Zhi-Li Zhang
University of Minnesota

- What GENI capabilities are most important?

-- GENI's capability of simulating big topologies and network performance when nodes are in different geographic locations, and the flexibility to change protocols according to experimenters' needs.

-- The ability of providing experimenters "machines" with layer 2 connections, on top of which anything could be altered

-- Experimenter ability to select resources without restriction from all the available GENI Aggregate Managers

-- Possibility to renew the expiration of reserved GENI resources without any restriction

-- To assign public IPs to hosts in GENI, to bound VMs to specifics PCs, select the type of VMs, installation of custom script/software and the possibility to create VM images (e.g. InstaGENI images) to reproduce experimenters in the future.

-- Possibility to connect resources at different GENI Aggregate Managers using stitching.

-- Continue to evolve the GENI architecture and softwarte/hardware, e.g., to support the future P4 capability of hardware/software switches

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

-- GENI Poster/Demo activities should perhaps continue. It allows experimenters to learn about how the research community is working with GENI. Furthermore, it allows experimenters to share "best practices" and develop collaboration with others GENI experimenters.

-- Continue to expand the GENI infrastructure and its coverage by supporting additional GENI clusters at various university campuses or even non-profit/for profit institutions/ISPs

- How should GENI be governed and sustained?

-- one model is to use a similar governance model as PlanetLab as a centralized management and support (e.g., via BBN), and distributed administration by various campuses which contribute resources

-- alternatively, an Internet2 type model can be used

- How can the GENI experience inform better research cyberinfrastructure?

The main limitations of GENI are the following:

a) once the resources are reserved, it's not possible to dynamically change the experiment topology (e.g. add or remove resources)

b) it takes very long time to obtain resources for relatively large topologies expanding more than one GENI Aggregate Manager (e.g. more than 15 nodes). Furthermore, it is very difficult to obtain resources for very large and distributed network topology (e.g. using more than 3 GENI aggregate managers across the US).

c) Perhaps more importantly, the current GENI infrastructure is still too closely tied to the existing Ethernet/IP/TCP based network architecture, with its reliance on existing generations of openflow/OVS switches or WiMax, etc. Most research projects supported on GENI focus on "higher-layer" innovations using GENI; few are on network innovations as the intended goal of GENI. Such "low layers" innovations are severely limited due to the reliance of existing hardware/software systems (e.g., openflow switches/OVS, WiMax, etc.)

Future research cyberinfrastructure should address the above limitations of the present GENI test-bed. They should allow experimenters to increase/decrease the resources in their experiment on-demand, as well as, decrease the time they spend to obtain the resources for their experiments. Furthermore, with the emergence of P4-capable switches and "white box" switches with programmable silicon merchants, hopefully GENI can be further evolved to incorporate such hardware/software capabilities to truly enable and inform cyberinfrastructure research and development of better cyberinfrastructure.