

GENI Experimenter and Educator Community Engagement

January 2016

Joe Mambretti, John Geske

Introduction

This section of the work group report consists of background materials on community engagement provided by the GENI Program Office (GPO) and materials generated by discussions on community engagement at the planning workshop.

Summary

GENI Experimenter and Educator Community Engagement workshop breakout section was provided with a comprehensive overview of current experimenter and educator community engagement activities by Vic Thomas, Niky Riga of the GENI Program Office. Subsequently, the group discussed the various GENI communities, methods of engagement, and potentials for enhanced engagement.

Recommendations:

Although no formal recommendations emerged from these discussions, there was general consensus on several issues. One was that it would be useful to further define the characteristics of the constituent GENI communities. Another closely related observation was that the requirements for each community should be made more explicit. Another was that addressing the needs of these constituencies should be placed in a priority context. While it is clear that the experimenters and educators are the primary constituencies, the rank order of the other communities was not well determined. Also, there was consensus on the need for enhanced processes for engaging all communities. While existing initiatives in this area are of very high quality, there are opportunities for improvement in these processes, particularly with regard to addressing the specific requirements for each.

GENI Communities

The GENI project engages with diverse communities including: (1) end users such as experimenters and educators who use GENI for research and teaching (the primary constituency for the GENI project), (2) campus academic departments that support GENI activities (3) IT departments that host and manage GENI resources, (4) academic research departments that host and manage GENI resources, (5) operators of national, regional and local R&E networks that provide wide-area network, (6) industry partners, (7) national and international cyber-infrastructures that federate with GENI, and (8) a variety of other interested partners, including those who are interested in using GENI type infrastructure for next generation services and application development.

GENI Transition

During the term of the GENI project, it has transitioned from activities related to designing and implementing the distributed research facility (the Global Environment for Network Innovations) To providing the services based on that platform to its constituent communities, the current focus of engagement activities. The workshop discussions centered on which current activities should be revised or discarded, and what new activities should be started. One observation was that GENI must continue to evolve. However, it is important to plan on the direction of that evolution. Specifically, that direction must be determined by constituency groups. Therefore, it is important to carefully define constituency groups, the GENI communities, and to reduce barriers for those communities if they would like to use GENI.

The three primary objectives of end-user community engagement activities are:

1. Growing the community of GENI users by informing researchers and educators about its capabilities and help them get started using it.
2. Supporting existing users so GENI is more useful for their research and teaching,
3. Fostering a sense of community to encourage users to help one another by sharing their GENI knowledge and experiences.

A major activity supporting growing the GENI community is the GEC series of events, which have been the largest community building events drawing from the experimenter, educator, network and campus operations and software developer communities. The GEC host, typically a faculty member at a university, handles all local arrangements, with active support from the GPO. These events have been supplemented by other community building events, including a GENI Network Innovators Community Events (NICE), CNERT, a community run workshop with a largely volunteer organizing committee., competitions for experimenters, and specialized education workshops.

1. Other existing Activities Related To Growing the GENI User Community Are:

1. Tutorials at GENI Engineering Conferences (GECs). Until recently, the GECs were the largest experimenter training venues with tutorials ranging from "Getting Started with GENI" to topics such as running "Hadoop in a Slice" and using tools such as geni-lib. Tutorials are proposed and led by members of the community and GPO staff.
2. GENI Regional Workshops. These workshops are intended to replicate the success of the tutorials at the GECs but at a smaller scale, lower cost and reduced organizational overhead. They are:

Two such workshops have been held, each hosted by a faculty member at a university who takes care of local arrangements such as tutorial rooms, catering and publicity. For both workshops, the GPO set the workshop agenda, arranged for instructors and speakers and led some of the tutorials. A few travel grants were given to attendees from outside the region and to instructors and speakers. In 2016 and 2017, a community member funded by the GPO will lead the organization of the workshops.

3. Summer and winter camps. Participants of these week-long camps learn start with in-

introductory GENI tutorials and move on to more advanced tutorials. Camps typically have a theme (e.g. software defined networking, systematic experimentation, etc.) that are covered in the advanced tutorials. Participants form teams early in the week, define and complete a GENI-based team project during the week and present results at the end of the camp. Teams are encouraged to continue to work together even after the camp and publish their results. Five such camps have been held to date. The organization of the 2016 and 2017 camps will continue to be led by GPO funded community members.

4. Tutorials at conferences and workshops. Over the years community members and GPO staff have organized tutorials at conferences and workshops. Community member led tutorials include those at SIGMETRICS, SIGCSE, ENC, Cloud Security Curriculum Development Workshop and Workshop on Large Scale Experimental Research Environments. GPO led tutorials include those at ICDCS, NSDI, TridentCom, IC2E and IEEE NFV-SDN. In addition to tutorials, GENI posters and demos at conferences have been used to spread the word about GENI.

5. Talks at conferences, workshops and university seminars. Members of the GENI community and the GPO have been invited to speak at numerous conferences, workshops, seminars and classes. Depending on the venue, the talk may include a demo of an experiment on GENI and even a simple hands-on exercise. In addition to explicit GENI-centric talks, overviews of GENI are often incorporated into presentations on other research and facility topics.

6. User Support: Online tutorials and documentation, information on special provisioning, reviews of potentially disruptive experiments, course modules for educators, support mailing lists, and support mailing lists.

Major Issues With Experimenter Communities

During the discussion, the question was asked, given these activities, why hasn't the facility attracted more researchers? Specifically, it was noted that many "top-notch" networking researchers do not use GENI now and have no plans to use GENI. Another question was how can these researchers be attracted to the facility. In response, one suggestion was that more community engagement could be planned at events frequently attended by network and distributed systems researchers. It was noted that most such researchers attend key annual events that are important to their activities. Another response was that although most researchers can not afford to set up their own testbeds, some top level researchers are able to obtain funding to set up their own labs, although they may require a mid-scale cyberinfrastructure for experiments that scale. Some researchers are addressing highly specialized topics and require infrastructure that is different from GENI. Some simply do not want to use shared testbeds. The point was made that there is a fundamental, qualitative difference between obtaining funding for a community for their own purposes as opposed to obtaining funding for a community that will build resources for another community. This was a key point, it was noted that many active members of the GENI are taking time from their own individual activities to work in partnership with others to build the GENI facility and the GENI community.

To address the issue of attracting additional researchers, a suggestion was a mechanism for establishing ways to attract additional researchers to view and become involved in interesting research, e.g., having facilities and software that others can build on. The dynamism of federation architecture is a major asset for this. Being able to integrate resources in the infrastructure in a more flexible way is key. Isolation is another nice benefit of GENI (in addition to federation). A shared tested makes it easier to share artifacts. Also, GENI could make it easy to access and use novel resources which would otherwise be difficult to obtain. It was noted that the reasons GENI has been useful/not useful for some in the networking research community is complicated., and perhaps a workshop should be organized on this topic.

One observation was that current research on GENI is being conducted by those who know about GENI. However, others are conducting on other types of facilities, for example, some researchers who are conducting reproducible network research (such as a group at Stanford that has virtual appliances for repeatable experiments including through developing "clickable papers") . GENI would have been a natural place to do this but it is not being used. One question asked was: Is GENI is too difficult to use? This question was followed by a discussion of the evolution of GENI interfaces from those that were fairly difficult to those that are more straightforward but that could perhaps be further improved, especially with regard to resource certification processes. Networking researchers need a large-scale testbed, with an easy on-ramp, an environment that is far more flexible than what they can build on their own.

Other Communities

A question was asked about whether GENI has spend too much time engaging core communities while ignoring other communities, such as more general campus communities. There may be tremendous potential in extending GENI sliceability to other resources so campuses can share innovations on their campuses. To some degree, advanced networking capabilities is being supported on campuses through the NSF's Campus Cyberinfrastructure programs. GENI may provide a mechanism to enable more distribution of innovations.

Domain Science Communities

A question was asked about the utility of GENI for domain science communities. GENI was initially developed to focused on networking/distributed systems research. It is not deployed where domain science researchers can easily access it, although there are campuses that have NSF CC* grants which mitigate barriers, places with science DMZs, significant bandwidth – eg, minimally 10 Gbps paths, large scale storage, and related resources). Some science domain communities have used GENI for proof-of-concept projects, experimental network techniques related to their science research, which they would later like to migrate to their production facilities. GENI is not a testbed for production domain science. It is a platform for "kicking the tires" of concepts/techniques. It is a good place to develop new techniques and methods. GENI can be used to show that the concept is viable done and then a production infrastructure that is suitable for the application can be built or the techniques/software/tools developed on GENI could be transitioned to existing production domain science infrastructure). Other domain science communities have expressed interest in this type of prototyping activity. This is an appropriate use of GENI and should be encouraged. With some minor changes, it would be

possible to more formally serve the domain science communities much better. This community needs (and understands) mid-scale distributed infrastructures. Although some domain scientists such as physicists already have much of the tools and infrastructure they need, other scientists such as bioinformatics don't have it---e.g., producing more data than they can process or move. Currently GENI is only one of many cyberinfrastructure facilities on campuses. No overall "glue," or umbrella orchestrator exists today. A question was asked about the potential for GENI being this mechanism.

Education Constituencies

The importance of the educational use of GENI was also stressed. Educational activities are a good use case for those who are not part of the core community (those not building GENI, not part of the current GENI community) but who may wish to be introduced to GENI. Some educational modules currently run on GENI and are being used by individuals do not know anything about GENI. The Mars Game is a STEM game for high-school education. They use GENI because they need a low-latency high-bandwidth. A recent event showed students Seattle, Emulab and GENI. GENI was by far the easiest to use for their projects, an overwhelming number of students preferred to use GENI.

Industry Partners

The GENI project has had several major industry partners, including Ciena Research Labs, which has been using GENI for proof of concept development projects, HP Research, Dell, and Cisco. GENI is a natural platform for virtual networking innovation -- virtualization is a major industry activity -- on GENI with 5 clicks of a mouse you can set up a virtual network---that is tremendous potential. Speed of processes (getting things done) is essential for industry. The question was asked: Should GENI do more than for industry community engagement?

Other Potential Community Partners

Another community that has been engaged with GENI is the US Ignite initiative. Questions were asked about the potential of establishing closer partnerships with distributed systems initiatives, including NSFCloud testbeds, Chameleon and CloudLab, both of which federate with GENI, international testbed communities, R&E networks, including regional networks, federal agency networks, especially ESnet, which has established a large scale SDN testbed, NASA networks, the NIH network. The ExoGENI project has been working with the national ESnet 100 Gbps testbed. Another question was asked about the potential to leverage regional R&E network meetings.

GENI and General Cyberinfrastructure

The question of whether testbeds can co-exist with general cyberinfrastructure was asked. It was noted that these two types of infrastructure have divergent objectives. Production cyberinfrastructure is hardened, locked down, application mission oriented, and is somewhat inflexible compared to testbeds, which are designed to allow for risk taking, including activities that may be highly disruptive to parts of the infrastructure. A suggestion was that perhaps there

should be two GENIs – a testbed GENI and a production GENI – paralleling the evolution of the NSFnet to the current Internet. A suggestion was made that perhaps parts of the infrastructure could transition from test mode to production mode – and back. Industry needs this type of flexible infrastructure to be able to implement high risk technologies without damaging production networks. A question was discussed about whether GENI was competing with commercial infrastructures. The consensus response was that if activities could be conducted on the general Internet and commercial clouds they should be. Such activities should not be supported by GENI.

Large Scale Cloud Providers

A question was raised about engaging the major cloud providers, especially those doing innovative networking with SDN. Some of architecture for their innovations are being placed in the public domain, for example, one that is removing UPSes in data centers by placing consumer batteries in power cords. Large cloud providers are putting their innovations in public domain so others can sell it to them. GENI can take advantage of this trend. Some of these cloud providers are moving fairly quickly – much faster than standards groups

Procedures for Communities

There was consensus agreement that for each of these communities that have been identified, the GENI value proposition would have to be identified, and a set of activities, processes, and communications need to be established to attract more of them. There was also consensus that a need existed to prioritize the communities and activities related to those communities.