

GENI Design Principles

GENI Planning Group



GENI can lead to a future Internet that is more secure, available, manageable, and efficient.

The Global Environment for Network Innovations (www.geni.net) is a major planned initiative of the US National Science Foundation to build an open, large-scale, realistic experimental facility for evaluating new network architectures. Unlike conventional testbeds, GENI is meant to

- support multiple experiments running in parallel,
- carry real traffic on behalf of end users, and
- connect to the existing Internet to reach external sites.

The facility's goal is to change the way we design networked and distributed systems, creating over time new paradigms that integrate rigorous theoretical understanding with compelling and thorough experimental validation. The research that GENI enables can lead to a future Internet that is more secure, available, manageable, and efficient, and better at handling mobile nodes.

GENI is intended to support two general kinds of activities:

- running controlled experiments to evaluate design, implementation, and engineering choices; and
- deploying prototype systems and

learning from observations of how they behave under real usage.

Classical science equates experimentation with the former, but computer science is different. We benefit from prototypes because building something, and watching it run, helps us identify implicit assumptions, the need for different functionality, surprising behavior, unexpected limitations, and so on.

In this sense, such “experimental systems” work is like constructing a building—engineering principles tell you whether a design is sound, but you must build it and use it to decide how well it serves its purpose. GENI must support both types of activities, and it must support moving from one to the other.

Designing GENI is a multiyear process, involving numerous people. In January 2006, the GENI Planning Group completed an initial Project Execution Plan, available online at the GENI Web site. GENI will continue to evolve, even after construction, as new technologies become available and new user requirements come into focus.

Having a well-articulated set of design principles is crucial for making decisions about GENI's design and construction to determine how to best manage and allocate limited resources such as construction budget and facility capacity. Here we outline the prin-

ciples and priorities that have guided the Planning Group and, with community feedback, can help direct the process going forward. A more detailed description of the design principles is available at the GENI Web site.

REQUIREMENTS

GENI comprises a collection of hardware resources, including compute nodes, backbone links, tail circuits, storage capacity, customizable routers, and wireless subnets. Each experiment using GENI will run on some subset of the GENI resources. The resources bound to a particular experiment are a *slice*. GENI includes management software that is used to allocate resources to slices, embed slices in these resources, and ensure that slices do not interfere with one another.

To support both controlled experiments and long-term deployment studies, GENI must satisfy several key requirements: sliceability, generality, fidelity, user access, controlled isolation, diversity and extensibility, wide deployment, observability, and federation and sustainability.

In addition to this set of requirements, which focus on GENI's unique characteristics as a research instrument, we also expect GENI to adhere to the same set of design goals that apply to any widely distributed system that serves a large user community: that it be secure, robust, efficient, easy to use, and manageable.

Sliceability

To be cost-effective, GENI must be a shared facility that researchers can use to support multiple experiments running on behalf of many independent research groups. We expect that on the order of 1,000 researchers will be utilizing GENI.

Virtualization is a key technology that supports this goal as it lets the facility be multiplexed across multiple researchers. Another approach is to partition resources among researchers, either in time (analogous to astronomers sharing a telescope) or in space (giving a single researcher exclusive access to some subset of resources).

Generality

GENI should give each researcher the flexibility needed to perform the desired experiment. This means that each component should be programmable, so that researchers aren't limited to experimenting with small changes to preexisting functionality.

Researchers should not be required to program their entire experiment from scratch—they should be able to take advantage of previously defined functionality and abstractions—but they shouldn't be restricted by such existing functionality.

Fidelity

GENI should permit experiments that correlate with what might be expected in a real network. This means individual components must expose functionality at the right abstraction level, and it must be possible to arrange these components into a representative network.

Clearly, much responsibility for conducting a meaningful experiment falls to the researcher, but one of the facility's goals is to not unduly limit a researcher's ability to conduct such an experiment.

User access

To support meaningful deployment studies, GENI must make it easy for a broad mix of users to "opt in" to experimental services. This means providing physical connectivity to a large user community, along with mechanisms that let users easily join one or more experimental services; allowing experiments to run continuously (as no user will want to use a service that is up for only a limited period of time each day); and connecting GENI to the legacy Internet (to gain the benefit of interacting with existing Internet services and their users).

Controlled isolation

GENI must support strong isolation between slices so that experiments don't interfere with one another. GENI's isolation mechanisms should be sufficiently robust to make reproducible

experiments possible; to the extent they are not, it should provide enough feedback about what resources a slice actually receives to enable researchers to evaluate their results' validity.

At the same time, GENI must support controlled interconnection of slices to each other and to the current Internet, letting researchers build directly on one another's work and draw on existing Internet users and resources. This implies mechanisms that enable user opt in and desirable data exchange between slices, while keeping undesirable outside factors from interfering with GENI experiments and containing such experiments so that they don't adversely affect the rest of the Internet.

Access to GENI can't be limited to only those few sites that host backbone nodes.

Diversity and extensibility

GENI must include a wide class of networking technologies that span the spectrum of wired and wireless technologies available today. It must also be extensible—with explicitly defined procedures and system interfaces—to make it easy to incorporate additional technologies, including those that don't exist today.

This will allow GENI to benefit a broad range of researchers, remain useful over a much longer lifespan, support GENI's role as a low-friction vehicle for deployment of new technologies by both academic researchers and industrial partners, and foster close collaboration between device and systems researchers.

Wide deployment

GENI must have as wide a reach as possible to support experimentation at scale and to maximize the opportunity to attract real users. Access can't be limited to only those few sites that host backbone nodes. Wide deployment

also implies a rich interconnection of the facility to the legacy Internet.

Observability

GENI must offer strong support for measurement-based quantitative research. This means that GENI's resources, along with all the network systems deployed on it, must be heavily instrumented. Researchers must be able to generate and archive data and develop analysis tools.

Federation and sustainability

GENI must be designed to last 15 to 20 years, extending well beyond a construction phase of 5 to 7 years. To ensure sustainability, participating institutions—including countries—should be able to contribute resources in return for access to the entire facility's resources. New research communities should be able to opt in by connecting their purpose-built networks, including dedicated transmission pipes and sensor networks, into GENI and running their applications and services in a slice of GENI.

Both of these scenarios imply the need to support federation. In addition, GENI must be designed with operational costs in mind, including hardware upgrades, software maintenance, and ongoing operational support.

TENSIONS

Many of these requirements are synergistic. For example, widespread deployment naturally supports greater user access, and making GENI extensible to accommodate new technologies is consistent with its support for federation to let communities and partners add their resources to GENI. On the other hand, intrinsic tensions exist among some of these requirements, as well as between various types of experiments that value the requirements differently.

Sliceability versus fidelity

Balancing sliceability and fidelity is one of the most fundamental challenges facing GENI. On the one hand, virtualizing the underlying hardware lets many researchers share a common

set of resources and can increase flexibility by synthesizing multiple or higher-function virtual environments from a single physical resource. On the other hand, virtualization

- allows for the possibility that one experiment might interfere with another experiment, and
- potentially hides certain capabilities and properties of the underlying hardware.

Both properties give the facility less fidelity than if a researcher had the resources exclusively.

On the surface, this particular conflict is easy to resolve—GENI should provide strong isolation between slices and the lowest level of virtualization that the technology allows. Any given component might not provide the desired level on day one, but advancing the state of the art in virtualization over GENI's lifetime is an ongoing objective.

Higher levels of abstraction should also be retained for those experiments that don't want to be exposed to low-level details, but virtualization should be pushed as low as technically possible, cost allowing.

However, some will argue that any amount of virtualization is too much and that their research requires access to "bare metal." This might be because of the need for access to a component-specific feature or because virtualization introduces too much unpredictability in timing measurements. There might also be resources that simply can't be virtualized.

GENI doesn't preclude the possibility that raw hardware elements can be allocated to some slices—partitioning is another way of implementing slices—but doing so is likely to come at one of two costs.

If resources are partitioned in time, a given resource might not sustain a real user workload, thereby limiting its appropriateness for deployment studies. Some fraction of GENI's resources can be shared in this way, as long as sufficient capacity is available to support deployment studies.

If resources are partitioned in space, only a limited number of researchers might be able to include a given resource in their slice. This might be necessary for certain high-cost resources that can't be easily virtualized, in which case the community will have to either prioritize its research or find ways to synthesize its many experimental systems into a few comprehensive systems.

While around 1,000 researchers might share GENI as a whole, only tens of research projects would share access to any high-cost, nonvirtualizable resource in this way.

No individual technology is fully validated until it has been shown to work with real users in a given context.

Generality versus fidelity

Designing GENI to be generally programmable is potentially at odds with perfect fidelity. For example, a researcher could argue that to faithfully evaluate a new function or protocol, it's necessary to experiment with a commercial implementation or possibly with a function-specific hardware implementation.

In practice, however, such an implementation is likely to expose a limited interface rather than be generally programmable. This kind of device has perfect fidelity for a narrow set of experiments, but less value to the larger research community. On the other hand, an open source, software-based implementation of the same function or protocol might run on a general-purpose component that other experimenters can share, but without the performance or fidelity of the special-purpose implementation.

Clearly, it should be possible to make a merit-based case for the special-purpose component that benefits a narrow set of researchers, but it's generally expected that some amount of fidelity will be sacrificed to support a general-purpose facility that serves a wide range

of research. Nevertheless, more narrowly defined communities should be allowed to connect their special-purpose components to GENI and make them available to interested researchers.

Related to the problem of generality versus fidelity is simplicity: Researchers want to work at a low enough level of abstraction so that important system details aren't hidden. At the same time, however, they don't want to work at such a low level that they have to reinvent what to them are uninteresting layers of software just to create an environment that lets them address their specific research problem.

This is actually a unique opportunity for GENI—it should support multiple levels of abstraction and, over time, build up a suite of shared code for commonly used functions. Researchers should be able to work at whatever level of abstraction best matches their needs.

Technology development versus architectural design

We expect an ongoing tension between researchers wanting to use GENI to test and evaluate new networking technologies and those wanting to evaluate new architectural designs that, among other things, take the capabilities of new technologies into account. The former tend to focus on single components, while the latter must take a more comprehensive, end-to-end, perspective.

GENI's policies should favor architectural research, broadly defined, that exploits the fact that the facility spans a diverse collection of hardware resources. No individual technology is fully validated until it has been shown to work with real users in a given context, and we're interested in exploring alternative architectures that are capable of integrating a diverse set of technologies.

However, there is value to component developers being able to evaluate their technology in the context of end-to-end architectures and under the realistic workloads GENI is expected to generate. GENI should allow such technologies to be plugged into the

facility once they're mature enough to support GENI users, although we expect early-stage technology development—both hardware and software—to happen outside of GENI. (There is also likely to be a transition path whereby a new technology is made available to early adopters in a GENI subset.)

New components added to GENI should support the interfaces defined by the management framework, be sufficiently programmable to give researchers the flexibility they need, and, to the extent possible, be sharable by multiple slices.

Networking versus applications research

Because GENI is neutral about what level of the network researchers focus on, it doesn't draw sharp lines between low-level network protocols, high-level network services, and end-user applications. GENI should support any research that benefits from widespread deployment, diverse network technologies, and realistic network conditions.

The critical point of tension is that GENI is designed to support research in networking and distributed systems—as opposed to simply providing bandwidth to end users—yet it also benefits from traffic generated by real users. It will be necessary to evaluate the research value of traffic that a given slice generates to decide if allocating resources to that slice is warranted, rather than merely providing an infrastructure service to some user community.

A research group could justify the value of traffic it's carrying by

- making traffic traces available to other researchers,
- providing a novel network service whose efficacy must be evaluated, and
- offering to run as part of (on top of) a novel network architecture.

New communities that find value in some capability of GENI—or some innovative service deployed on

GENI—are free to augment the facility with enough capacity to carry their user traffic, independent of other research considerations.

Design studies versus measurement studies

GENI is being designed primarily to let researchers experiment with new network architectures and services not available today, and this purpose will largely determine how to prioritize among various design choices and resource allocation decisions.

Our hope and intention, however, is that the facility will also provide a new capability for monitoring the current Internet. We believe such dual use is possible because both capabilities require wide deployment, rich interconnection to the existing Internet, and heavy instrumentation. Using GENI as a platform to monitor the current Internet is a secondary goal that will also inform its design.

Deployment studies versus controlled experiments

We don't view the two primary GENI usage models as conflicting—a research group might naturally progress from a series of controlled experiments to a long-term deployment study—but there is an important difference in how the two models stress the facility. Both are related to security.

A controlled experiment attempts to eliminate all outside, uncontrolled influences from affecting the experiment. It also needs to keep the experiment from impacting the rest of the world. This requires strong containment mechanisms so that, for example, an experiment that measures the effectiveness of a new malware-prevention architecture can't escape onto the Internet. Because such a breach could have catastrophic effects, experiments will likely need to be reviewed to evaluate the risks.

In contrast, a deployment study necessarily involves an experimental service interacting with real users, including individuals trying to abuse the network in some way as well as

those trying to use the network to transport illegal content. GENI must be willing to carry such traffic; it can't be isolated for the sake of security.

Thus, like an ISP in today's Internet, GENI must be responsive to complaints when they are raised. This means it must include auditing mechanisms that let operators identify badly behaving experiments so that they can be quickly isolated or shut down. In general, it must be possible to rapidly bring the facility as a whole into a safe and controlled state.

GENI is a unique facility meant to support both controlled experiments and long-term deployment studies with new network architectures. Although its design will continue to evolve, we hope that these principles—and the need to weigh them carefully against one another—will help guide GENI's evolution in the coming years. ■

The GENI Planning Group consists of Larry Peterson, Princeton University (Chair); Tom Anderson, University of Washington; Dan Blumenthal, University of California, Santa Barbara; Dean Casey, Ngenet Research; David Clark, Massachusetts Institute of Technology; Deborah Estrin, University of California, Los Angeles; Joe Evans, University of Kansas; Dipankar Raychaudhuri, Rutgers University; Mike Reiter, Carnegie Mellon University; Jennifer Rexford, Princeton University; Scott Shenker, University of California, Berkeley; and John Wroclawski, University of Southern California/Information Sciences Institute. Visit www.geni.net/mail.php to join the GENI mailing list.

This article was largely motivated by questions asked at GENI Town Hall meetings. Fred Schneider and Paul Barford also contributed comments and suggestions.

Editor: Simon S.Y. Shim, Department of Computer Engineering, San Jose State University; sishim@email.sjsu.edu