The future of GENI research, operations and governance

Ilya Baldin
RENCI/UNC Chapel Hill
(ibaldin@renci.org)

Paul Ruth
RENCI/UNC Chapel Hill
(pruth@renci.org)

Jeff Chase
Duke University
(chase@cs.duke.edu)

Overview

The development of GENI as a federated, heterogeneous and deeply networked operational CI facility has come to fruition for the Computer Science community. We believe the main impact of these efforts may be best realized by transitioning this success to serve the needs of a broader set of communities in research, education, industry, and government. GENI has the potential to be more than just an existence proof of such a facility operated by its developers: like the early Internet, current and future GENI technologies may provide a starting point that alters the technological foundation of our society. Thus further efforts in GENI operations, research and governance must be designed to realize this potential. This white paper is written with this perspective in mind.

Research, operations and governance all have their roles to play in this process. Research activities should be aimed at making GENI more elastic, scalable and user friendly. The latter is particularly important when it comes to attracting non-CS users, who will provide the critical bridge to adopting GENI technologies in other settings. The operations plan should support continuity of existing facilities as well as growth of the GENI federation, and development of methods and best practices for managing federated, virtualized and deeply networked infrastructure at large scale. Given the nature of GENI, governance structure should give a place at the table for infrastructure operators, architect-developers and users of different types. It should facilitate discussions around architecture, APIs and implementations and help solidify the gains around points of consensus by shepherding those through a standards process.

Special consideration should be given in all three those activities to interfacing with the industry, as part of technology transfer. At the same time, in order to preserve its innovative spirit and ideas, GENI should avoid the burden of engaging in industry-driven standards processes or chasing commercial market drivers. Rather, development should focus on leveraging the unique capabilities that GENI provides as a cooperative geographically dispersed multi-provider infrastructure service operated in a decentralized manner for the public good, with deep interconnection to campuses and other research sites and international research networks. For example, the “GENI Of The Future” (™) should support secure interconnection of campuses with CI resources via stitched virtual networks with manageable off-by-default connectivity to the open Internet. It should support long-running elastic services that can interconnect and compose securely by mutual consent. The seeds of these capabilities are present in the GENI architecture, and should be nurtured and developed further.

GENI should aim to serve as a “uniter” to federate additional research resources on campuses and in government and commercial labs. A key step is to enable a flexible multi-federation structure that allows peering arrangements to grow organically, generalizing the current top-down trust structure. Operators of large CI installations (clusters, supercomputers, networks) have a critical role to play. By using GENI to federate, they can enhance resource sharing for their users and also increase the scale and diversity of resources available using GENI protocols and tools. From there it is a natural step to broadening their user communities and increasing the scale of resources available to GENI users. They must also provide support for virtual infrastructure operators, by partitioning resources as needed to support virtual provider demands. Virtual providers will then compete in features and APIs and ability to federate novel resources, thus creating greater competition and attracting new communities of users and thus helping shape the future deployments of large CI installations.

The small number of these large providers will exercise significant power in making deployment choices over the architectures of their facilities. We must balance this power against the broader needs of the community to support more generalized architectural choices and not allow these deployment decisions to dictate broad architectural and implementation decisions. This is another reason to support virtual providers, who can offer alternative and sometimes competing implementations, creating a rich marketplace of ideas.

The separation of concerns can maximize return on public investment by creating a more flexible, agile, and responsive infrastructure that can serve multiple purposes and adapt as needs change. It can unleash innovation of new services by establishing powerful platforms that can support experimentation and production deployments, limited only by the capabilities of the hardware. It properly focuses user communities on the services they use, rather than the facilities they run on, and (we hope) enables these services to evolve in tighter cooperation with the communities they serve. It properly places infrastructure providers in a common carrier role: users of a service are not locked into a particular hosting facility, and facility owners should not gain undue monopoly power from the services they host.

GENI should influence the architectures of these large deployments by, on the one hand proving by example the feasibility and utility of its architecture and implementations and on the other by attracting and influencing user communities who can shape these future deployments. This is particularly true of many computational science domains, which today are served by large purpose-built supercomputers or clusters. These communities have a strong influence over architectures and features of future machines and by attracting these users to GENI we can ensure their machines can be federated with each other using GENI technologies. This will serve the user communities by allowing them to have elastic computational resources tightly coupled with storage and network functions that more fluidly move data and computations among the federated infrastructure and offering them sophisticated authorization controls. It will also serve GENI by broadening the adoption of its technologies.

Finally, it's important to emphasize that GENI should become an incubator of multiple federations based on the developed technologies. While the continuing existence of GENI federation as an artifact is important from the point of view of providing existence proof and being a platform for further experimentation on federation technologies, the creation of multiple federations will allow to attract new user communities with different usage models and policies without forcing them under the same policy/resource umbrella. The CI providers will be able to choose which GENI-like federations to opt in and under what conditions. An ideal proof of GENI capabilities will have some of the providers participating in different federations with different usage policies and user communities simultaneously.

The following sections discuss in further detail aspects of research, operations and governance that are critical to the future of GENI.

Future GENI Research

After several years of operating GENI it is clear there are many unanswered questions and interesting topics of research that involve building the infrastructure itself. The GENI effort made some progress as a proof of concept, but it avoided the research challenges by focusing on building operational capability around expedient standards.

From our perspective it is important to emphasize the two types of research needed: research on novel uses of GENI, to spur further development of features and tools and the study of scaling properties of the different aspects of implementations, and research into GENI itself, which during the initial construction of GENI were largely neglected in favor of speeding up the construction of the facility. The latter should help (re)define and crystallize the details of APIs, policies, description and authorization mechanisms that generally fall under the rubric of 'architecture.

The topic of using GENI as a domain science instrument stands on its own and is critically important. It will provide the important bridge for new user communities to accept GENI approaches and find new uses for the capabilities, like tight integration of various types of resources with network functions, resource elasticity, repeatability and dynamic provisioning, flexible authorization controls. These is discussed in the following section. In this section we primarily discuss a few directions for research into GENI architecture, tools and policies.

Scalable and flexible resource orchestration is required by many different user communities. Only the most advanced of users typically want to deal with explicitly constructing their own topologies, picking deployment sites and so on. More capable approaches are needed to support sophisticated partitioning and stitching of requests based on generally-specified user preferences for e.g. latencies, resource types, complexity of the topology involved. This requires research and development of tools, algorithms and aggregate APIs.

Of special mention is the mismatch between queue-oriented APIs that most domain scientists are used to and the (mostly) immediate scheduling supported in GENI. The meshing of these

two paradigms in a single architecture is difficult, especially without losing the deep integration with network function, usually requiring on-line control, yet carries great benefits in making GENI attractive to a wider set of user communities.

Scaling accounting, authorization mechanisms and policies to large numbers of users, facilities and virtual providers. In general the study of scaling of the different aspects of GENI architecture remains an empirical exercise. A stricter approach to these issues and a disciplined evaluation is required to make judgements about the choices made so far. Support for resource accounting stands out as a piece that is completely missing from the current architecture, yet without it the integration into GENI of substrates such as public clouds remains a one-off capability exercise. The lack of accounting will also serve as a barrier to adoption by many user communities used to accounting the aggregate use of the resources by projects or users.

Integration of exchange of consideration for services. Once proper accounting mechanisms are in place, the next natural step is providing the ability for users to exchange 'consideration' with resource and service providers. This consideration can take many forms, from tokens, to money, to various information about the users. This problem area has a strong intersection with authorization policies and has also been neglected.

Deep(er) network integration. The intersection of IaaS (Infrastructure-as-a-Service), NFV (Network Function Virtualization) and SDN provides a fertile environment for research on the intersection of policies governing federated resource allocation and control over networked virtual resources (an exemplar of which can be some types of SDX) and further integration with GENI architecture. Also in this rubric is integration with non-GENI campus infrastructure, not in the form of federation, but rather by establishing dynamically controlled interconnects between slices and campus networks.

User-facing interfaces. Relatively little attention has been paid to usability studies in GENI. By most accounts, while GENI tools are slowly improving in terms of their ease of use, the process of improvement is still empirical, without much foundation in e.g. human interface design.

Domain Science Research

The currently deployed NSFCloud mid-scale infrastructure is a significant step forward because it is not just a testbed but offers sufficient scale and is composed of hardware capable of producing scientifically-significant results for domain sciences. Modern domain science computational models and simulations often require amounts and types of resources that were not available on the initial GENI testbeds. These applications are known to scale to thousands (sometimes tens of thousands) of compute cores often connected with high-bandwidth low-latency network fabrics (e.g. Infiniband). The mid-scale NSFCloud have enough computational and network power to support all but the largest domain science applications. This is a significant step forward, however the deployment of GENI mid-scale infrastructures must be thought of as progress toward an end but not as the goal itself. The ultimate goal is for

GENI technologies to be considered a requirement in all new leadership-class computational infrastures.

In order to reach this goal the GENI community must use the opportunity created by the new mid-scale infrastructures to demonstrate the effectiveness of GENI technologies for domain science. Although the primary goal of this effort is to expose domain scientists to the novel capabilities developed by GENI, the ultimate success of this project will be measured by the willingness of computational science communities to adopt these technologies. This means that the focus of these efforts must be on appealing to domain scientists.

Make it free. Currently GENI and NSFCloud resources are available primarily to computer science researchers. Although domain science applications are allowed to run on these resources it is always with the assumption that we are not providing free cycles and that there needs to be a compute science reason for running the application. One of the best ways to attract computational domain scientists is to offer free “cycles”. Once attracted by free “cycles”, they will experiment with the new novel capabilities which will convert them to new computational models enabled by GENI technologies.

Make it easy. Several domain scientists have been introduced to GENI and have worked with computer scientists to run their applications on GENI resources. The primary complaint that domain scientist have when using GENI is that it is still too difficult to use. GENI needs an interface for domain scientists that is similar to what they are currently using (HTCondor, PBS, SLURM, etc.), yet still provides access to its deep network integration and resource scheduling features. Another possibility is to enable a simple integration of GENI with domain science portals such that GENI become a back-end infrastructure to existing portals.

Make it fast. Performance and scalability in GENI have not attracted the focus required to be a viable alternative for domain science applications. In order to attract domain scientists GENI must enable the new NSFCloud infrastructures to provide performance and scalability that is comparable or exceeds that of existing HPC resources. This includes avoiding oversubscription of compute resources and providing QoS guarantees on network bandwidth and latency.

Make it better. After the domain scientists have been attracted to GENI because it is free, easy, and fast, they should be shown how novel GENI technologies enhance their ability to produce real scientific results. Technologies that should be emphasised to domain scientists include: custom elastic compute environments, bandwidth provisioned long reach circuits, and integration with private data repositories and computational facilities. These novel GENI technologies augment the scientists’ traditional compute environments with features that enable efficiently moving large amounts of data, move computation to optimal physical resources that are then used in familiar ways.

Operations

GENI operations should reflect the federated architecture and enable the research described above. This means avoiding consolidation of operations in one organization, supporting provider autonomy and their participation in multiple GENI-like federations, and developing tools and procedures to enable those. Participating providers should be empowered to approve or deny any request according to their policies, using the tools developed in GENI.

The critical aspect here is getting campus providers more engaged in the process, such that operating federated infrastructures becomes an ingrained part of the campus IT operations. This will likely require extensive outreach and training with campus personnel. As with most other aspects of GENI discussed in this document, users are the forcing function in this process.

GENI operations structure should also support virtual providers, who will use portions of resources of other providers to create their own unique offerings within the federation, distinguished by unique programmability features and targeted at specific communities of users. This will attract new users and help serve the research needs of GENI, by allowing to experiment with tools, architectures and APIs, while at the same time providing sufficient scale to be tested with realistic workloads and quantities of users.

Governance

It is important to get the governance structure right, especially when it comes to balancing the needs of user communities against the power of the resource providers. At the same time it is important to maintain the interest of developers and architects by leaving sufficient freedom to deploy novel solutions and avoiding solidifying too many decisions in an attempt to operationalize the infrastructure.

Governance structure should give power to new groups, joined via outreach and campus engagement and should also support interfacing with multiple federations built on similar principles and technologies, but for different user communities. These communities should not merely be consumers of GENI technologies, but have a direct stake in GENI success and a voice when it comes to various architectural and operational decisions.

Finally, the governance structure should be reflective of GENI relationship with industry and standards organizations, by providing them a seat at the table, preferably on a reciprocal basis.