V. K. Cody Bumgardner
University of Kentucky

- What GENI capabilities are most important?

I think the greatest benefit of GENI is the self-service model (and of course price) used to provide resources. Having worked in industry, healthcare, and now academics/research and I can say the velocity of infrastructure acquisition and operation, regardless of funding, is rather slow in research by comparison. Unfortunately, much research that could benefit from advanced cyberinfrastructure is conducted from the laptop, drastically limiting (in both capacity and mindset) the scope of problems. The self-service aspects of GENI should be maintained regardless of direction.

- What activities should GENI continue, expand, or wind down?

The first GEC was held in 2007, only a few months after the first iPhone was launched. How has the use of cyberinfrastructure changed in research, on campuses, or in society in general over the last 8 years? In a mobile-first cloud-enabled world the network could not be more important, yet the lower-level (L1-3) aspects of networking are out of reach for the majority of users. I would argue that the majority of applied network development is taking place on the higher-levels of the network stack.

One could continue to focus on providing an environment to support lower-level protocols, or start focusing on application-based networking experimentation further up the stack (L3+). Perhaps both can be done simultaneously if proper isolation is put in place. I contend that if services are to be provided for application-based network experimentation, the lower-levels of the network and related computational resources would have to be operated as a production system. What we think of now as GENI experiments, would turn in to GENI projects in this paradigm. This would mean a different operational profile than what is currently in place.

I would continue to provide the end-user aspect of computational (VM) and network (SDN) resources for experimentation and/or projects. Additionally, I would expand GENI in the following areas:

1) In addition to VMs, container-based (Docker) technologies should be provided for application-based network experimentation/projects. Container-based isolation is based on system-level virtualization and does not have the overhead of traditional machine-level virtualization. For instance, a machine that supports run 25-50 VMs could support 750-1000 containers. In addition, containers provide methods to deploy application-layer dependencies and can be thought of more as an application delivery and isolation framework than traditional VMs. Traditionally, the systems-level isolation of containers limited their use in network-focused applications. However, we and many others are working with Docker to integrate OpenFlow and containers, which would use to use them in GENI for application-level projects.

2) GENI was well ahead of its time in terms of distributed infrastructure orchestration in 2007. However, there now exist several infrastructure and application-based orchestration engines and description languages. One of the most common description languages is the AWS CloudFormation (https://aws.amazon.com/cloudformation/) template format used by both AWS and the OpenStack Heat (https://wiki.openstack.org/wiki/Heat) project. Another orchestration engine is the one provided by Ubuntu called Juju (https://jujucharms.com/).

One could provide application-based orchestration in GENI in many ways including:

2.1) Native operation: develop procedures and RSpec(s) to deploy VM- and/or container-based application-level orchestration engines. Technically, I think this is what Rob is doing with CloudLab, with his full OpenStack install. There are perhaps more light-weight (with limitations) ways to do this with containers.

2.2) API translation: develop an API translation layer between GENI and one or more common infrastructure APIs. For instance, if a subset of AWS or OpenStack APIs could be translated into something usable by GENI, so existing templates and orchestration engines could be used. I expect this as well could be done easier with containers (I.e. https://docs.docker.com/machine/).

- How should GENI be governed and sustained?

This depends on the focus of operations. If we focus on application-layer experimentation or even production services, the lower layers of the infrastructure must be treaded like a production environment. Much like I2 provides network services, you would need to tightly control rack standardization, updates, and forced decommissioning (or service tier isolation). Additional support would need to be provided to sustain 24/7 operations. Monitoring (obviously I am biased here) of both infrastructure, infrastructure, and experiment/project instrumentation should be mandatory, at least for production tiers. In this mode of operation GENI could look much like I2 or other similar organizations, with membership, elected board of directors, and a full-time executive director.

What campus, city, or national resource does GENI provide that could be consider a cost avoidance? I think that in the current case, GENI provides new resources, but does not displace old resources paid by users. However, if the lower-levels of the infrastructure were supported from a centralized production prospective, other workloads could be trusted to GENI services. Sustainability will not be a problem if GENI can provide both new resources, but also displace existing more costly resources. Providing a lower cost Amazon-like (self-service) experience for users would be very valuable and I expect sustaining.

I don't think costing based on time-based or individual resource-based usage (Amazon approach) will work for our experimentation, but it might for production projects. I expect the best way to start recovery is to charge an operating cost per rack or per resources within the rack. Racks without paid support would not be consider production, but could remain part of the network. An allocation committee, policy and accounting system (already part of monitoring) would have to be developed to account for production resource distributions.

- How can the GENI experience inform better research cyberinfrastructure?

If feel the following are key aspects:

1) Account/System Federation: time saved dealing with common access issues can translate into productive time.

2) User self-service: putting appropriate resources in the hands of people that can do the work is one of the most important things we can do. We must also change with technology and focus on areas that provide the most promise. If you provide infrastructure resources to the broader scientific community, those that can use infrastructure will, those that can't won't. If you provide a distributed application-based project platform, then a broader group of people can participate. In 2007 all researchers had phones, but how many had a smart phone? How many have one now? Perhaps instead of providing dial-tone we should be providing smarter services.

3) Sustainability: Would (as many) people pay $200k for Internet2 port if it was in support (from campus prospective) of research alone? I2 services displace some services that could be provided by more costly commercial providers. Future infrastructure projects must aim to provide core production services that can displace more costly services currently being provided in other ways. Through early adoption of disruptive technologies like OpenFlow, OpenStack, and containers, we can often provide services as effectively as our commercial counterparts at lower cost. In addition, greater integration into campus/city IT departments, as suggested by NSF infrastructure programs is also needed. The proof of this will be in service adoption and the displacement of alternatives, which could produced a sustainable system.