GENI Futures - Community Perspective of the Current State of GENI

Kuang-Ching Wang[1] and Jim Bottum[2]
[1]Department of Electrical and Computer Engineering, Clemson University
[2]Chief Information Officer and Vice Provost, Clemson University

Introduction

Over the course of the project, GENI has "touched" a large fraction of the US academic community via multiple fronts. To **computer science researchers and educators**, GENI is a collection of deeply programmable compute and networking (wired and wireless) testbed infrastructure with a national and international footprint (through federation) and a testbed control and instrumentation software suite enabling a single portal and API for invocation of computer science experiments. To **domain science researchers and educators**, particularly those in computation and data intensive disciplines, GENI-proven software and hardware components (software defined networks, GENI racks) have either been incorporated into production cyberinfrastructure (via NSF CC-*IE programs and regional/national software defined networks) or stimulated domain-specific workflow integration efforts (e.g., NSF CIF21 DIBBS: Tripal Gateway, NSF Cloud: CloudLab). To **university IT and research computing support organizations**, the impact is broad and is driven by all the above, with researchers needing their support for operating and accessing hardware and software over the campus cyberinfrastructure for both experimentation and production uses. A reality check on most campuses leveraging GENI resources today, professional IT staff is involved in their operation, and the same staff is often also involved in CC-*IE and other sponsored research projects to deploy, and in some cases create, novel GENI-inspired services into campus. Essentially, to the diverse range of communities on campus, "GENI" can mean a wide range of infrastructure and services, which all require support by production IT staff in operation, in development, and in use.

Clemson Involvement with GENI

Clemson University has participated in GENI since 2009 through engineering subcontracts on Campus OpenFlow Trials, Programmable Radio, GENI WiMAX, experiment subcontract on GENI Cinema, and EAGER grants for experiments on security and p2p. Clemson also has a CC-NIE award building a production SDN as a distributed science DMZ across 20 buildings. Clemson is a member of the NSF CloudLab project team. Clemson also has a CISE CRI project that builds a compute cluster for data intensive research leveraging SDN solution derived from GENI.

At Clemson, IT support for GENI and relevant activities has evolved and matured over time. In 2009 as our first GENI subcontract launched, the faculty PI requested a set of campus VLANs to be provisioned across three buildings where experimental switches are hosted. In 2012, as Clemson's CC-NIE project launched, IT planned and deployed the 20-building network following production workflow and standards. In 2014, as our ACI-REF project launched, a dedicated FTE (50%) was assigned to facilitate a wide range of research and education uses of advanced networking accessing the full spectrum of advanced cyberinfrastructure resources, including the GENI resources.

Current Needs

Much like the initial efforts to build major scientific and cyberinfrastructure-related communities, GENI is seeking to build a community-based organization that can sustain itself and leverage existing and future investments by NSF through developing a core program at the national level that prompts action at the campus level. This can be done in a variety of ways, as arguably a large number of campuses have a need for this type of community resource. With GENI being a NSF CISE infrastructure, a model is needed for campus IT organizations to sustainably support this infrastructure and couple with it long-term technology planning.

Sustainability – Possible Models

As of today, GENI has a distributed infrastructure across university campuses, municipals (US Ignite), and regional and national networks, a central clearinghouse based access mechanism, a central user support and monitoring team (GPO and GMOC), and a distributed infrastructure support structure (campus IT). It is clear that the future GENI governance, operation, and support framework must account for both the centralized and distributed (on campus) functions.

Maintaining and growing sustainable distributed infrastructure is not a foreign concept to universities. Take the NSFNet [1] example, in 1986, the National Science Foundation invested in a national 56 Kbps backbone network, initially connecting the five main supercomputer centers across the US[1]. In addition, NSF investment in regional networks over the next several years was the catalyst for universities to build out of campus networks. The regional networks connected to the NSFnet Backbone and the price of admission for campuses was an integrated campus LAN. Fueled by increasing demand on campuses by researchers to connect to these national resources located at supercomputing centers, this initiative prompted campus-level action and the need for a sustainable model for the future.

GENI, in our opinion, is both similar to and different from NSFNet. The similarity is obvious – GENI is a NSF seeded infrastructure and GENI provides/connects compute resources of value to campus researchers. The differences are, on the other hand, more subtle yet significant – the GENI infrastructure is distributed and embedded in campuses; the GENI infrastructure requires on-campus support to continue operation; the use of GENI has proven to benefit tremendously to have "human" guidance to achieve the best outcome given the breadth, depth, and novel nature of the technology, as well as the diversity of the user communities. Summing it all up – what GENI needs most to achieve sustainable and maximal impact, is campus investment in the distributed operation and support. Building a community framework for a common strategic mission while leveraging collective wisdom on the breadth of technology and campus issues is not an easy task. Fortunately, the recent NSF-funded ACI-REF project shed some light and initial success towards building such a university-driven distributed framework for community governance, support, and collective advancement of advanced cyberinfrastructure and its use.

The ACI-REF Project/Consortium

Funded in 2014, the ACI-REF Project (NSF ACI-1341935) and its associated group of institutions have built and operated a coordinated network of Advanced

Cyberinfrastructure Research and Education Facilitators (ACI-REFs) whose mission is to leverage existing national and campus resources and "make a difference" in supporting local research communities. The group is dedicated to the adoption of models and strategies to advance our nation's research & scholarly achievements through the transformation of campus computational capabilities and enhanced coupling to the national infrastructure. The project is based in part upon Clemson University's experience with funding an initial "facilitator" on its campus in 2010, and a small number of other campuses began to centrally-fund these positions soon after. In the wake of these experiences and those with the ACI-REF project, more campuses have begun to fund 'facilitators' within their own IT organizations as the researcher need grew along with a demand for facilitator expertise. The 2014 NSF-funded ACI-REF Pilot proposal (Phase I) was an initial approach to a national-scale program that helped campuses get started in the area of facilitation and served as a mechanism to begin building this distributed community.

As the initial pilot phase of the Advanced Cyberinfrastructure Research and Educational Facilitation (ACI-REF) program nears the two-year mark, the project has attracted national attention, mainly from campuses requesting information on how they can get involved with the project's activities, and also from inclusion in several campus CI plans in recent NSF proposals where part of their local strategy to evolve campus cyberinfrastructure is to join the ACI-REF effort. In response to this growing community interest, the project team is forming the ACI-REF Consortium, which will become a persistent, self-sustaining enterprise based upon the model developed under the NSF-funded ACI-REF project to which interested campuses can 'plug in' to the many ACI-REF activities and base of knowledge and expertise. This consortium will be self-sustaining after a bootstrap period through a membership structure that provides a sustainable financial model, and through a governance structure that allows for shared community governance through campus representation on the ACI-REF Council, which will be loosely modeled on the governance model of the Open Science Grid. The ACI-REF Consortium and ACI-REF Project will be separate but linked on parallel tracks.

GENI can learn from the ACI-REF experience and leverage it as a model to build community support. The key driver for much of the success of the project, and ultimately the community-based consortium, will be the CIO/Institutional-Level commitment, as this ensures that the organization's mission and activities are being viewed at the highest levels. This also ensures that the top-level responsibility for funding of a membership model, any technical refreshes or new deployments, and other budget-related topics.

References

1. "The Launch of NSFNet." https://www.nsf.gov/about/history/nsf0050/internet/launch.htm.