

## GENI Racks

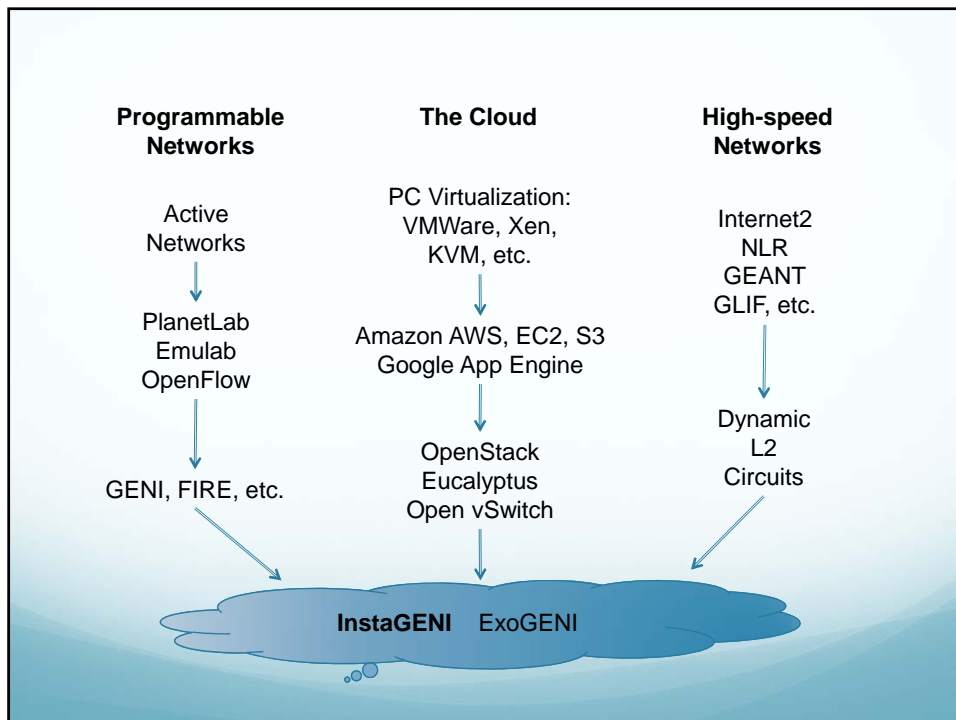
### InstaGENI Design and Deployment, Changes to PlanetLab Networking Model

Andy Bavier  
Princeton Univ.

Joe Mambretti  
Northwestern Univ.

Rick McGeer  
HP Labs

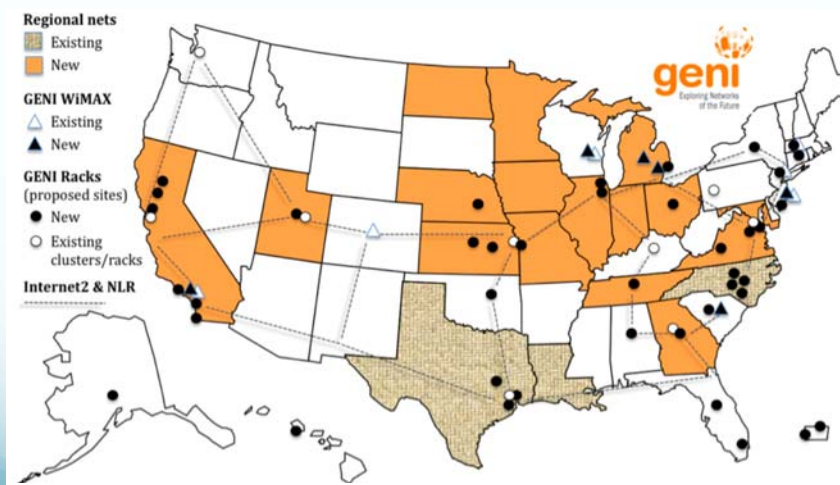
Rob Ricci  
Univ. of Utah



## GENI: Key Ideas

- *Sliceable* compute and network resources
  - Per-experiment sandbox
  - Rich programming environment within sandbox
- Deeply programmable network stack
- High-speed connectivity via R&E networks
- A possible model for future cyberinfrastructure

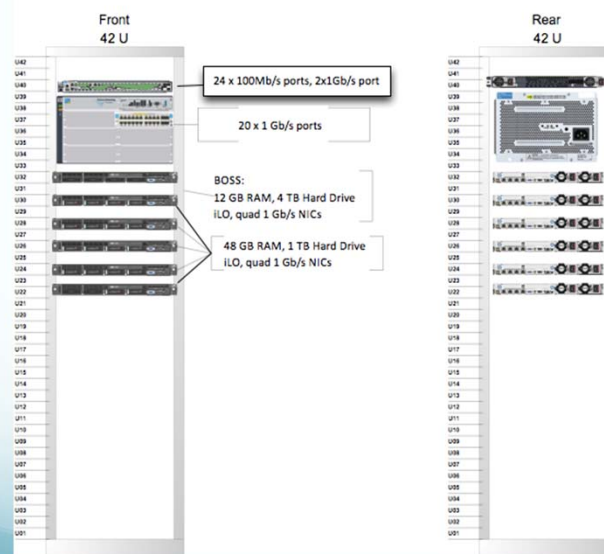
## GENI Deployment



## InstaGENI Goals

- Quickly deploy systems that already work
  - ProtoGENI (Emulab)
  - PlanetLab
  - OpenFlow, FlowVisor
- Lots of points of presence
  - 8 sites in 2012
  - 24 sites in 2013
  - Expandable COTS hardware configuration

## InstaGENI Rack

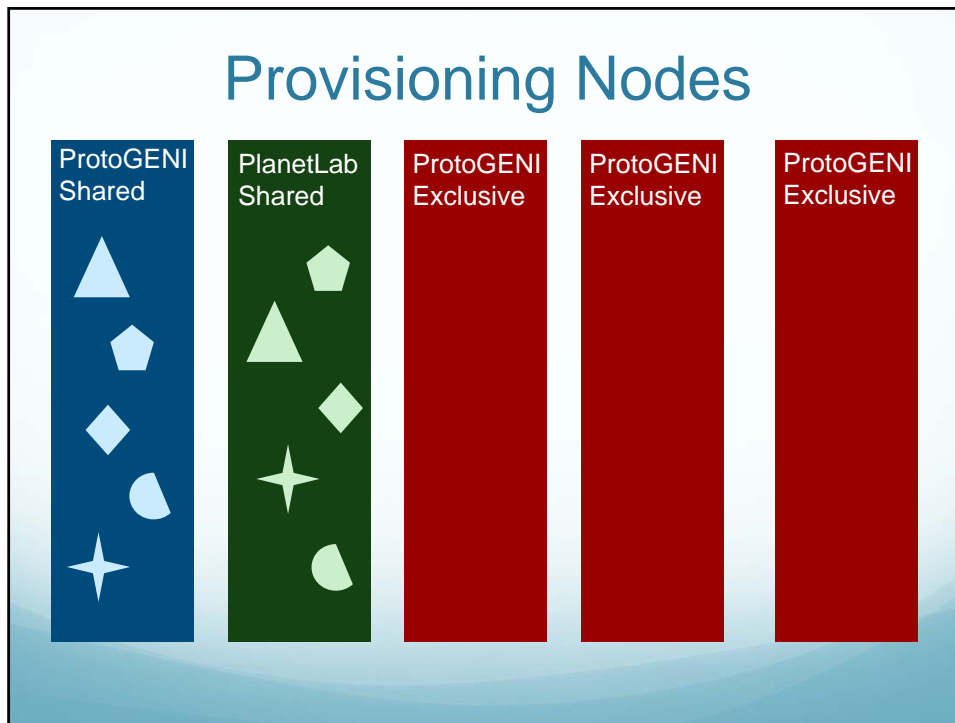


## InstaGENI Hardware

- Control node
  - Control software: ProtoGENI AM, FOAM, ProtoGENI boss
  - Image storage
  - Permanent storage for users
- Five experiment nodes
  - HP ProLiant DL360G8
  - 12 cores, 48GB RAM, 1TB disk, 4x1Gbps NIC
- HP E 5406 OpenFlow-enabled switch
  - Experiment dataplanes
- Control switch

## InstaGENI Software

- It's just a mini ProtoGENI cluster
  - Supports Hardware as a Service (HaaS)
    - Users can install custom OS images on bare metal
  - Integrated with FlowVisor OpenFlow Agg Mgr (FOAM)
    - OpenFlow programmability of experiment dataplane
  - OpenVZ image for shared nodes
  - Can run a PlanetLab node image
    - Familiar, Linux-based experiment environment
    - Platform as a Service (PaaS)
- Deploying systems that predate GENI



## PlanetLab on InstaGENI

- Deploy PlanetLab nodes on InstaGENI racks
  - Lightweight slicing of InstaGENI compute resources
  - Familiar, widely-used experiment environment
- Distinct infrastructure from public PlanetLab
  - Bring up InstaGENI-PlanetLab Central (IG-PLC)
    - Use MyPLC (PlanetLab software release)
    - Separate administration
  - Federate with other GENI testbeds
    - Easily allocate resources across all InstaGENI racks

## PlanetLab? Really?

- But PlanetLab is no longer cutting-edge!
  - Some implementation decisions are a decade old
  - Many new virtualization technologies available
- But PlanetLab is designed for overlay networks!
  - GENI is about deep network programmability, PlanetLab only supports L4+ networking
  - GENI requirements
    - Sandbox experiments at L2
    - OpenFlow-enabled experiment dataplanes

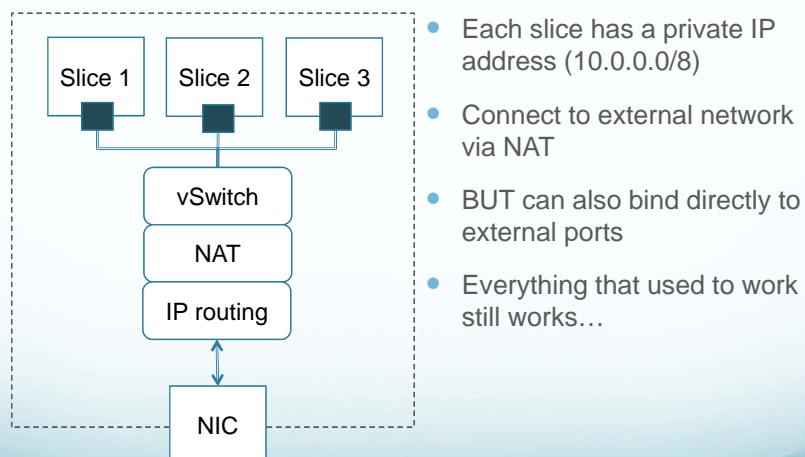
## PlanetLab Refresh

- InstaGENI will run next-gen PlanetLab node image
- PlanetLab has used VServer for lightweight slicing
  - At time VServer was most mature technology
  - VServer has few developers, requires kernel patch
- Moving to Linux Containers (LXC)
  - Similar to VServers: lightweight OS containers
  - In mainline Linux kernel
    - Better hardware support, security fixes, etc.
  - Contribute our expertise to making LXC more robust

## PlanetLab Networking

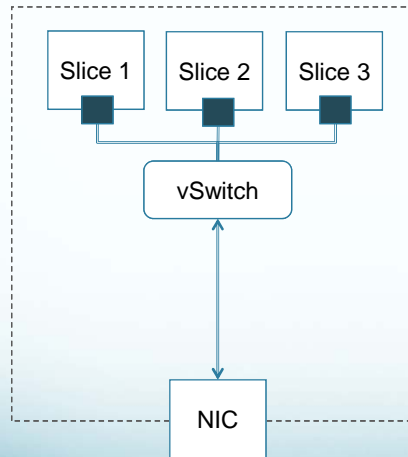
- VServer has minimal network virtualization
  - All slices have same IP address, share port space
  - Single shared stack, per-slice traffic isolated only
- LXC allows customization of network stack
  - Each slice has its own network namespace
  - Can customize at L2 and above: per-slice virtual devices, routing table, firewall, bridging, TCP parms
- Leverage Open vSwitch
  - Flexibly supports connecting VMs to the network

## PlanetLab Classic



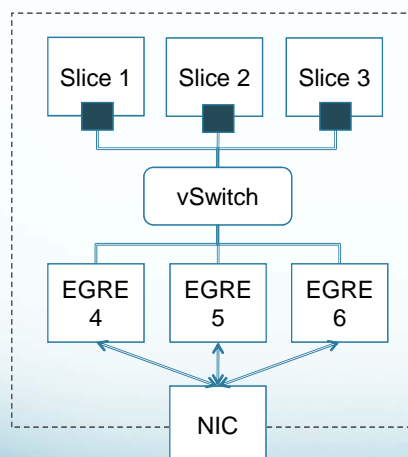
- Each slice has a private IP address (10.0.0.0/8)
- Connect to external network via NAT
- BUT can also bind directly to external ports
- Everything that used to work still works...

## Measurement Lab



- Gather data on ISP performance
- Joint project with Google and New America Foundation
- Each slice has a public IP address
- Virtual interfaces bridged at L2 to external NIC

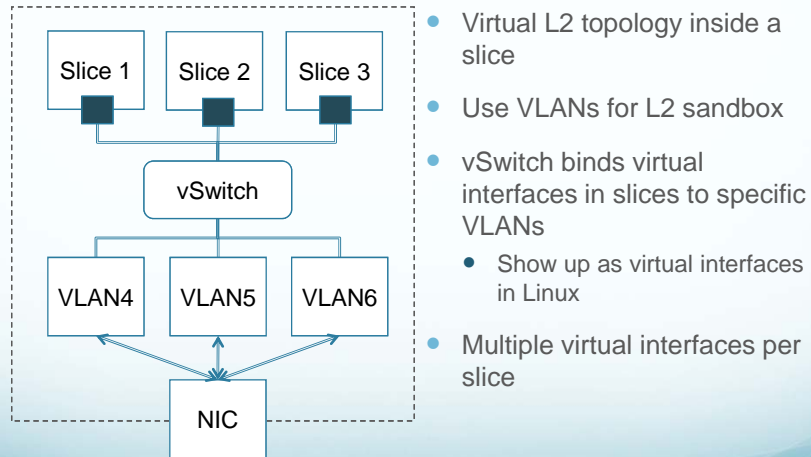
## VINI



- Virtual L2 topology inside a slice
- Tunnel L2 over L3 (EGRE)
- vSwitch binds virtual interfaces in slices to specific EGRE tunnels
  - Show up as virtual interfaces in Linux
- Multiple virtual interfaces per slice
- SIGCOMM 2006



## InstaGENI



- Virtual L2 topology inside a slice
- Use VLANs for L2 sandbox
- vSwitch binds virtual interfaces in slices to specific VLANs
  - Show up as virtual interfaces in Linux
- Multiple virtual interfaces per slice

## Current Status

- InstaGENI racks
  - One operational (Utah), shipping to other sites now
- Next-gen PlanetLab image (LXC + OvS)
  - In testing on VICCI
  - Only “PlanetLab Classic” model currently supported
  - Expect IG-PL image to be ready in September

# Thanks!

- GENI racks / InstaGENI:
  - <http://groups.geni.net/geni/wiki/GENIRacksHome>
- PlanetLab-based testbeds
  - <http://www.planet-lab.org>
  - <http://www.measurementlab.net>
  - <http://www.vicci.org>
  - <http://www.vini-veritas.net>