)

[R1900] GENI shall include a vehicular network for research in radio and MAC layer performance, channel management, coexistence of critical and infotainment traffic, protocol design and testing and interfacing to the Internet and legacy wireless technologies. (3.3.1.4)

[R1910] GENI shall support long-running location services to new mobile or sensor applications. (3.4.4)

2.4.6 Integration of Optical Technologies

[R1920] GENI shall include switched optical components in GENI to allow researchers to experiment with algorithms for rapid reconfiguration of aggregates. (3.2.1)

[R1930] GENI shall permit on-demand allocation of additional capacity based on router and/or experiment demand. (3.3.2)

[R1940] GENI shall make bufferless all-optical switches available for experiments in congestion control. (3.2.4)

[R1950] GENI shall include a reconfigurable optical layer permitting TDM and WDM circuits and entire fibers to be allocated under control of an experiment. (3.3.2)

2.4.7 Development of Distributed Applications

[R1960] GENI shall permit development of a coherent architecture at the level of information management and dissemination, and allow a range of transport mechanisms to support it. (3.1.2)

[R1970] GENI shall permit research in performing content distribution at scale, placing new functionality within the network to exploit localized knowledge and optimization opportunities. (3.1.2)

[R1980] GENI shall contain distributed infrastructure for content distribution research. (3.1.2)

[R1990] GENI nodes shall have significant computational and storage capabilities to permit experimental validation of new dependability techniques that rely on large-scale replication of code and data. (3.1.7)

[R2000] GENI shall support development of distributed stream query engine architectures that can run at Internet scales, with enormous aggregate volumes of data being generated across thousands or millions of sites. (3.4.1)

[R2010] GENI shall support collective monitoring of the network from its constituent end-hosts. (3.4.1)

[R2020] GENI will need to include components which may be used to model data centers. (3.4.2)

[R2030] GENI shall include a prototyping infrastructure to allow for the empirical analysis of a variety of alternative data center architectures, from as fine a grain as the "many-core" chip level up through the cluster level, to as coarse a grain as the federation of geographically-distributed data centers connected by long-haul links. (3.4.2)

3 GENI Top Level Requirements

3.1 Scope

This section presents GENI's identified "A-Spec", or customer requirements.

The requirements in this section are all captured directly from Section 5 the GENI Research Plan. Unlike the previous section, the requirements in this section are more specific and, in many cases, quantifiable. Where the previous section primarily described research that GENI needed to support, this section speaks more specifically to what the infrastructure should look like. Section 5 of the GRP gives the justification of how the research described in Section 2 (above) maps to the infrastructure requirements into this section.

The organization of the requirements below is the same as Section 5 of the GRP to facilitate review. The organization of the requirements is expected to change.

In future versions of this document, further requirements may be derived from additional sources. If developed as a separate document, a GENI Operations Plan would be one likely source for such requirements.

3.2 Related Documents

GRP, "GENI Research Plan" version 4.3, January 2, 2007

3.3 Requirements

3.3.1 Description

GENI is an experimental facility being planned by the NSF, in collaboration with the research community. It's goal is to enable the research community to invent and demonstrate a global communications network and related services that will be qualitatively better than today's Internet.

To support multiple, long running experiments, the designers have based GENI on the concept of **slices.** The concept of slices is that the resources of GENI can be divided up among many different researchers in such a way that each can run his own experiment. One approach to slices is **virtualization,** an idea that has been a part of CS research for decades. Virtualization takes a physical resource (e.g. a processor) and creates the illusion that it is multiple processors, identical to the original except that each runs slower. We have experience today in how to virtualize a processor (indeed, there are virtualized routers available as products today), and in how to virtualize a fiber link. The concepts necessary to virtualize or otherwise slice a wireless system are less mature, but there are good proposals. So GENI puts these ideas together to build a new class of facility, a virtualized infrastructure for network. [[[may be more appropriate to say just refer to experiments rather than slices for the A-spec]]]

The rationale for these requirements is contained in the GENI Research Plan, version 4.3, referred to below as GRP.

3.3.2 System Requirements

3.3.2.1 Multiple simultaneous experiments (RP 5.1.1)

[R2000] Number of Concurrent Experiments. The GENI facility shall support at least 1000, continuous, concurrent experiments. Each experiment may run for weeks or months. (GRP 5.1.1)

[R2010] Inter-slice connectivity. To permit high-performance, multi-slice experiments (e.g. multi-domain routing experiments), GENI shall include cross-slice connectivity at throughputs consistent with the backbone capacity of the GENI facility. (GRP 5.1.1)

[R2020] Controlled Isolation. The GENI facility shall support strong isolation between slices so that experiments do not interfere with each other. GENI's isolation mechanisms must be sufficiently robust to make reproducible experiments possible. (GRP 5.1.1)

[R2030] Controlled Slice/Internet Interconnection. The GENI facility 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. (GRP 5.1.1)

[R2040] Resources-Received Feedback. The GENI facility shall provide feedback about what resources a slice actually receives to enable researchers to evaluate the validity of their results. (GRP 5.1.1)

3.3.2.2 Generality (RP 5.1.2)

[R2050] Packet Formats. The GENI facility shall permit experimentation with network packet formats that materially differ from those of the Internet. (GRP 5.1.2)

[R2060] Network-level Node Functions. GENI facility nodes – in-network packet processing elements – must be capable of examining, storing, and resending data. (GRP 5.1.2)

[R2070] Link-layer Node Functions. Nodes must be capable of reassembling and reformatting low-level elements such as packets. (GRP 5.1.2)

[R2080] Multiplexing-free Paths. The GENI facility shall have the capability to provide end-to-end data circuits which are devoid of any framing or multiplexing along the path. (GRP 5.1.2)

[R2090] Reconfigurable Optical Circuits. The GENI facility shall include rapidly reconfigurable optical circuits in the core. (GRP 5.1.2)

[R2100] Current Representative Technologies. The GENI facility shall include a wide class of representative networking technologies, spanning the spectrum of wired and wireless technologies available today. (GRP 5.1.2)

[[[Proposed additional text:

Technologies should be selected based on their ability expand the diversity of the facility in the following respects:

- transmission speed and range of transmission speed
- shared vs isolated medium? (point to point, single channel multicast-broadcast, fully shared)
- dynamics and switchability - how fast can a path be created/destroyed/rerouted? (months? minutes? milliseconds? per packet?)
- variability - is a channel fixed capacity, or subject to variability / noise / interference?
- cost per bit (or cost per some other useful metric)? (In the case of multicast, perhaps cost per bit per receiver)
- functional capabilities:
    - cost/feasibility of buffering?
    - cost/feasibility of data examination (ie, looking at headers)
    - cost/feasibility of data insertion/deletion (adding or subtracting data)
    - cost/feasibility of data manipulation (changing data that is there) **]]]**

[R2110] Future Representative Technologies. The GENI facility shall include explicitly defined procedures and system interfaces to facilitate incorporation of additional technologies, including those that *do not exist* today. (GRP 5.1.2)

[R2120] Experimenter Access to Technology Features. The GENI facility shall permit experimenter access to important technology-specific features. (GRP 5.1.2)

3.3.2.3 Support for real applications (RP 5.1.3)

3.3.2.4 Support for real users (RP 5.1.4)

[R2130] Attractive to Users. The system must provide low barrier to entry and recovery protection for people using experimental services. [[[design goal: Some experiments may be designed so that users can exploit their services without taking explicit action. For this class of experiment, it would be beneficial if the GENI infrastructure could be involved in the mechanisms by which users join in the experiments, so that the users can be switched out again if the experiment crashes. (GRP 5.1.4)]]]

[R2140] User Reach. The GENI facility shall include apparatus that is located on the campus of research facilities, with connectivity all the way to the end-node computers used by the target users. (GRP 5.1.4)

[R2150] Tools for End-system Modification. The GENI facility shall include tools which enable experimenters to install new software in popular operating systems, such as Linux or Windows. (GRP 5.1.4)

[R2160] Diversity of Path Characteristics. The GENI facility must include paths which have wide characteristic variation to permit experimentation in user-selected routing. (GRP 5.1.4)

[R2170] Internet-Connected Users. The GENI facility shall permit access by users who connect via the Internet. (GRP 5.1.4)

[R2180] Directly-Connected Users. GENI shall include connectivity to an adequate pool of potential users that have their end-node computers directly connected to, and a part of, the GENI infrastructure. (GRP 5.1.4)

3.3.2.5 Fidelity (RP 5.1.5)

[R2190] Path Diversity. The GENI facility shall be capable of generating physically node-disjoint and link-disjoint topologies. (GRP 5.1.5)

[R2200] Underlying Fiber Paths. GENI facility backbone nodes shall be placed in locations with multiple fiber-optic connections. (GRP 5.1.5)

[R2210] Topology-Related Delays. The GENI facility should have end-to-end delays no more than 2x (TBR) the geographic distance between end-points. (GRP 5.1.5)

[R2220] Major Interconnection Points – ISP. The GENI facility shall have backbone elements existing interconnection points where other ISPs have their backbone sites. (GRP 5.1.5)

[R2230] Major Interconnection Points – Exchange Points. The GENI facility shall have backbone elements at major exchange points for acquiring connectivity to the legacy Internet. (GRP 5.1.5)

[R2240] Major Interconnection Points – GigaPops. The GENI facility shall have backbone elements at major aggregation points (such as the GigaPops) to facilitate efficient, low-cost connectivity to edge sites, such as university campuses. (GRP 5.1.5)

[R2250] Physical Distribution. GENI shall provide a realistic platform to test systems that range from centralized, to distributed on a regional, campus or end-node basis. (GRP 5.1.5)

[R2260] Scale. The GENI facility shall permit experimental systems that are at least in the same league as what the commercial world has today, perhaps to within an order of magnitude of currently deployed distributed systems. (GRP 5.1.5)

[[[The Akamai system, for example is over 20,000 servers on 1,000 networks (many with multiple points of operation), currently in 71 countries. – This requirement *could* read "The GENI facility shall permit experimental systems with at least 2,000 servers on as many as 100 networks, in up to 7 countries." That probably 3 requirements, actually.]]]

[R2270] Virtual Failure On Command. The GENI facility shall support intentional failure on command of any virtualized components. Failures may be single or en-masse to support simulations of massive infrastructure outages. (GRP 5.1.5)

[R2280] Failure Notification. Experimenters must be made aware of the impact real-component failure will have on virtualized components. (GRP 5.1.5)

[R2290] Virtualization Realism. Unrealistic behavior in GENI virtualized systems, such as timing jitter, shall be identified, minimized, and specifically documented. (GRP 5.1.5)

3.3.2.6 Support for all aspects of a new network architecture (RP 5.1.6)

[R2300] Virtualized Management Interfaces. Each virtual device created as a slice of a physical device shall present a full management interface within the slice. (GRP 5.1.6) For example:

- o If the device has operating modes or states, these shall be separately settable for each slice.
- o It shall be possible to bring up and shut down a slice of a component.
- o If a device has a physical management interface that is used by system operators, this too shall be virtualized to the extent possible.

[R2310] Secure Management Technologies. Specialized security technology, such as hardware-specific un-forgeable identity tags, key generators, or physical hardware interfaces, shall be used for secure management. (GRP 5.1.6)

[R2320] Security Fidelity. The GENI facility must be stable and secure to an adequate degree, so that that experiments that claim enhanced security or availability can actually demonstrate these virtues. (GRP 5.1.6)

[R2330] Isolation for Security. The GENI facility shall contain mechanisms for isolation among slices must to ensure system attacks in one portion of the system cannot "escape" and attack other experiments. (GRP 5.1.6)

[[[Feature Extensibility. As we envision the facilities that should be included in GENI, we must remember that in 10 years, there may be features that will be commonplace then, but are not yet realized in any effective way. (GRP 5.1.6)]]]

3.3.2.7 Support for experimenters (RP 5.1.7)

Isolation from the Internet. The GENI facility shall contain network behavior to prevent attack today's Internet. (GRP 5.1.7)

[R2340] Ease of Use. The GENI facility shall provide a rich set of tools for configuring, monitoring, and debugging experiments, a rich set of common utilities to be used by experimenters, and predictable and repeatable behavior for experiments running on the system. (GRP 5.1.7)

[R2350] Power Users. The GENI facility shall provide access to the full set of capabilities of the system for "power users." (GRP 5.1.7)

[R2360] Observability. The GENI facility shall support on-line measurements in support of measurement-based quantitative research. [[[what measurements?]]] (GRP 5.1.7)

[R2370] Traceback. The GENI facility must permit network activity to be traced back to the responsible experiment (& experimenter). (GRP 5.1.7)

[R2380] Fail-safe. Should GENI enter a period where activities of some components cannot be adequately monitored or controlled, GENI shall be capable of restricting 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). (GRP 5.1.7)

3.3.2.8 Federation & Sustainability (RP 5.1.8)

[R2390] **Operational Lifetime.** The GENI facility shall be designed for a 15-20 year lifetime. (GRP 5.1.8)

[R2400] **Facility Opt-in.** The GENI facility shall permit connection of purpose-built networks (including dedicated transmission pipes and sensor networks) into GENI and running their applications and services in a slice of GENI. (GRP 5.1.8)

[R2410] **Upgrade In-Place.** The GENI facility shall permit insertion new technology to GENI while it is in operation. (GRP 5.1.8)

4 GENI Design

This section has been excised from the GENI Facility Design [[[ref]]], building on the research plan and facility requirements presented elsewhere [GDD-06-28], describes the design of the GENI facility. It introduces the key hardware and software elements that make up GENI, along with the system architecture that combines these elements into a coherent facility. This design meets the requirements in the previous section and provides a basis for deriving additional top-level requirements which will follow.

The overall facility architecture can be divided into three levels, as illustrated in Figure 1. At the bottom level, GENI provides a set of physical facilities (e.g., routers, processors, links, wireless devices), which we refer to as the *physical substrate*. The design of this substrate is concerned with ensuring that physical resources, layout, and interconnection topology are sufficient to support GENI's research objectives. As technology evolves, it is expected that the set of network technologies included in the substrate will evolve as well. A key goal of the overall GENI architecture is to enable and accommodate the evolution of these technologies throughout the lifetime of the facility.

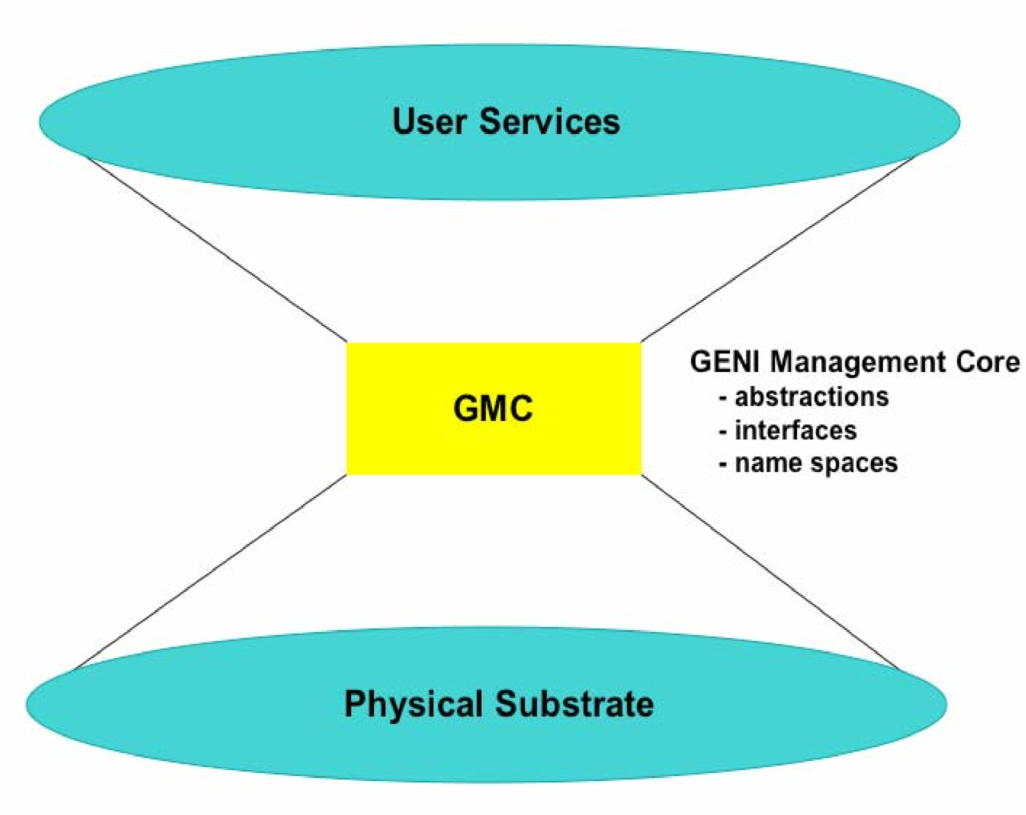

Figure 1: Overview of the GENI Architecture

At the top level, GENI's *user services* provide a rich array of user-visible support services intended to make the facility accessible and effective in meeting its research goals. As researcher requirements change and GENI's sophistication advances, it is expected that the set of available user services will change and advance as well. A key goal of the overall GENI architecture is to enable and foster the evolution of these services throughout the lifetime of the facility.

Sitting between the physical substrate and the user services is the *GENI Management Core*, or GMC. The purpose of the GMC is to define a stable, predictable, long-lived framework—a set of abstractions, interfaces, and name spaces—to bind together the GENI architecture. Because GENI's physical substrate and user services will develop and evolve rapidly as the facility is constructed and used, the GMC is designed to provide a narrowly defined set of mechanisms that both *support and foster* this development and *isolate* developmental change in one part of the system from that in other parts, so that independent progress may be made.

By analogy with today's Internet architecture, the GENI architecture conforms with the hour-glass model: the GMC corresponds to the IP layer of the Internet stack with its attendant addressing routing and service model (i.e., it defines the "narrow waist" of the GENI hourglass); the set of high-level user services corresponds to the additional functionality needed to make the Internet a complete system (e.g., HTTP, WWW, Skype); and the GENI substrate corresponds to the collection of computing and networking devices that make up the physical Internet.

For GENI to function, the GMC must be reduced to a specific implementation. We refer to this implementation as the *GENI Management Core Implementation*, or GMCI. The GMCI provides two key elements of the overall GENI system. These are the small set of core services that are necessary for the system to operate, and an underlying messaging and remote operation invocation framework needed for elements of the GENI system to communicate with each other. An important aspect of the GMCI is its ability to allow different implementations and versions of non-core user services to be easily added to the system as appropriate.

5 Requirements Derived based on the Architecture

5.1 Scope

The requirements in this section are derived from the top level requirements given the associated architecture described in Section 4.

5.2 Related Documents

GDD-07-XX GENI Facility Design document (FDD) Version 0.7

5.3 Requirements

5.3.1 Description

<blah blah>

5.3.2 System Requirements

The requirements in this section are organized roughly in the order in which they were derived from within the Facility Design Document for ease of review. The organization may change in the future.

5.3.2.1 System Architecture (FDD 1.1)

[B1000] **User Services Subsystem.** The GENI system shall include a subsystem responsible for providing services to researchers. (FDD 1.1)

[B1010] **Evolvability**. GENI shall accommodate changes in the physical substrate and user services during development and after commencement of operations. (FDD 1.1)

[B1020] **GENI Management Core.** The GENI system shall include a framework of abstractions, interfaces, and namespaces capable of configuring the GENI physical substrate. (FDD 1.1)

[B1030] **Open Interfaces.** The GMC shall permit different implementations and versions of non-core user services to be easily added to the system as appropriate. (FDD 1.1)

5.3.2.2 Physical Substrate (FDD 1.2)

[B1040] **Physical Substrate**. The GENI system shall include a collection of building block components permitting the creation of virtual networks. (FDD 1.2)

[B1050] **Expandable Substrate.** The GENI physical substrate shall be expandable, permitting new components to be added after system construction. (FDD 1.2)

[B1060] **Computational Resources for Services.** The physical substrate shall include the computational resources necessary to build wide-area services and applications, as well as initial implementation of new network elements.(FDD 1.2)

[B1070] **Programmable Core Nodes.** The physical substrate shall include core network data processing functions for high-speed, high volume traffic flows.(FDD 1.2)

[B1080] **Programmable Edge Nodes.** The physical substrate shall include data forwarding functionality at the boundary between access networks and a high-speed backbone. (FDD 1.2)

[B1090] **Programmable Wireless Nodes**. The physical substrate shall include proxies and other forwarding functionality within a wireless network. (FDD 1.2)

[B1100] **Mobile Devices**. The physical substrate shall include mobile devices intended to run applications that give end-users access to experimental services available on the combined wired/wireless substrate. (FDD 1.2)

[B1110] **National Fiber Facility**. The physical substrate shall include light path interconnection between GENI core nodes, forming a nationwide backbone network supporting rates between 10-40Gbps. (FDD 1.2)

[B1120] **Edge-to-Core Tail Circuits**. The physical substrate shall include tail circuits to connect GENI Edge Sites to the GENI Core. [[[number? rates?]]] (FDD 1.2)

[B1130] **Core-to-Internet Tail Circuits**. The physical substrate shall include tail circuits to connect the GENI Core nodes to the Internet. [[number? Rates?]]] (FDD 1.2)

[B1140] **Internet Exchange Points**. The physical substrate shall include nodes at Internet Exchange Points. [[[number? Rates?]]] (FDD 1.2)

[B1150] **Urban 802.11-based Mesh Wireless Subnets**. The physical substrate shall include at least one 802.11-based mesh wireless subnet to provide real-world experimental support for ad-hoc and mesh network research based on an emerging generation of short-range radios. (FDD 1.2)

[B1160] **Wide-Area Suburban 3G/WiMax-based Wireless Subnets**. The physical substrate shall include at least one Wide-Area Suburban 3G/WiMax-based Wireless Subnet intended to provide open-access 3G/WiMax radios for wide area coverage, along with short-range 802.11 class radios for hotspot and hybrid service models. (FDD 1.2)

[B1170] **Cognitive Radio Subnets**. The physical substrate shall include at least one cognitive radio subnet intended to support experimental development and validation of emerging spectrum allocation, access, and negotiation models. (FDD 1.2)

[B1180] **Sensor Subnets**. The physical substrate shall include at least one Application-Specific Sensor Subnet capable of supporting research on both underlying protocols and specific applications of sensor networks. (FDD 1.2)

[B1190] **Emulation Grids.** The physical substrate shall include at least one Emulation Grid that allow researchers to introduce and utilize controlled traffic and network conditions within an experimental framework. (FDD 1.2)

5.3.2.3 User Services (FDD 1.3)

[B1200] **Support for Owners.** GENI User Services shall support the needs of owners, those responsible for the externally visible behavior of GENI equipment and who establish high-level utilization policies. (FDD 1.3)

[B1210] Support for Administrators. GENI User Services shall support the needs of administrators, those responsible for keeping the platform running, providing a service to researchers, and preventing malicious or otherwise damaging activity exploiting the platform. (FDD 1.3)

[B1220] Support for Developers. GENI Users Services shall support the needs of developers, those who build on the GMC functionality to implement services of general use to the GENI community. (FDD 1.3)

[B1230] Support for Researchers. GENI User Services shall support the needs of researchers, those employing GENI in their work, for running experiments, deploying experimental services, measuring aspects of the platform, and so on.(FDD 1.3)

[B1240] Support for End Users. GENI Users Services shall support the needs of End users not affiliated with GENI, but who access services provided by research projects that run over GENI. (FDD 1.3)

[B1250] Support for Third Parties. GENI User Services shall support the needs of Third parties that may be confused or concerned about the effect that GENI-hosted experiments and services are having on their own enterprises. (FDD 1.3)

[B1260] Resource Allocation & Usage Policy Mechanism. The GENI system shall allow owners to declare resource allocation and usage policies for substrate facilities under their control, and to provide mechanisms for enforcing those policies. (FDD 1.3)

[B1270] Substrate Management. The GENI system shall allow administrators to manage the GENI substrate, which includes installing new physical plant and retiring old or faulty plant, installing and updating system software, and monitoring GENI for performance, functionality, and security. (FDD 1.3)

[B1280] Decentralized Management. The GENI system shall permit more than one organization to administer disjoint collections of GENI sites, ranging from individual owners managing their own machines, to a small number of large organizations that federate at a coarse grain. (FDD 1.3)

[B1290] Experimenter Services. The GENI system shall allow researchers to create and populate experiments, allocate resources to them, and run experiment-specific software. (FDD 1.3)

[B1300] Execution Environment. The GENI system shall expose an execution environment to researcher's applications, experiments, and services. These execution environments must be flexible (i.e., support a wide-range of program behavior) and perform satisfactorily (i.e., not unduly interfere with or distort the measurements or results). (FDD 1.3)

[B1310] Low-level Information Exposure. The GENI system shall expose low-level information about the state of the GENI substrate to developers, so they can implement high-level monitoring, measurement, auditing, and resource discovery services. (FDD 1.3)

5.3.2.4 System Architecture (FDD 2.0)

[B1320] (deleted)

[B1330] Design for Asynchronous System Development. The GMC shall include a narrowly defined set of mechanisms that both support and foster physical substrate and user service development and isolate developmental change in one part of the system from that in other parts, so that independent progress may be made. (FDD 2.0)

5.3.2.5 Components (FDD 2.1.1)

[B1340] Component Managers. Each GENI component shall include an abstraction of operations available to user-level services known. (FDD 2.1.1)

[B1350] CM Interface Definitions. Interfaces for component management shall be well-defined and remotely accessible. (FDD 2.1.1)

[B1360] Resource Allocation Policy Control. GENI component resource allocation policies shall be determined by the component's owner. (FDD 2.1.1)

[B1370] Component Sharing. GENI physical substrate components shall be shared or "sliced" among multiple users. This may be done through either partitioning access in time or through virtualization. The portion of a slice that exists on a single component is known as a sliver. (FDD 2.1.1)

[B1380] Slice Isolation. Each GENI physical substrate component shall include hardware or software mechanisms that isolate shared slivers from each other. (FDD 2.1.1)

[B1390] Self-limitation of Sharing. The GMC shall permit physical substrate components to establish limits on the number of simultaneous users it can support. In particular, a component that is incapable of being shared shall be able to limit itself to supporting one active user at any given time. (FDD 2.1.1)

[B1400] Behavior Restriction. GENI shall be capable of restricting the behavior of a component-and the slivers it hosts-in the surrounding network. Restrictions shall include:

- the rate at which a component can send traffic into the network,
- the other components and slivers with which it can communicate
- the rate at which a component can initiate communications with new components
- packet-format restrictions on the ability spoof source addresses protocol-specific restrictions. (FDD 2.1.1)

[B1410] Audits of Internet-bound Traffic. Physical substrate components shall include permit audits of packets transmitted to the Internet so that responsibility for any unwanted or disruptive traffic can be traced back to the responsible sliver. (FDD 2.1.1)

[B1420] Disconnection for Misbehavior. GENI components must be capable of rapid disconnection from the network permitting problem diagnosis (including potential security compromises). (FDD 2.1.1)

[B1430] Safe State. GENI components must be able to be brought into a safe state from which problems can be diagnosed (including potential security compromises). (FDD 2.1.1)

[B1440] Virtual Socket Interfaces. GENI shall permit component processes to communicate using TCP or UDP connections or IP tunnels. (FDD 2.1.1)

[B1450] Virtual Link Interfaces. GENI shall permit component processes to communicate using virtualized link interface that exposes transmission queues and link failures. (FDD 2.1.1)

[B1460] Constant Bit Rate. GENI shall provide constant bit rate service to component processes. (FDD 2.1.1)

[B1470] Point-to-point Virtual Links. GENI shall permit component processes to communicate using point-to-point virtual links. (FDD 2.1.1)

[B1480] Virtual Radio Interface: GENI shall permit component processes to communicate using a virtual radio interface that exposes the characteristics of radio devices at the PHY/MAC layers. (FDD 2.1.1)

[B1490] Radio Interface Configuration. The virtual radio shall be configurable to select a specific physical transmission rate, adjust the transmit power within a range, set threshold values of certain radio parameters, and so on, depending on the type of radio device. (FDD 2.1.1)

[B1500] Virtual Wire Interface: GENI shall permit component processes to communicate using a virtualized wire interface that permits the dynamic creation, termination, and provisioning of end-to-end circuits. (FDD 2.1.1)

[B1510] Virtual Wire Types The following types of virtual wire interfaces shall be supported: (1) the underlying capacity corresponds to a framed TDM electrical circuit; (2) the underlying capacity corresponds to an unframed TDM electrical circuit; or (3) the underlying capacity corresponds to an unframed optical circuit (e.g., a wavelength). (FDD 2.1.1)

[B1520] Component Names. GENI components shall have unique, human-readable names. (FDD 2.1.1)

[B1530] Scalable Name Systems. GENI names systems must be scalable, using hierarchical structures where possible. (FDD 2.1.1)

5.3.2.6 Slices (FDD 2.1.2)

[B1540] Multi-point Virtual Links. GENI shall permit component processes to communicate using multi-point virtual links. (FDD 2.1.2)

[B1550] (deleted)

[B1560] Slice Names. GENI slices shall have human-readable names. (FDD 2.1.2)

[B1570] User Vetting. GENI shall vet users before assigning resources. (FDD 2.1.2)

[B1580] Naming External Components. GENI naming systems shall permit federation with non-GENI components. (FDD 2.1.2)

[B1590] User-Slice Mapping. GENI shall map slices to human users. (FDD 2.1.2)

[B1600] Resource Database. GENI shall maintain a database describing all resources found in all slices. (FDD 2.1.2)

[B1610] **Experiment Configuration**. GENI shall permit users to change the configuration within a slice, e.g., changing executable code or component configuration for different runs of an experiment. (FDD 2.1.2)

[B1620] **User Services**. GENI shall maintain services to help users to manage their slices and implement experiments. (FDD 2.1.2)

5.3.2.7 Aggregates (FDD 2.1.3) & Users (FDD 2.1.4)

[B1630] **(deleted)**

[B1640] **User Identity**. Each user shall be identified by a certificate and key pair issued by one of the GENI authorities. Rights to operate on slices are associated with user credentials. (FDD 2.1.4)

[B1650] **Credential Delegation**. The ability to credential users shall able to be delegated. (FDD 2.1.4)

5.3.2.8 Interfaces (FDD 2.2)

[B1660] **Identifier System.** GENI shall include a system of cryptographic identifiers such that an entity that possesses one is able to confirm that the it was issued by GENI and has not been forged, and to authenticate that the object claiming to correspond to the identifier is the one to which it was actually issued. (FDD 2.2.1)

[B1670] **Common Resource Specification data structure.** GENI shall export a common, extensible data structure capable of specifying all resources including processing capabilities, network interfaces, and instrumentation. (FDD 2.2.2)

[B1680] **Transferable Slice Control**. Any authorized user shall be capable of starting, stopping, or destroying a slice, even if that user is not the one that created the slice. (FDD 2.2.4)

5.3.2.9 Security (FDD 2.3)

[B1690] **Explicit Trust.** Privileges in GENI shall be managed explicitly and formally. (FDD 2.3.1)

[B1700] **Least Privilege**. Each GENI component shall be given exactly the authority it needs to perform its tasks and no more (FDD 2.3.1)

[B1710] **Revocation**. GENI shall be capable of quickly revoking and replacing keys, suspending all permissions (e.g., slices) derived from a compromised key, and resetting each node to a known secure state. (FDD 2.3.1)

[B1720] **Auditability**. GENI shall include mechanisms that identify which slice is responsible for each packet, and also be able to determine the entire chain of responsibility (from central administrator to local administrator to local user) that gave the user a specific capability that was misused. (FDD 2.3.1)

[B1730] **Autonomy**. GENI shall include the ability to federate autonomous facilities. A GENI site should be able to authenticate and authorize requests from users in other sites, support delegation of rights, and it should be able to do so without requiring centralized trust. (FDD 2.3.1)

[B1740] **Security Usability**. GENI shall include intuitive and easy interfaces for users to create roles, restrict rights, etc. (FDD 2.3.1)

[B1750] **Local Management of Security Information**. GENI security information (such as certificates delegating rights to a specific set of users) shall be managed locally as far as possible. (FDD 2.3.1)

5.3.2.10 Canonical Components (FDD 2.4.4)

[B1760] **Sliver Creation.** All GENI components shall be able to create and destroy slivers, bind a set of resources to a sliver, and reclaim those resources. (FDD 2.4.4)

[B1770] **Sliver Isolation**. All GENI components shall isolate slivers from each other, such that

- the resources consumed by one sliver do not unduly affect the performance of another sliver,
- one sliver cannot, without permission, eavesdrop on network traffic to or from another sliver,
- one sliver cannot access objects (e.g., files, ports, processes) belonging to another sliver, and
- users are allowed to install software packages in their sliver without consideration for the packages installed in other slivers running on the same component. (FDD 2.4.4)

[B1780] **Sliver Secure Login**. All GENI components shall allow users to securely log into a sliver that has been created on their behalf. (FDD 2.4.4)

[B1790] **Reboot Signal**. All GENI components shall deliver signals to slivers, including a "reboot" signal that is delivered whenever the sliver starts up. (FDD 2.4.4)

[B1800] **Sliver Interposition**. All GENI components shall grant privileged operations to select slivers, including the ability of one sliver to access private state associated with another sliver (thereby supporting sliver interposition). (FDD 2.4.4)

[B1810] **Component Rate Limiting**. All GENI components shall rate-limit the network traffic generated by a sliver, as well as by all slivers running on the component. (FDD 2.4.4)

[B1820] **Internet Filtering**. All GENI components supporting a connection to the existing Internet shall limit (filter) how a sliver interacts with (exchanges packets with) the Internet. (FDD 2.4.4)

[B1830] **Auditability**. All GENI components supporting a connection to the existing Internet shall support a mechanism to audit all packet flows transmitted by slivers to the Internet, and determine what sliver (slice) is responsible for a given packet. (FDD 2.4.4)

[B1840] **Safe State**. All GENI components shall disconnect the component from the network and bring it into a safe state, and subsequently reboot the component into a correct configuration. (FDD 2.4.4)

[B1850] **Component Power Control & Systems Errors**. All GENI components shall power the component on and off, and monitoring its hardware and other software systems for errors. (FDD 2.4.4)

[B1860] **Traffic Monitor**. GENI components shall include a Traffic Monitor responsible for authorizing, shaping, and auditing network traffic that is sent and received by a component. (FDD 2.4.4)

[B1870] **Secure Boot Monitor**. GENI components shall include a Secure Boot Monitor that interacts with the GMC to install and boot a node with the correct software. (FDD 2.4.4)

[B1880] **Hardware Monitor**. GENI components shall include a Hardware Monitor that provides a reliable method of remotely accessing and controlling a component. (FDD 2.4.4)

5.3.2.11 Edge Sites (FDD 3.1)

[B1890] **Number of GENI Edge Sites.** GENI shall be sized to include at least 200 edge sites. (FDD 3.1)

[B1900] **Dedicated Tail Circuits**. Edge sites hosting GENI components shall be connected to the GENI backbone by dedicated tail circuits. (FDD 3.1.2)

[B1910] **Internet Connectivity**. Edge sites hosting GENI components shall be connected to the Internet using commodity Internet access services. (FDD 3.1.2)

[B1920] **Site Locations**. Each GENI site shall be co-located within the campus network infrastructure of some hosting organization (e.g., a University or industrial research lab). (FDD 3.1.2)

[B1930] **Peering Co-location Model**. Co-location designs shall permit GENI sites to peer directly with GENI backbone and with the hosting organizations campus network. (FDD 3.1.2)

[B1940] **DMZ Co-location Model**. Co-location designs shall permit GENI sites outside the hosting organization's campus firewall but within the campus border router or gateways, on a boundary network between the campus and the Internet. (FDD 3.1.2)

[B1950] **Embedded Co-location Model**: Co-location designs shall permit GENI components to be embedded within the campus network itself, such within a departmental network. (FDD 3.1.2)

[B1960] **Dark Fiber Tail Circuits**. GENI shall support dark fiber tail circuits between edge site and backbone nodes. (FDD 3.1.2)

[B1970] **Shared Dark Fiber Tail Circuits**. GENI shall support shared dark fiber tail circuits, e.g., using DWDM, between edge sites and backbone nodes. (FDD 3.1.2)

[B1980] Leased PHY Tail Circuits. GENI shall support leased layer 1 service, e.g., SONET or pont-to-point Ethernet, for between edge sites and backbone nodes. (FDD 3.1.2)

[B1990] L2 Tail Circuits. GENI shall support dedicated and shared layer 2 circuits such as Ethernet, Infiniband, and Fiber Channel between edge sites and backbone nodes. (FDD 3.1.2)

[B2000] L3 Tail Circuits. GENI shall support shared best effort routed IP and dedicated service (e.g., MPLS) between edge sites and backbone nodes. (FDD 3.1.2)

[B2010] Maximum Tail Circuit Capacity. The maximum supported capacity for GENI tail circuits shall be 10Gbps. (FDD 3.1.2)

5.3.2.12 Backbone Network (FDD 3.2)

[B2020] Backbone Connectivity between GENI Edge Sites: The GENI backbone shall provide high-bandwidth, low-latency connectivity between GENI edge sites for experiments with distributed services and wireless/sensor subnets. (FDD 3.2)

[B2030] Backbone-enabled Network Experiments: The GENI backbone shall permit experimentation with novel routing, forwarding, signaling, and addressing schemes. (FDD 3.2)

[B2040] Backbone Internet Connectivity: The GENI backbone shall include direct connections to Internet Exchange Points. (FDD 3.2)

[B2050] Backbone Packet Processing. GENI backbone nodes shall provide high-speed programmable packet-processing. (FDD 3.2)

[B2060] Backbone Programmable Framers. GENI backbone nodes shall provide programmable framers and optical switches available to the experiments. (FDD 3.2)

[B2070] Packet-switched & Circuit-switched Backbone Experiments. GENI backbone nodes shall support packet-switched and circuit-switched experiments. (FDD 3.2)

[B2080] Realistic Traffic. GENI shall enable subjecting of experimental network architectures to realistic traffic and network conditions. (FDD 3.2.2)

[B2090] Backbone Path Diversity: The GENI backbone topology shall include multiple paths between sites. Link-disjoint or node-disjoint paths shall be available. (FDD 3.2.2)

[B2100] Fiber-hub Cities. GENI backbone nodes shall be placed in cities with multiple fiber-optic connections. (FDD 3.2.2)

[B2110] Backbone sites at IXPs. GENI backbone nodes shall be located at major Internet exchange points. (FDD 3.2.2)

[B2120] Maximum End-to-end Latency. The maximum one-way, end-to-end latency across the GENI backbone shall be 100 msec. (FDD 3.2.2)

[B2130] Maximum Hop-by-hop Latency. The maximum one-way latency between any two connected GENI backbone sites shall be 2x the line-of-site propagation time (i.e.,"air miles") between the sites. (FDD 3.2.2)

[B2140] **Number of Backbone Nodes**. The GENI backbone shall have 25 nodes. (FDD 3.2.2)

[B2150] **Backbone Node Fan-out**. GENI backbone nodes shall connect to up to four other nodes. (FDD 3.2.2)

[B2160] **Internet Packet Logging**. GENI backbone and edge nodes shall log data packets sent to and from the Internet including sufficient information to map the packets back to the associated slice. (FDD 3.2.3)

[B2170] **Internet Rate & Packet Control**. GENI backbone packet processors and edge nodes shall perform access control and rate limiting to control traffic sent to the Internet. (FDD 3.2.3)

[B2180] **Tunnels**. GENI shall support the use of tunnels to permit resources to interconnect to GENI over the commodity Internet. (FDD 3.2.3)

[B2190] **BGP Gateway**. GENI shall permit experimenters to inject routes into the Internet using their own IP addresses, while also protecting the ISP from malfunctioning or malicious experiments.(FDD 3.2.3)

5.3.2.13 Wireless Network (FDD 3.3)

[B2200] **Untethered Wireless Nodes**. GENI shall support tethered and untethered wireless nodes. (FDD 3.3.2)

[B2210] **Remote Management for Wireless Node**. GENI shall support wireless nodes where the component management interface resides on a another node. (FDD 3.3.3)

[B2220] **Non-GENI Wireless Devices**. GENI shall support access by wireless devices which are not GENI nodes. (FDD 3.3.3)

[B2230] **GENI Wireless "Opt-in".** GENI shall support access by a mobile device which is completely independent of GENI, but configured to "opt into" an experiment running on GENI. That is to say, the device does not run a CM or support slices that can be configured to support various experiments, but rather, an application running on the device can be configured to send or receive data via an experimental service running on one or more GENI components. (FDD 3.3.3)

[B2240] **Wireless Spectrum Management**. Each GENI Wireless Subnet shall operate a management process that provides the RF spectrum for the entire deployment including checks for compatible allocation of frequencies, power, and so on, across nodes and slices in the subnet. (FDD 3.3.4)

[B2250] **RF Activity Profiling**. Each GENI Wireless Subnet shall profile of RF activity during experimental activities, whether generated by the experiment itself or externally induced. (FDD 3.3.4)

[B2260] **802.11 Subnet**. GENI shall include an 802.11 Urban Mesh Subnet allows experimentation with wireless mesh architectures, as well as with multi-hop radio networks and mobile ad hoc networks (MANETs). It also provides access to GENI-wide experiments and services from moving vehicles and other mobile devices. (FDD 3.3.4)

[B2270] **Cellular/WiMax-WiFi Hybrid Subnet**. GENI shall include a wide-area suburban wireless network with open-access Cellular/ WiMax radios for wide-area coverage, along with short-range 802.11 radios for hotspot and hybrid service models (Figure 3.10). This wireless scenario is of particular importance for the future Internet as cellular phone and data devices are expected to migrate from vertical protocol stacks such as GSM, CDMA and UMTS towards an open Internet protocol model. (FDD 3.3.4)

[B2280] **Cognitive Radio Subnet**. GENI shall include a cognitive radio demonstrator network to perform experimentation on architectures that inherently build upon larger network services. (FDD 3.3.4)

[B2290] **Sensor Network Kit**. GENI shall include a sensor deployment kit consisting of network gateways (from sensor radios to 802.11 or cellular), sensor modules and related platform software to enable users to build additional sensor networks for specific research or application objectives. (FDD 3.3.4)

[B2300] **Local Sensor Subnets**. GENI shall include up to 50 local sensor subnets, each containing up to 100 nodes. (FDD 3.3.4)

[B2310] **Sensornet Outdoor Deployment**. GENI shall contain 2-3 metropolitan-scale outdoor sensornet deployments, each with at least two application-specific sensor subnetworks of up to 1000 sensors with 10's of network gateways. (FDD 3.3.4)

[B2320] **Sensornet Indoor Deployment**. GENI shall contain two indoor sensornets, each of 1000 nodes. (FDD 3.3.4)

[B2330] **External Sensornet Interconnect.** GENI shall develop an interface permitting existing sensornets to connect. (FDD 3.3.4)

[B2340] E**mulation Grid**s. GENI shall include networked wireless emulation grids to permit repeatable experimentation with wireless subsystems. (FDD 3.3.4)

5.3.2.14 Operator Portal (FDD 4.1)

[B2350] **Operations Staff**. GENI shall be designed to be managed by a small, full-time operations staff. (FDD 4.1)

[B2360] **Automated Problem Detection & Resolution**. GENI shall include tools which permit problem detection and resolution with minimal human intervention. (FDD 4.1)

[B2370] **Fault Management System**. GENI shall detect problems with system hardware and software, and initiate the process of repairing these faults. (FDD 4.1.1)

[B2380] **Fault Data Acquisition**. GENI components and aggregates shall provide fault data to the Fault Management System. (FDD 4.1.1)

[B2390] **Fault Management Rules**. GENI shall include an expandable set of rules used to map observed faults to possible root-causes. (FDD 4.1.1)

[B2400] **FMS-Operator Interface**. GENI shall have the ability to alert an operator of a possible problem and include the ability for the operator to inform the system that a problem has been resolved. (FDD 4.1.1)

[B2410] **Administrator Database**. GENI shall maintain a database of people and institutions associated with GENI resources, so the GENI Fault Management System can notify the appropriate people of problems at their site. (FDD 4.1.1)

[B2420] **Configuration Management System**. GENI shall include a Configuration Management System to facilitate the orderly introduction of new nodes, links, and sensors into GENI. The GENI Configuration Management System keeps track of where hardware components are located, what versions of software they're running, which patches have been installed, and whether they're in a valid configuration. (FDD 4.1.2)

[B2430] **Provisioning New Components**. The GENI Configuration Management System shall provision, configure, and validate new components. (FDD 4.1.2)

[B2440] **Accounting Management System**. GENI shall include an GENI Accounting Management System is used to map GENI accounts to real people and institutions, and to control access to GENI resources. By tracking users and accounts, the GENI Accounting Management System can ensure that only authorized users and experiments are permitted to consume GENI resources, and express policies about how many resources they may consume. If an experiment causes problems for other users, the GENI Accounting Management System can help the GENI Operations Staff determine the person responsible for the faulty experiment and their institution. (FDD 4.1.3)

[B2450] **Resource Allocation Policies**. The GENI Accounting Management System shall provide an interface through with resource allocation policies can be set. (FDD 4.1.3)

[B2460] **Performance Management System**. GENI shall provide a Performance Management System is to provide fine-grained tracking of resource usage. The data collected by this system must be accessible to the GENI Operations Staff and researchers. Additionally, the GENI Performance Management System must also provide an interface for adaptive experiments to collect performance data and change their behavior as a result. (FDD 4.1.4)

[B2470] **Slice Performance Policing**. The GENI Performance Management System shall monitor the resources consumed by each slice to determine when a slice is operating outside of its resource profile. (FDD 4.1.4)

[B2480] **Security Management System**. GENI shall provide a Security Management System is to log all security-related events in an auditable trail. Doing so helps the GENI Operations Staff determine when a user or an experiment is violating the GENI Acceptable Use Policy. Given the scale and level of programmability of GENI, and its potential for being used to attack other GENI users or other networks such as the Internet, it's critical that potentially bad behavior be detected and stopped as quickly as possible. The GENI Security Management System, working closely with the GENI Fault Management System and the GENI Performance Management System, can be used to detect security-related problems and alert GENI Operations Staff before serious damage is done. (FDD 4.1.5)

[B2490] **Account Creation Oversight**. The GENI Security Management System shall monitor the ongoing setup and configuration of accounts in the GENI Accounting Management System. Repeated attempts to create new accounts or change the privileges of existing accounts may indicate a potential security violation. When these problems are detected, the GENI Operations Staff should be alerted. (FDD 4.1.5)

[B2500] Problem Tracking. GENI shall include a mechanism for tracking GENI bugs, user questions, maintenance requests, and other problems. (FDD 4.1.6)

5.3.2.15 Researcher Portal (FDD 4.2)

[B2510] Resource Allocation Portal. GENI shall include a resource allocation portal permitting experimenters to acquire, schedule, and release system resources. (FDD 4.2)

[B2520] Slice Embedding Portal. GENI shall include tools for slice embedding, instantiating a researcher's slice on some set of components. (FDD 4.2)

[B2530] Experimenter's Workbench. GENI shall include a portal permitting researchers to create, configure, and control their experiments. (FDD 4.2)

[B2540] Fine Granularity Admission Control. GENI shall support admission control at a fine granularity; i.e., dedicated access to portions of multiple machine's CPU, memory, network, and storage resources. (FDD 4.2.1)

[B2550] Best Effort by Default. By default, GENI slices shall run in best-effort mode to ensure a high-level of resource utilization and to reduce the barrier-to-entry for running experiments on GENI. (FDD 4.2.1)

[B2560] Experiment Priority Policy. GENI shall use policy to determine the amount of resources that can be reserved by an application, the relative priority of best-effort applications, and the proportion of global resources subject to admission control. (FDD 4.2.1)

[B2570] Resource Allocation Duration. GENI's resource allocation mechanism shall support a variety of policies for the duration of resource bindings, but guaranteed resource allocations must always be time-limited. (FDD 4.2.1)

[B2580] Resource Discovery. GENI shall permit users to issue commands to find subsets of GENI components that match some criteria, such as constraints on network capacity or processor load. (FDD 4.2.2)

[B2590] Resource Monitoring: GENI shall monitor the state of substrate components using sensors and other measurement processes. (FDD 4.2.2)

[B2600] Query processing (matching): GENI shall match researchers' requests with the available resources and select some set of components to host the slice. (FDD 4.2.2)

[B2610] Support for Gradual Experiment Refinement. GENI shall provide a smooth implementation path from simulation to deployment, with a single, integrated toolset. (FDD 4.2.3)

[B2620] Flexible programming environment: GENI shall export experimentation functionality in various forms to permit use by a wide variety of users. A beginner might desire a single monolithic tool to address all aspects of experimentation, such as resource discovery, experiment setup, experiment monitoring, and the gathering of results. But with additional experience, the programmer might desire greater flexibility and not be constrained by tools that prescribe a specific sequence of operations in setting up. It should support the use of existing tools to automate for all aspects of the experiment. (FDD 4.2.3)

[B2630] Incremental refinement and integration: The GENI Experimenter's Toolkit shall abstract mechanisms and services; e.g., for resource discovery, allocation, and content distribution, as and when they are made available. (FDD 4.2.3)

[B2640] Sophisticated fault handling: The experiment support toolkit shall enable the gradual "hardening" of distributed services by first allowing service execution in controlled settings, where faults can be deterministically and gradually administered and programs can be spared from having to face the full brunt of faults. (FDD 4.2.3)

5.3.2.16 Other Services (FDD 4.3)

[B2650] Bulk Data Transfer Service. GENI shall provide a bulk data transfer services to load experiment code onto a distributed set of machines. (FDD 4.3.1)

[B2660] Small Message Dissemination Service. GENI shall provide a small message dissemination (e.g., application level scalable reliable multicast) to reliably and efficiently distribute GENI control messages to a distributed set of machines. (FDD 4.3.1)

[B2670] Log/Sensor Data Collection Service. GENI shall provide log/sensor data collection from a distributed set of machines to the GENI central storage repository/analysis engine. (FDD 4.3.1)

[B2680] GENI-Internet Topology Information. GENI shall provide an information plane to provide topology information about GENI and the legacy Internet, to allow the efficient implementation and optimization of these communication services. (FDD 4.3.1)

[B2690] Node Storage. GENI shall provide local storage on the nodes for experimenter use (e.g., to store results, logs, or the data their experiments operate on). (FDD 4.3.1)

[B2700] Remote Access to Node Storage. GENI shall provide experimenters remote access to node storage resources, to debug their experiments or examine results. (FDD 4.3.2)

[B2710] High-Performance Experimenter Storage. GENI shall provide experimenters higher performance storage resources within a cluster of GENI machines for the data that their experiment operates on (e.g., cached web pages, network measurement data, and so on). (FDD 4.3.2)

[B2720] Storage for Management Use. GENI shall provide reliable storage for audit trails, resource use accounting, and tracking GENI resources such as equipment and network connectivity. (FDD 4.3.2)

[B2730] Direct Access to Node Storage. GENI shall provide direct access to storage resources by experimenters running code on an individual machine. (FDD 4.3.2)

[B2740] Database Access to Storage. GENI shall provide access through databases for more convenient access to structured data storage. (FDD 4.3.2)

[B2750] Storage Access for System Management. GENI shall provide convenient remote access for experiment and system management. (FDD 4.3.2)

[B2760] **Flexible Access to Storage**. GENI shall allow services to rapidly and efficiently access storage resources on a node both when they are running locally on that node and when they are remote. (FDD 4.3.2)

[B2770] **High-Performance Access to Storage**. GENI shall provide a high performance clustered file system for use within a single, geographically co-located collection of nodes. (FDD 4.3.2)

[B2780] **Parallel-enabled Storage**. GENI shall make it easy for services to access the parallelism available on the machines in its clusters. (FDD 4.3.2)

[B2790] **Best-Effort Storage**. GENI shall allocate long-term storage to projects when necessary, and not reclaim the storage until that time has expired. GENI management must create and follow guidelines about the frequency of intentional storage erasures (e.g., due to upgrades or routine maintenance). Finally, services must use replication or other strategies to ensure that they can survive accidents or erasures within the GENI parameters. (FDD 4.3.2)

[B2800] **Storage Service Extensibility**. GENI storage services shall be easily extensible. (FDD 4.3.2)

[B2810] **Storage Security**. GENI storage services shall provide strong security. (FDD 4.3.2)

[B2820] **Virtualized BGP**. GENI will provide a virtualized version of BGP for use by experimenters. (FDD 4.3.3)

[B2830] **Virtualized Data Plane**. GENI will provide a virtualized data plane packet processing and forwarding. (FDD 4.3.3)

[B2840] **Virtualized HTTP**. GENI shall provide a virtual version of HTTP for use by experimenters. (FDD 4.3.3)

[B2850] **Virtualized DNS**. GENI shall provide a virtual version of DNS for use by experimenters. (FDD 4.3.3)

[B2860] **Client Opt-In**. GENI shall provide downloadable software for popular operating systems (e.g., Symbian, Vista, XP, MacOS, Linux, and Windows CE) that allows them to interact with GENI experiments. (FDD 4.3.3)

[B2870] **Distributed Dynamic NAT**. GENI shall provide network addressing service that return packets to be delivered back to the originating slice. (FDD 4.3.3)

5.3.2.17 Instrumentation & Data Repository (FDD 4.4)

[B2880] **Ubiquitous Measurement Capability**. All GENI components shall be capable of collecting measurements. (FDD 4.4)

[B2890] **Unobtrusive Measurements**. Measurements shall have no measurable impact on GENI experiments. (FDD 4.4)

[B2900] **Measurement Extensibility**. GENI shall be capable of adding new instrumentation and/or new measurement synthesis capability. (FDD 4.4)

[B2910] Measurement System Reliability. GENI measurements shall be at least as available as the system upon which they are measuring. (FDD 4.4)

[B2920] Measurement Accuracy. GENI shall have the ability to measure detailed activity with high accuracy and precision from physical layer through application layer (includes the ability to calibrate measurements). (FDD 4.4)

[B2930] Experimenter Commanded Measurements. GENI experimenters shall be able to specify required measurements for an experiment in a slice (using either standard measurement types from a library or defining user specific measurements) and then having these measurements initialized in the infrastructure when an experiment is activated. (FDD 4.4)

[B2940] Measurement Access Control. GENI shall have the ability to specify what data is available from a particular device or collection of devices, to whom, and for how long. (FDD 4.4)

[B2950] Anonymized Measurements. GENI shall contain a secure central repository in which collected data can be anonymized and made available to researchers. (FDD 4.4)

[B2960] Data Analysis Warehouse. GENI shall include a "data analysis warehouse" where tools for visualizing, interpreting and reporting measurement data can be developed and made openly available. (FDD 4.4)

6 Requirements Traceability Table

The following two tables show how Research Plan requirements are associated with derived requirements. They are useful tools for assessing the completeness of both sets of requirements and for understanding how changes in customer requirements will impact the design and vice-versa.

[These tables are very rough. In particular, there tends to be a greater association between requirements than is likely deserved. It is expected that as the wording (and associated scope) of both the A-spec and B-spec requirements evolves, the associations will become more clear that this table will become more accurate.]

6.1 A→B Mapping

This table shows the association of B-spec requirements to each A-spec parent, permitting an evaluation of the completeness of the derivation. A complete derivation requires that all aspects of each A-spec requirement are reflected in the set of derived B-spec requirements.

A-Spec Top-Level Requirement	B-Spec Derived Requirement
[A1000] Number of Concurrent Experiments.	[B1040] Physical Substrate [B1370] Component Sharing [B1760] Sliver Creation [B1890] Number of GENI Edge Sites [B2550] Best Effort by Default [B2570] Resource Allocation Duration [B2790] Best-Effort Storage
[A1010] Inter-slice connectivity.	[B1110] National Fiber Facility [B1120] Edge-to-Core Tail Circuits. [B1800] Sliver Interposition [B2010] Maximum Tail Circuit Capacity [B2630] Incremental refinement and integration
[A1020] Controlled Isolation.	[B1070] Programmable Core Nodes. [B1080] Programmable Edge Nodes [B1380] Slice Isolation [B1400] Behavior Restriction [B1770] Sliver Isolation [B1810] Component Rate Limiting [B2470] Slice Performance Policing [B2540] Fine Granularity Admission Control
[A1030] Controlled Slice/Internet Interconnection.	[B1070] Programmable Core Nodes. [B1080] Programmable Edge Nodes [B1400] Behavior Restriction [B1420] Disconnection for Misbehavior [B1800] Sliver Interposition [B1690] Explicit Trust [B1700] Least Privilege [B1710] Revocation [B1730] Autonomy [B1780] Sliver Secure Login [B1820] Internet Filtering [B1860] Traffic Monitor [B2080] Realistic Traffic [B2170] Internet Rate & Packet Control
[A1040] Resources-Received Feedback.	[B1230] Support for Researchers [B1390] Self-limitation of Sharing
[A1050] Packet Formats.	[B1060] Computational Resources for Services. [B1070] Programmable Core Nodes. [B1080] Programmable Edge Nodes
[A1060] Network-level Node Functions.	[B1060] Computational Resources for Services [B1070] Programmable Core Nodes. [B1080] Programmable Edge Nodes [B1300] Execution Environment [B1440] Virtual Socket Interfaces [B2000] L3 Tail Circuits [B2030] Backbone-enabled Network Experiments [B2050] Backbone Packet Processing [B2070] Packet-switched & Circuit-switched Backbone Experiments

[A1070] Link-layer Node Functions.	[B1060] Computational Resources for Services. [B1070] Programmable Core Nodes. [B1080] Programmable Edge Nodes [B1450] Virtual Link Interfaces [B1990] L2 Tail Circuits [B2060] Backbone Programmable Framers
[A1080] Multiplexing-free Paths.	[B1890] Number of GENI Edge Sites [B1960] Dark Fiber Tail Circuits [B1970] Shared Dark Fiber Tail Circuits [B1980] Leased PHY Tail Circuits
[A1090] Reconfigurable Optical Circuits.	[B1110] National Fiber Facility [B1270] Substrate Management [B1900] Dedicated Tail Circuits [B1960] Dark Fiber Tail Circuits [B1970] Shared Dark Fiber Tail Circuits [B1980] Leased PHY Tail Circuits [B2070] Packet-switched & Circuit-switched Backbone Experiments
[A1100] Current Representative Technologies.	[B1100] Mobile Devices. [B1110] National Fiber Facility [B1150] Urban 802.11-based Mesh Wireless Subnets [B1160] Wide-Area Suburban 3G/WiMax-based Wireless Subnets [B1170] Cognitive Radio Subnets [B1180] Sensor Subnets. [B2200] Untethered Wireless Nodes [B2220] Non-GENI Wireless Devices [B2230] GENI Wireless "Opt-in". [B2260] 802.11 Subnet [B2270] Cellular/WiMax-WiFi Hybrid Subnet [B2280] Cognitive Radio Subnet [B2290] Sensor Network Kit [B2300] Local Sensor Subnets [B2310] Sensornet Outdoor Deployment [B2320] Sensornet Indoor Deployment [B2330] External Sensornet Interconnect
[A1110] Future Representative Technologies.	[B1050] Expandable Substrate.
[A1120] Experimenter Access to Technology Features.	[B1230] Support for Researchers [B1310] Low-level Information Exposure [B1340] Component Managers [B1490] Radio Interface Configuration [B1850] Component Power Control & Systems Errors [B1880] Hardware Monitor
[A1130] Attractive to Users.	[B1000] User Service Subsystem [B1190] Emulation Grids [B1220] Support for Developers [B1290] Experimenter Services [B2340] Emulation Grids [B2580] Resource Discovery [B2600] Query processing (matching)
[A1140] User Reach.	[B1110] National Fiber Facility [B1130] Core-to-Internet Tail Circuits [B1260] Resource Allocation & Usage Policy Mechanism
[A1150] Tools for End-system Modification.	[B1230] Support for Researchers [B1240] Support for End Users [B2180] Tunnels [B2860] Client Opt-In
[A1160] Diversity of Path Characteristics.	[B1460] Constant Bit Rate [B2030] Backbone-enabled Network Experiments [B2550] Best Effort by Default
[A1170] Internet-Connected Users.	[B1130] Core-to-Internet Tail Circuits [B2180] Tunnels

[A1180]	Directly-Connected Users	[B1890] [B2020] [B2030]	Number of GENI Edge Sites Backbone Connectivity between GENI Edge Sites Backbone-enabled Network Experiments
[A1190]	Path Diversity.	[B2090] [B2140] [B2150]	Backbone Path Diversity Number of Backbone Nodes Backbone Node Fan-out
[A1200]	Underlying Fiber Paths.	[B1110] [B1120] [B2100] [B2140]	National Fiber Facility Edge-to-Core Tail Circuits. Fiber-hub Cities Number of Backbone Nodes
[A1210]	Topology-Related Delays.	[B2120] [B2130]	Maximum End-to-end Latency Maximum Hop-by-hop Latency
[A1220]	Major Interconnection Points – ISP	[B1110] [B1130] [B1910] [B1930] [B2040] [B2080] [B2190] [B2870]	National Fiber Facility Core-to-Internet Tail Circuits Internet Connectivity Peering Co-location Model Backbone Internet Connectivity Realistic Traffic BGP Gateway Distributed Dynamic NAT
[A1230]	Major Interconnection Points – Exchange Points.	[B1110] [B1130] [B1910] [B1930] [B2040] [B2080] [B2110] [B2190] [B2870]	National Fiber Facility Core-to-Internet Tail Circuits Internet Connectivity Peering Co-location Model Backbone Internet Connectivity Realistic Traffic Backbone sites at IXPs BGP Gateway Distributed Dynamic NAT
[A1240]	Major Interconnection Points – GigaPops.	[B1110] [B1130] [B1920]	National Fiber Facility Core-to-Internet Tail Circuits Site Locations
[A1250]	Physical Distribution.	[B1070] [B1080] [B1110] [B1190] [B1890] [B1920] [B1930] [B1940] [B1950] [B2340]	Programmable Core Nodes. Programmable Edge Nodes. National Fiber Facility Emulation Grids Number of GENI Edge Sites Site Locations Peering Co-location Model DMZ Co-location Model Embedded Co-location Model Emulation Grids
[A1260]	Scale.	[B1110] [B1530] [B1650] [B1750] [B1890] [B2010] [B2870]	National Fiber Facility Scalable Name Systems Credential Delegation Local Management of Security Information Number of GENI Edge Sites Maximum Tail Circuit Capacity Distributed Dynamic NAT
[A1270]	Virtual Failure On Command.	[B1070] [B1080] [B1300]	Programmable Core Nodes. Programmable Edge Nodes Execution Environment
[A1280]	Failure Notification.	[B1850] [B2370] [B2380] [B2390] [B2400]	Component Power Control & Systems Errors Fault Management System Fault Data Acquisition Fault Management Rules FMS-Operator Interface
[A1290]	Virtualization Realism.	[B1230] [B1300] [B1310] [B1390] [B2460] [B2470]	Support for Researchers Execution Environment Low-level Information Exposure Self-limitation of Sharing Performance Management System Slice Performance Policing

[A1300] Virtualized Management Interfaces.	[B1020] GENI Management Core [B1220] Support for Developers [B1270] Substrate Management [B1310] Low-level Information Exposure [B1340] Component Managers [B1490] Radio Interface Configuration [B1880] Hardware Monitor [B2210] Remote Management for Wireless Node [B2240] Wireless Spectrum Management
[A1310] Secure Management Technologies.	[B1210] Support for Administrators [B1270] Substrate Management [B1570] User Vetting [B1590] User-Slice Mapping [B1640] User Identity [B1650] Credential Delegation [B1660] Identifier System [B1870] Secure Boot Monitor
[A1320] Security Fidelity.	[B1200] Support for Owners [B1210] Support for Administrators [B1250] Support for Third Parties [B1680] Transferable Slice Control [B1690] Explicit Trust [B1700] Least Privilege [B1710] Revocation [B1730] Autonomy [B1740] Security Usability [B2440] Accounting Management System [B2480] Security Management System [B2490] Account Creation Oversight [B2810] Storage Security [B2940] Measurement Access Control
[A1330] Isolation for Security.	[B1200] Support for Owners [B1210] Support for Administrators [B1250] Support for Third Parties [B1380] Slice Isolation [B1400] Behavior Restriction [B1420] Disconnection for Misbehavior

[A1340] Ease of Use.	[B1220] Support for Developers [B1230] Support for Researchers [B1290] Experimenter Services [B1350] CM Interface Definitions [B1440] Virtual Socket Interfaces [B1450] Virtual Link Interfaces [B1470] Point-to-point Virtual Links [B1480] Virtual Radio Interface [B1500] Virtual Wire Interface [B1510] Virtual Wire Types [B1520] Component Names [B1540] Multi-point Virtual Links [B1560] Slice Names [B1610] Experiment Configuration [B1680] Transferable Slice Control [B1740] Security Usability [B1790] Reboot Signal [B2500] Problem Tracking [B2510] Resource Allocation Portal [B2520] Slice Embedding Portal [B2530] Experimenter's Workbench [B2580] Resource Discovery [B2610] Support for Gradual Experiment Refinement [B2620] Flexible programming environment [B2630] Incremental refinement and integration [B2640] Sophisticated fault handling [B2650] Bulk Data Transfer Service [B2660] Small Message Dissemination Service [B2670] Log/Sensor Data Collection Service [B2680] GENI-Internet Topology Information [B2690] Node Storage [B2700] Remote Access to Node Storage [B2710] High-Performance Experimenter Storage [B2730] Direct Access to Node Storage [B2740] Database Access to Storage [B2760] Flexible Access to Storage [B2770] High-Performance Access to Storage [B2780] Parallel-enabled Storage [B2820] Virtualized BGP [B2830] Virtualized Data Plane [B2840] Virtualized HTTP [B2850] Virtualized DNS [B2960] Data Analysis Warehouse
[A1350] Power Users.	[B1220] Support for Developers [B1230] Support for Researchers [B1290] Experimenter Services [B1600] Resource Database [B1670] Common Resource Specification data structure [B2620] Flexible programming environment
[A1360] Observability.	[B1230] Support for Researchers [B1310] Low-level Information Exposure [B2590] Resource Monitoring [B2880] Ubiquitous Measurement Capability [B2890] Unobtrusive Measurements [B2900] Measurement Extensibility [B2910] Measurement System Reliability [B2920] Measurement Accuracy [B2930] Experimenter Commanded Measurements [B2940] Measurement Access Control [B2950] Anonymized Measurements

[A1370] Traceback.	[B1200] Support for Owners [B1210] Support for Administrators [B1250] Support for Third Parties [B1410] Audits of Internet-bound Traffic [B1720] Auditability [B1830] Auditability [B2160] Internet Packet Logging [B2250] RF Activity Profiling [B2720] Storage for Management Use
[A1380] Fail-safe.	[B1200] Support for Owners [B1210] Support for Administrators [B1250] Support for Third Parties [B1420] Disconnection for Misbehavior [B1430] Safe State [B1790] Reboot Signal [B1840] Safe State
[A1390] Operational Lifetime.	
[A1400] Facility Opt-in.	[B1030] Open Interfaces [B1260] Resource Allocation & Usage Policy Mechanism [B1270] Substrate Management [B1280] Decentralized Management [B1360] Resource Allocation Policy Control [B1580] Naming External Components [B1650] Credential Delegation [B1930] Peering Co-location Model [B1940] DMZ Co-location Model [B1950] Embedded Co-location Model [B2330] External Sensornet Interconnect [B2410] Administrator Database
[A1410] Upgrade In-Place	[B1010] Evolvability [B1050] Expandable Substrate. [B1270] Substrate Management [B1330] Design for Asynchronous System Development [B2430] Provisioning New Components [B2800] Storage Service Extensibility

B-specs with no A-spec parent:

[B2420] Configuration Management System

[B2450] Resource Allocation Policies

[B2560] Experiment Priority Policy

6.2 B→A Mapping

This table, a reverse mapping, is also useful since many B-spec requirements have multiple A-spec parents. The B-to-A mapping is useful for evaluating the impact on customer requirements should the design requirements change.

B-Spec Derived Requirement	A-Spec Top-level Requirement
[B1000] User Service Subsystem	[A1130] Attractive to Users.
[B1010] Evolvability	[A1410] Upgrade In-Place
[B1020] GENI Management Core	[A1300] Virtualized Management Interfaces.
[B1030] Open Interfaces	[A1400] Facility Opt-in.
[B1040] Physical Substrate	[A1000] Number of Concurrent Experiments.
[B1050] Expandable Substrate.	[A1110] Future Representative Technologies. [A1410] Upgrade In-Place
[B1060] Computational Resources for Services	[A1060] Network-level Node Functions. [A1070] Link-layer Node Functions. [A1050] Packet Formats.
[B1070] Programmable Core Nodes.	[A1020] Controlled Isolation. [A1030] Controlled Slice/Internet Interconnection. [A1050] Packet Formats. [A1060] Network-level Node Functions. [A1070] Link-layer Node Functions. [A1250] Physical Distribution. [A1270] Virtual Failure On Command. [A1020] Controlled Isolation. [A1030] Controlled Slice/Internet Interconnection. [A1050] Packet Formats. [A1060] Network-level Node Functions. [A1070] Link-layer Node Functions. [A1270] Virtual Failure On Command. [A1250] Physical Distribution.
[B1100] Mobile Devices.	[A1100] Current Representative Technologies.
[B1110] National Fiber Facility	[A1010] Inter-slice connectivity. [A1090] Reconfigurable Optical Circuits. [A1100] Current Representative Technologies. [A1140] User Reach. [A1200] Underlying Fiber Paths. [A1220] Major Interconnection Points – ISP [A1230] Major Interconnection Points – Exchange Points. [A1240] Major Interconnection Points – GigaPops. [A1250] Physical Distribution. [A1260] Scale.
[B1120] Edge-to-Core Tail Circuits.	[A1010] Inter-slice connectivity. [A1200] Underlying Fiber Paths.
[B1130] Core-to-Internet Tail Circuits	[A1140] User Reach. [A1170] Internet-Connected Users. [A1220] Major Interconnection Points – ISP [A1230] Major Interconnection Points – Exchange Points. [A1240] Major Interconnection Points – GigaPops.
[B1150] Urban 802.11-based Mesh Wireless Subnets	[A1100] Current Representative Technologies.
[B1160] Wide-Area Suburban 3G/WiMax-based Wireless Subnets	[A1100] Current Representative Technologies.
[B1170] Cognitive Radio Subnets	[A1100] Current Representative Technologies.
[B1180] Sensor Subnets.	[A1100] Current Representative Technologies.
[B1190] Emulation Grids	[A1130] Attractive to Users. [A1250] Physical Distribution.
[B1200] Support for Owners	[A1320] Security Fidelity. [A1330] Isolation for Security. [A1370] Traceback. [A1380] Fail-safe.

[B1210] Support for Administrators	[A1310] Secure Management Technologies. [A1320] Security Fidelity. [A1330] Isolation for Security. [A1370] Traceback. [A1380] Fail-safe.
[B1220] Support for Developers	[A1130] Attractive to Users. [A1300] Virtualized Management Interfaces. [A1340] Ease of Use. [A1350] Power Users.
[B1230] Support for Researchers	[A1040] Resources-Received Feedback. [A1120] Experimenter Access to Technology Features. [A1150] Tools for End-system Modification. [A1290] Virtualization Realism. [A1340] Ease of Use. [A1350] Power Users. [A1360] Observability.
[B1240] Support for End Users	[A1150] Tools for End-system Modification.
[B1250] Support for Third Parties	[A1320] Security Fidelity. [A1330] Isolation for Security. [A1370] Traceback. [A1380] Fail-safe.
[B1260] Resource Allocation & Usage Policy Mechanism	[A1140] User Reach. [A1400] Facility Opt-in.
[B1270] Substrate Management	[A1090] Reconfigurable Optical Circuits. [A1300] Virtualized Management Interfaces. [A1310] Secure Management Technologies. [A1400] Facility Opt-in. [A1410] Upgrade In-Place
[B1280] Decentralized Management	[A1400] Facility Opt-in.
[B1290] Experimenter Services	[A1130] Attractive to Users. [A1340] Ease of Use. [A1350] Power Users.
[B1300] Execution Environment	[A1060] Network-level Node Functions. [A1270] Virtual Failure On Command. [A1290] Virtualization Realism.
[B1310] Low-level Information Exposure	[A1120] Experimenter Access to Technology Features. [A1290] Virtualization Realism. [A1300] Virtualized Management Interfaces. [A1360] Observability.
[B1330] Design for Asynchronous System Development	[A1410] Upgrade In-Place
[B1340] Component Managers	[A1120] Experimenter Access to Technology Features. [A1300] Virtualized Management Interfaces.
[B1350] CM Interface Definitions	[A1340] Ease of Use.
[B1360] Resource Allocation Policy Control	[A1400] Facility Opt-in.
[B1370] Component Sharing	[A1000] Number of Concurrent Experiments.
[B1380] Slice Isolation	[A1020] Controlled Isolation. [A1330] Isolation for Security.
[B1390] Self-limitation of Sharing	[A1040] Resources-Received Feedback. [A1290] Virtualization Realism.
[B1400] Behavior Restriction	[A1020] Controlled Isolation. [A1030] Controlled Slice/Internet Interconnection. [A1330] Isolation for Security.
[B1410] Audits of Internet-bound Traffic	[A1370] Traceback.
[B1420] Disconnection for Misbehavior	[A1030] Controlled Slice/Internet Interconnection. [A1330] Isolation for Security. [A1380] Fail-safe.
[B1430] Safe State	[A1380] Fail-safe.
[B1440] Virtual Socket Interfaces	[A1060] Network-level Node Functions. [A1340] Ease of Use.
[B1450] Virtual Link Interfaces	[A1070] Link-layer Node Functions. [A1340] Ease of Use.
[B1460] Constant Bit Rate	[A1160] Diversity of Path Characteristics.

[B1470] Point-to-point Virtual Links	[A1340] Ease of Use.
[B1480] Virtual Radio Interface	[A1340] Ease of Use.
[B1490] Radio Interface Configuration	[A1120] Experimenter Access to Technology Features. [A1300] Virtualized Management Interfaces.
[B1500] Virtual Wire Interface	[A1340] Ease of Use.
[B1510] Virtual Wire Types	[A1340] Ease of Use.
[B1520] Component Names	[A1340] Ease of Use.
[B1530] Scalable Name Systems	[A1260] Scale.
[B1540] Multi-point Virtual Links	[A1340] Ease of Use.
[B1560] Slice Names	[A1340] Ease of Use.
[B1570] User Vetting	[A1310] Secure Management Technologies.
[B1580] Naming External Components	[A1400] Facility Opt-in.
[B1590] User-Slice Mapping	[A1310] Secure Management Technologies.
[B1600] Resource Database	[A1350] Power Users.
[B1610] Experiment Configuration	[A1340] Ease of Use.
[B1640] User Identity	[A1310] Secure Management Technologies.
[B1650] Credential Delegation	[A1260] Scale. [A1310] Secure Management Technologies. [A1400] Facility Opt-in.
[B1660] Identifier System	[A1310] Secure Management Technologies.
[B1670] Common Resource Specification data structure	[A1350] Power Users.
[B1680] Transferable Slice Control	[A1320] Security Fidelity. [A1340] Ease of Use.
[B1690] Explicit Trust	[A1030] Controlled Slice/Internet Interconnection. [A1320] Security Fidelity.
[B1700] Least Privilege	[A1030] Controlled Slice/Internet Interconnection. [A1320] Security Fidelity.
[B1710] Revocation	[A1030] Controlled Slice/Internet Interconnection. [A1320] Security Fidelity.
[B1720] Auditability	[A1370] Traceback.
[B1730] Autonomy	[A1030] Controlled Slice/Internet Interconnection. [A1320] Security Fidelity.
[B1740] Security Usability	[A1320] Security Fidelity. [A1340] Ease of Use.
[B1750] Local Management of Security Information	[A1260] Scale.
[B1760] Sliver Creation	[A1000] Number of Concurrent Experiments.
[B1770] Sliver Isolation	[A1020] Controlled Isolation.
[B1780] Sliver Secure Login	[A1030] Controlled Slice/Internet Interconnection.
[B1790] Reboot Signal	[A1340] Ease of Use. [A1380] Fail-safe.
[B1800] Sliver Interposition	[A1010] Inter-slice connectivity. [A1030] Controlled Slice/Internet Interconnection.
[B1810] Component Rate Limiting	[A1020] Controlled Isolation.
[B1820] Internet Filtering	[A1030] Controlled Slice/Internet Interconnection.
[B1830] Auditability	[A1370] Traceback.
[B1840] Safe State	[A1380] Fail-safe.
[B1850] Component Power Control & Systems Errors	[A1120] Experimenter Access to Technology Features. [A1280] Failure Notification.
[B1860] Traffic Monitor	[A1030] Controlled Slice/Internet Interconnection.
[B1870] Secure Boot Monitor	[A1310] Secure Management Technologies.
[B1880] Hardware Monitor	[A1120] Experimenter Access to Technology Features. [A1300] Virtualized Management Interfaces.
[B1890] Number of GENI Edge Sites	[A1000] Number of Concurrent Experiments. [A1080] Multiplexing-free Paths. [A1180] Directly-Connected Users [A1250] Physical Distribution. [A1260] Scale.
[B1900] Dedicated Tail Circuits	[A1090] Reconfigurable Optical Circuits.
[B1910] Internet Connectivity	[A1220] Major Interconnection Points – ISP [A1230] Major Interconnection Points – Exchange Points.

[B1920] Site Locations	[A1240] Major Interconnection Points – GigaPops. [A1250] Physical Distribution.
[B1930] Peering Co-location Model	[A1220] Major Interconnection Points – ISP [A1230] Major Interconnection Points – Exchange Points. [A1250] Physical Distribution. [A1400] Facility Opt-in.
[B1940] DMZ Co-location Model	[A1250] Physical Distribution. [A1400] Facility Opt-in.
[B1950] Embedded Co-location Model	[A1250] Physical Distribution. [A1400] Facility Opt-in.
[B1960] Dark Fiber Tail Circuits	[A1080] Multiplexing-free Paths. [A1090] Reconfigurable Optical Circuits.
[B1970] Shared Dark Fiber Tail Circuits	[A1080] Multiplexing-free Paths. [A1090] Reconfigurable Optical Circuits.
[B1980] Leased PHY Tail Circuits	[A1080] Multiplexing-free Paths. [A1090] Reconfigurable Optical Circuits.
[B1990] L2 Tail Circuits	[A1070] Link-layer Node Functions. [A1060] Network-level Node Functions.
[B2010] Maximum Tail Circuit Capacity	[A1010] Inter-slice connectivity. [A1260] Scale.
[B2020] Backbone Connectivity between GENI Edge Sites	[A1180] Directly-Connected Users
[B2030] Backbone-enabled Network Experiments	[A1060] Network-level Node Functions. [A1160] Diversity of Path Characteristics. [A1180] Directly-Connected Users
[B2040] Backbone Internet Connectivity	[A1220] Major Interconnection Points – ISP [A1230] Major Interconnection Points – Exchange Points.
[B2050] Backbone Packet Processing	[A1060] Network-level Node Functions.
[B2060] Backbone Programmable Framers	[A1070] Link-layer Node Functions.
[B2070] Packet-switched & Circuit-switched Backbone Experiments	[A1060] Network-level Node Functions. [A1090] Reconfigurable Optical Circuits.
[B2080] Realistic Traffic	[A1030] Controlled Slice/Internet Interconnection. [A1220] Major Interconnection Points – ISP [A1230] Major Interconnection Points – Exchange Points.
[B2090] Backbone Path Diversity	[A1190] Path Diversity.
[B2100] Fiber-hub Cities	[A1200] Underlying Fiber Paths.
[B2110] Backbone sites at IXPs	[A1230] Major Interconnection Points – Exchange Points.
[B2120] Maximum End-to-end Latency	[A1210] Topology-Related Delays.
[B2130] Maximum Hop-by-hop Latency	[A1210] Topology-Related Delays.
[B2140] Number of Backbone Nodes	[A1190] Path Diversity. [A1200] Underlying Fiber Paths.
[B2150] Backbone Node Fan-out	[A1190] Path Diversity.
[B2160] Internet Packet Logging	[A1370] Traceback.
[B2170] Internet Rate & Packet Control	[A1030] Controlled Slice/Internet Interconnection.
[B2180] Tunnels	[A1150] Tools for End-system Modification. [A1170] Internet-Connected Users.
[B2190] BGP Gateway	[A1220] Major Interconnection Points – ISP [A1230] Major Interconnection Points – Exchange Points.
[B2200] Untethered Wireless Nodes	[A1100] Current Representative Technologies.
[B2210] Remote Management for Wireless Node	[A1300] Virtualized Management Interfaces.
[B2220] Non-GENI Wireless Devices	[A1100] Current Representative Technologies.
[B2230] GENI Wireless "Opt-in".	[A1100] Current Representative Technologies.
[B2240] Wireless Spectrum Management	[A1300] Virtualized Management Interfaces.
[B2250] RF Activity Profiling	[A1370] Traceback.
[B2260] 802.11 Subnet	[A1100] Current Representative Technologies.
[B2270] Cellular/WiMax-WiFi Hybrid Subnet	[A1100] Current Representative Technologies.
[B2280] Cognitive Radio Subnet	[A1100] Current Representative Technologies.
[B2290] Sensor Network Kit	[A1100] Current Representative Technologies.
[B2300] Local Sensor Subnets	[A1100] Current Representative Technologies.
[B2310] Sensornet Outdoor Deployment	[A1100] Current Representative Technologies.
[B2320] Sensornet Indoor Deployment	[A1100] Current Representative Technologies.
[B2330] External Sensornet Interconnect	[A1100] Current Representative Technologies. [A1400] Facility Opt-in.

[B2340] Emulation Grids	[A1130] Attractive to Users. [A1250] Physical Distribution.
[B2370] Fault Management System	[A1280] Failure Notification.
[B2380] Fault Data Acquisition	[A1280] Failure Notification.
[B2390] Fault Management Rules	[A1280] Failure Notification.
[B2400] FMS-Operator Interface	[A1280] Failure Notification.
[B2410] Administrator Database	[A1400] Facility Opt-in.
[B2430] Provisioning New Components	[A1410] Upgrade In-Place
[B2440] Accounting Management System	[A1320] Security Fidelity.
[B2460] Performance Management System	[A1290] Virtualization Realism.
[B2470] Slice Performance Policing	[A1020] Controlled Isolation. [A1290] Virtualization Realism.
[B2480] Security Management System	[A1320] Security Fidelity.
[B2490] Account Creation Oversight	[A1320] Security Fidelity.
[B2500] Problem Tracking	[A1340] Ease of Use.
[B2510] Resource Allocation Portal	[A1340] Ease of Use.
[B2520] Slice Embedding Portal	[A1340] Ease of Use.
[B2530] Experimenter's Workbench	[A1340] Ease of Use.
[B2540] Fine Granularity Admission Control	[A1020] Controlled Isolation.
[B2550] Best Effort by Default	[A1000] Number of Concurrent Experiments.
[B2550] Best Effort by Default	[A1160] Diversity of Path Characteristics.
[B2570] Resource Allocation Duration	[A1000] Number of Concurrent Experiments.
[B2580] Resource Discovery	[A1130] Attractive to Users. [A1340] Ease of Use.
[B2590] Resource Monitoring	[A1360] Observability.
[B2600] Query processing (matching)	[A1130] Attractive to Users.
[B2610] Support for Gradual Experiment Refinement	[A1340] Ease of Use.
[B2620] Flexible programming environment	[A1340] Ease of Use.
	[A1350] Power Users.
[B2630] Incremental refinement and integration	[A1010] Inter-slice connectivity.
	[A1340] Ease of Use.
[B2640] Sophisticated fault handling	[A1340] Ease of Use.
[B2650] Bulk Data Transfer Service	[A1340] Ease of Use.
[B2660] Small Message Dissemination Service	[A1340] Ease of Use.
[B2670] Log/Sensor Data Collection Service	[A1340] Ease of Use.
[B2680] GENI-Internet Topology Information	[A1340] Ease of Use.
[B2690] Node Storage	[A1340] Ease of Use.
[B2700] Remote Access to Node Storage	[A1340] Ease of Use.
[B2710] High-Performance Experimenter Storage	[A1340] Ease of Use.
[B2720] Storage for Management Use	[A1370] Traceback.
[B2730] Direct Access to Node Storage	[A1340] Ease of Use.
[B2740] Database Access to Storage	[A1340] Ease of Use.
[B2760] Flexible Access to Storage	[A1340] Ease of Use.
[B2770] High-Performance Access to Storage	[A1340] Ease of Use.
[B2780] Parallel-enabled Storage	[A1340] Ease of Use.
[B2790] Best-Effort Storage	[A1000] Number of Concurrent Experiments.
[B2800] Storage Service Extensibility	[A1410] Upgrade In-Place
[B2810] Storage Security	[A1320] Security Fidelity.
[B2820] Virtualized BGP	[A1340] Ease of Use.
[B2830] Virtualized Data Plane	[A1340] Ease of Use.
[B2840] Virtualized HTTP	[A1340] Ease of Use.
[B2850] Virtualized DNS	[A1340] Ease of Use.
[B2860] Client Opt-In	[A1150] Tools for End-system Modification.
[B2870] Distributed Dynamic NAT	[A1220] Major Interconnection Points – ISP [A1230] Major Interconnection Points – Exchange Points. [A1260] Scale.
[B2880] Ubiquitous Measurement Capability	[A1360] Observability.
[B2890] Unobtrusive Measurements	[A1360] Observability.
[B2900] Measurement Extensibility	[A1360] Observability.
[B2910] Measurement System Reliability	[A1360] Observability.
[B2920] Measurement Accuracy	[A1360] Observability.

[B2930] Experimenter Commanded Measurements	[A1360] Observability.
[B2940] Measurement Access Control	[A1320] Security Fidelity. [A1360] Observability.
[B2950] Anonymized Measurements	[A1360] Observability.
[B2960] Data Analysis Warehouse	[A1340] Ease of Use.
	[A1390] Operational Lifetime.