Global Environment for Network Innovations

Jennifer Rexford

Princeton University

http://www.cs.princeton.edu/~jrex

Outline

- Networking research challenges
  - Security, economic incentives, management, layer-2 technologies
- Importance of building and deploying
  - Bridging the gap between simulation/testbeds and real deployment
- Global Environment for Network Innovations (GENI)
  - Major NSF initiative to support experimental networking research
  - Key ideas: virtualization, programmability, and user opt-in
- GENI backbone design
  - Programmable routers, flexible optics, and connection to Internet
  - Example experiments highlighting the capabilities
- Conclusions

Is the Internet broken?

- It is great at what it does.
  - Everyone should be proud of this.
  - All sorts of things can be built on top of it.
- But…
  - Security is weak and not getting better.
  - Availability continues to be a challenge.
  - It is hard to manage and getting harder.
  - It does not handle mobility well.
  - A long list, once you start…

Challenges Facing the Internet

- Security and robustness
  - Naming and identity
  - Availability
- Economic incentives
  - Difficulty of providing end-to-end services
  - Commoditization of the Internet infrastructure
- Network management
  - No framework in the original Internet design
  - Tuning, troubleshooting, accountability, …
- Interacting with underlying network technologies
  - Advanced optics: dynamic capacity allocation
  - Wireless: mobility, dynamic impairments
  - Sensors: small embedded devices at large scale

FIND: Future Internet Design

- NSF research initiative
  - Requirements for global network of 10-15 years out?
  - Re-conceive the network, if we could design from scratch?
- Conceive the future, by letting go of the present:
  - This is not change for the sake of change
  - Rather, it is a chance to free our minds
  - Figuring out where to go, and then how to get there
- Perhaps a header format is not the defining piece of a new architecture
  - Definition and placement of functionality
  - Not just data plane, but also control and management
  - And division between end hosts and the network

The Importance of Building

- Systems-oriented computer science research needs to build and try out its ideas to be effective
  - Paper designs are just idle speculation
  - Simulation is only occasionally a substitute
- We need:
  - Real implementation
  - Real experience
  - Real network conditions
  - Real users
  - To live in the future

Need for Experimental Facility

Goal: Seamless conception-to-deployment process

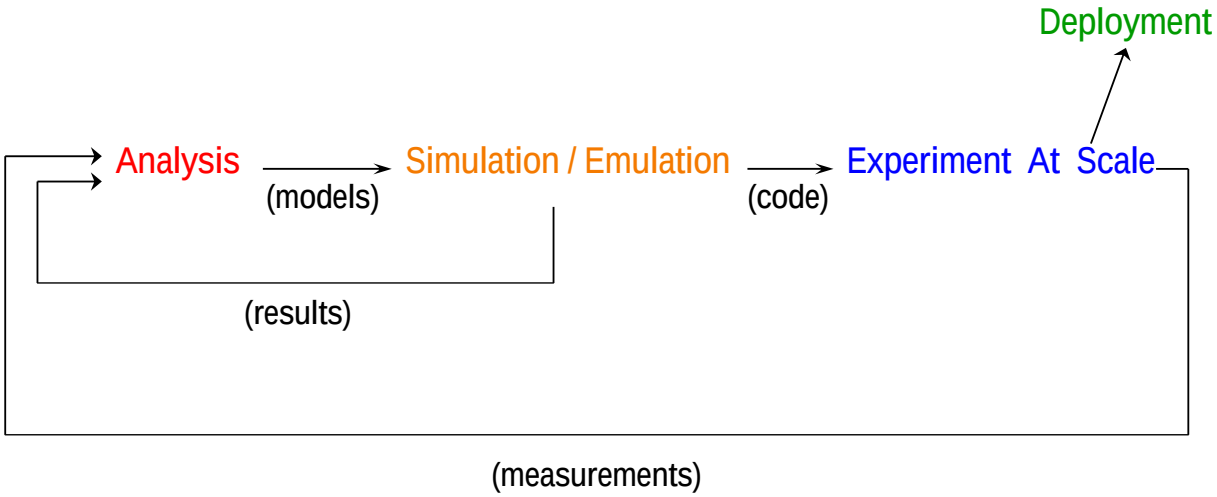

Existing Tools

- Simulators
  - ns
- Emulators
  - Emulab
  - WAIL
- Wireless testbeds
  - ORBIT
  - Emulab
- Wide-area testbeds
  - PlanetLab
  - RON
  - X-bone
  - DETER

Today's Tools Have Limitations

- Simulation based on simple models
  - Topologies, administrative policies, workloads, failures…
- Emulation (and "in lab" tests) are similarly limited
  - Only as good as the models
- Traditional testbeds are targeted
  - Not cost-effective to test every good idea
  - Often of limited reach
  - Often with limited programmability
- Testbed dilemma
  - Production network: real users, but hard to make changes
  - Research testbed: easy to make changes, but no users

Bridging the Chasm

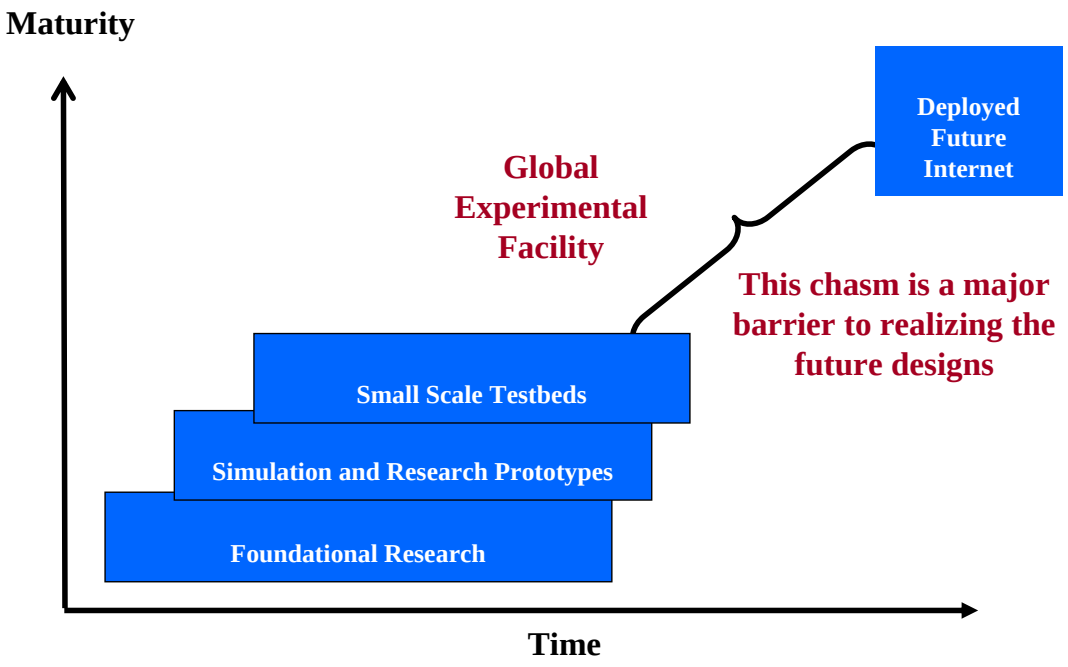

Goals for the Experimental Facility

- Broader impact
  - Positive influence on the design of the future Internet
  - Network that is more secure, reliable, efficient, manageable, usable
- Intellectual progress
  - Network science
    - Experimentally answer questions about complex systems
    - Better understanding of dynamics, stability, evolvability, etc.
  - Network architecture
    - Evaluate and compare alternative architectural structures
    - Reconcile the contradictory goals an architecture must meet
  - Network engineering
    - Evaluate engineering trade-offs in a controlled, realistic setting
    - Test theories of how different elements might be designed

GENI

- Experimental facility
  - MREFC proposal to build a large-scale facility
  - Jointly from NSF's CS directorate, & research community
  - We are currently at the "Conceptual Design" stage
  - Will eventually require Congressional approval
- Global Environment for Network Innovations
  - Prototyping new architectures
  - Realistic evaluation
  - Controlled evaluation
  - Shared facility
  - Connecting to real users
  - Enabling new services

See http://www.geni.net

Three Key Ideas in GENI

- Virtualization
  - Multiple architectures on a shared facility
  - Amortizes the cost of building the facility
  - Enables long-running experiments and services
- Programmable
  - Enable prototyping and evaluation of new architectures
  - Enable a revisiting of today's "layers"
- Opt-in on a per-user / per-application basis
  - Attract real users
    - Demand drives deployment / adoption
  - Connect to the Internet
    - To reach users, and to connect to existing services

Slices

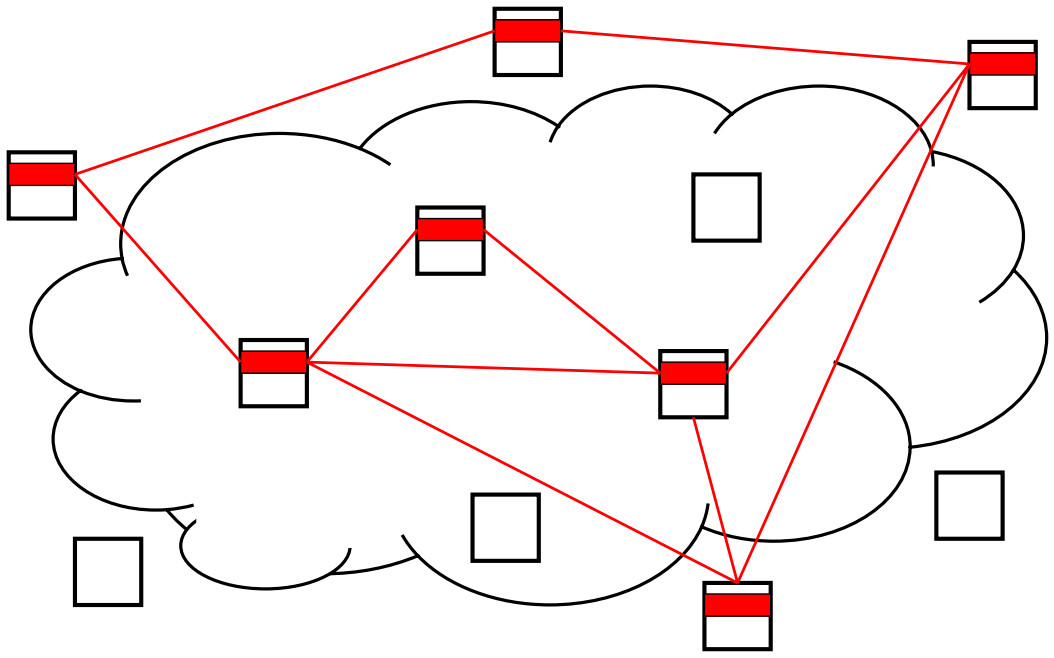

Slices

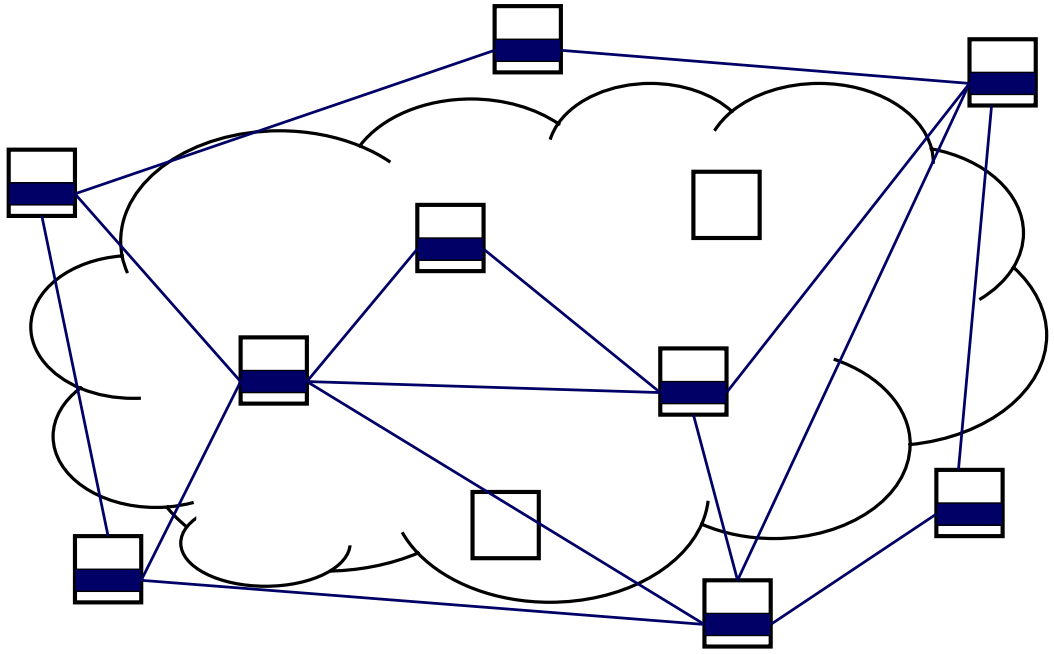

User Opt-in

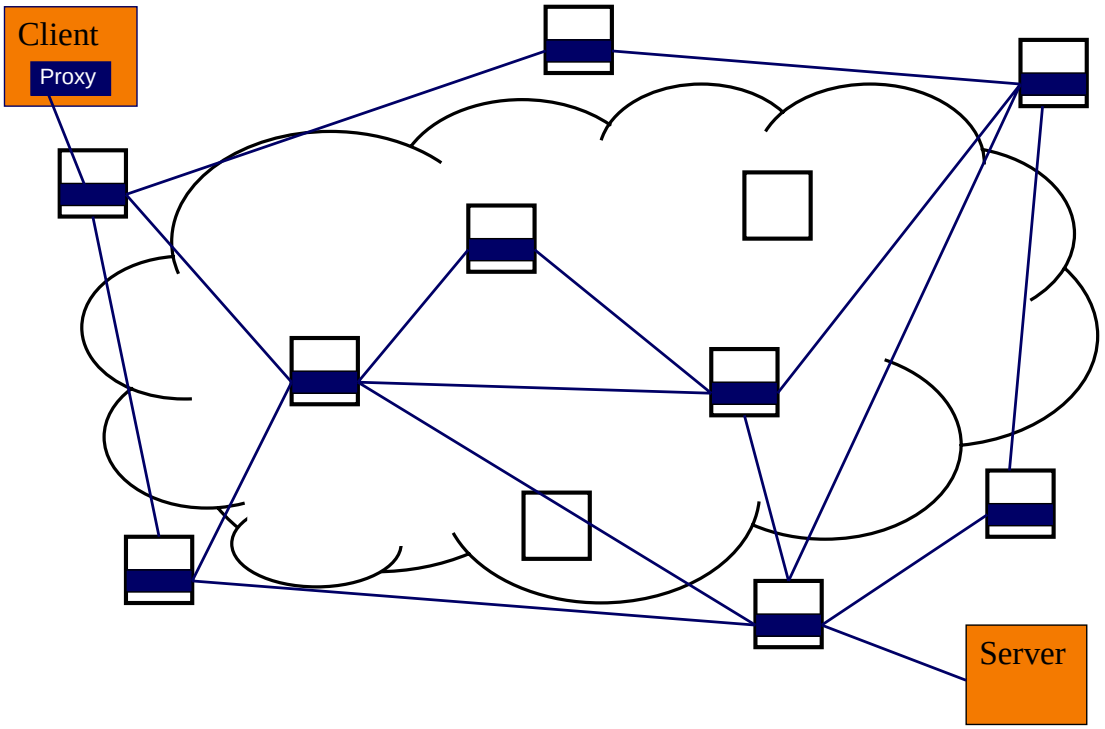

Realizing the Ideas

- Slices embedded in a substrate of resources
  - Physical network substrate
    - Expandable collection of building block components
    - Nodes / links / subnets
  - Software management framework
    - Knits building blocks together into a coherent facility
    - Embeds slices in the physical substrate
- Builds on ideas in past systems
  - PlanetLab, Emulab, ORBIT, X-Bone, …

National Fiber Facility

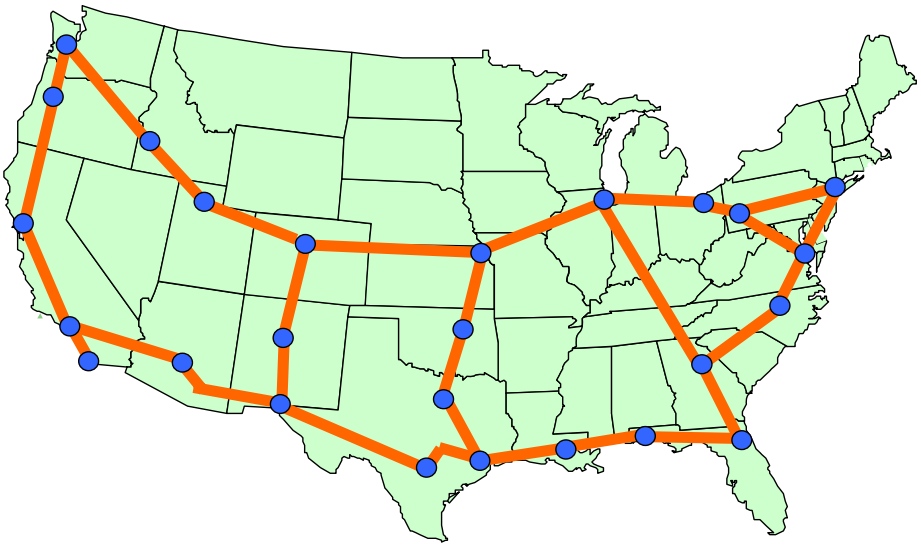

+ Programmable Routers

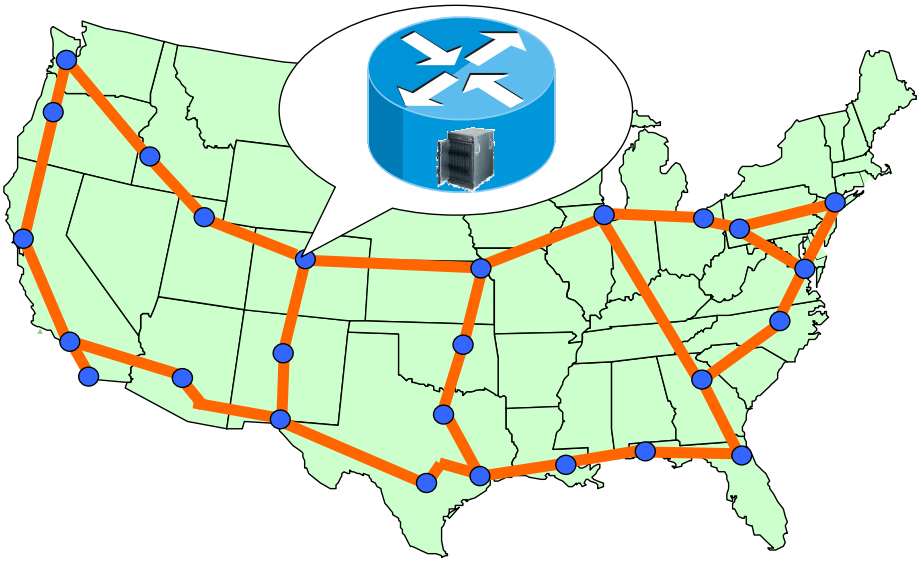

+ Clusters at Edge Sites

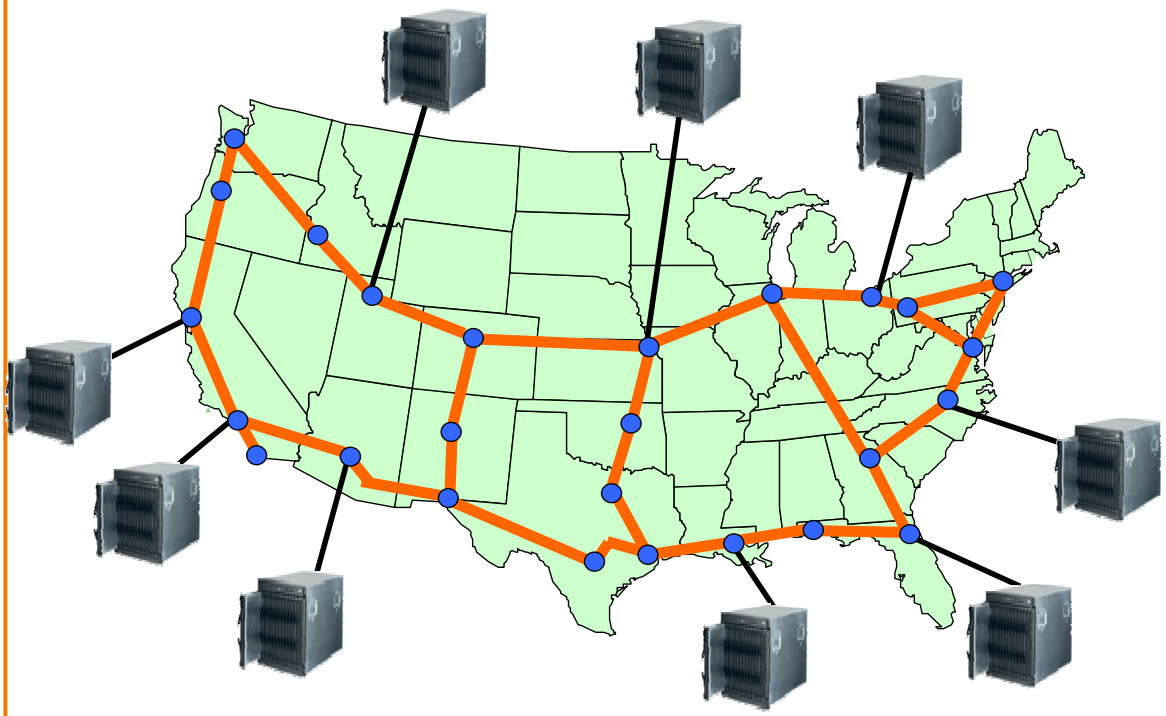

+ Wireless Subnets

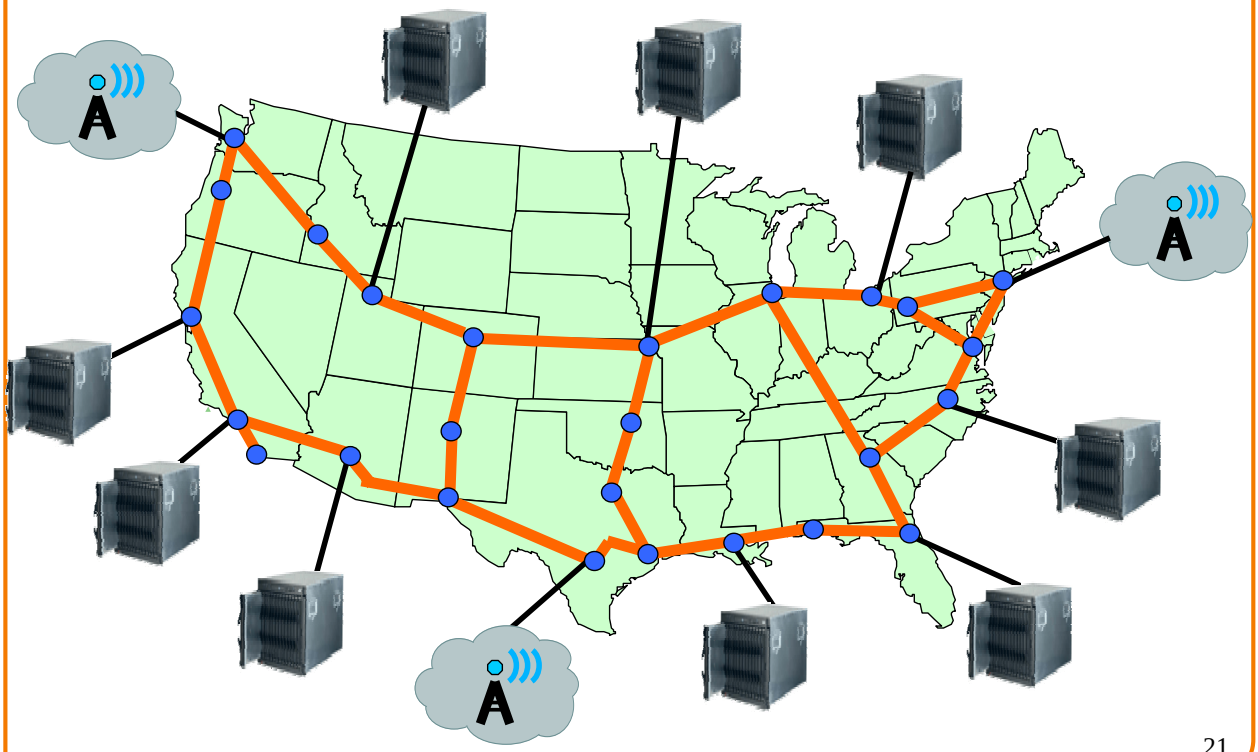

+ ISP Peers

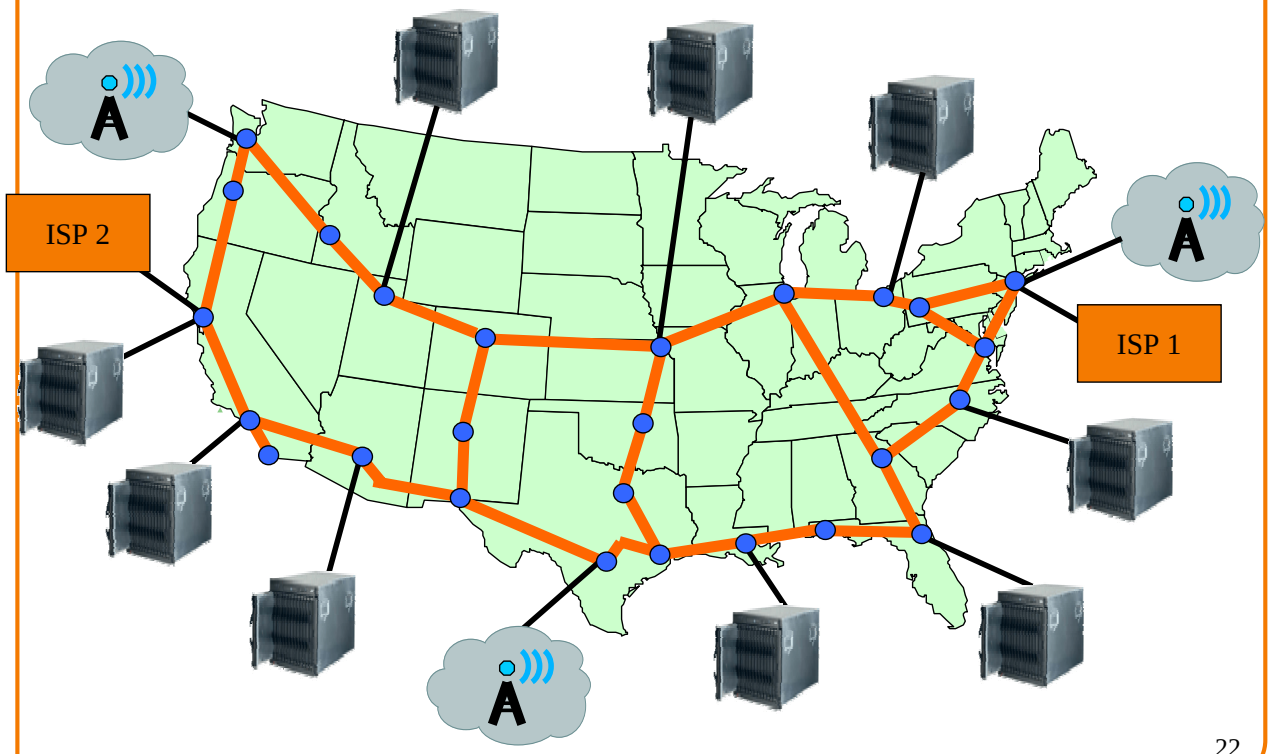

Closer Look

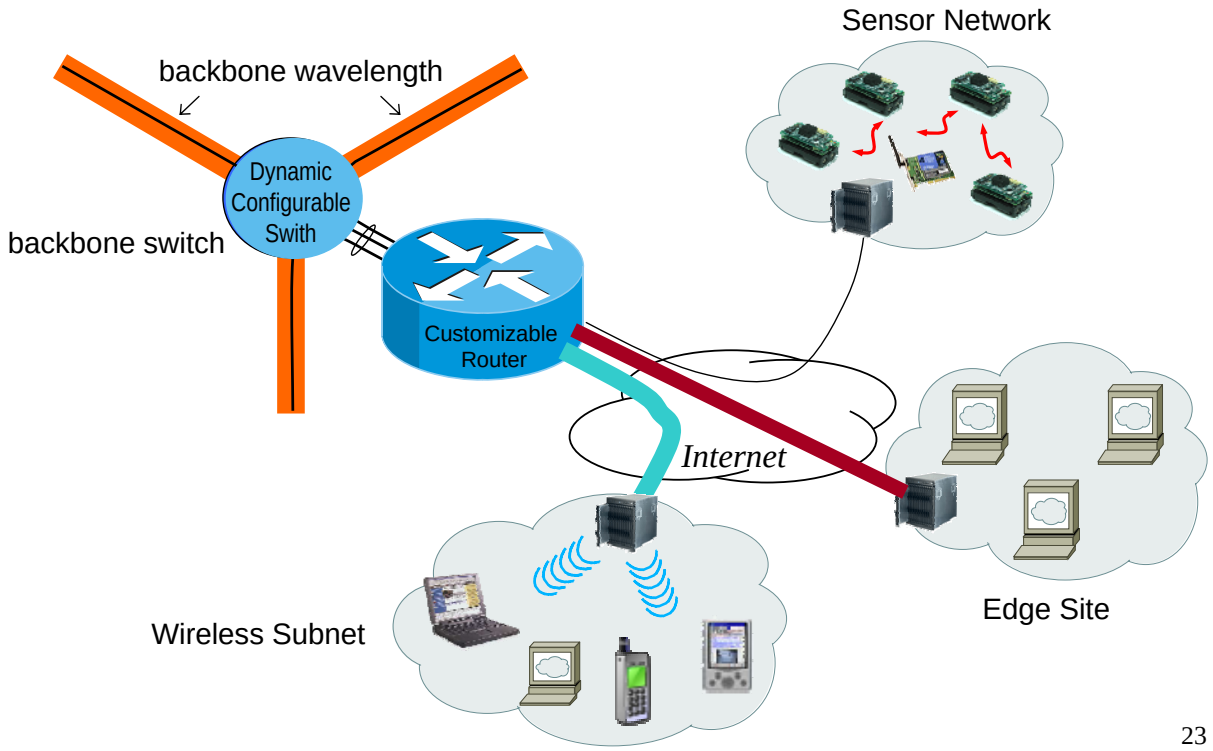

GENI Substrate: Summary

- Node components
  - Edge devices
  - Customizable routers
  - Optical switches
- Bandwidth
  - National fiber facility
  - Tail circuits
- Wireless subnets
  - Urban 802.11
  - Wide-area 3G/WiMax
  - Cognitive radio
  - Sensor net
  - Emulation

GENI Management Core

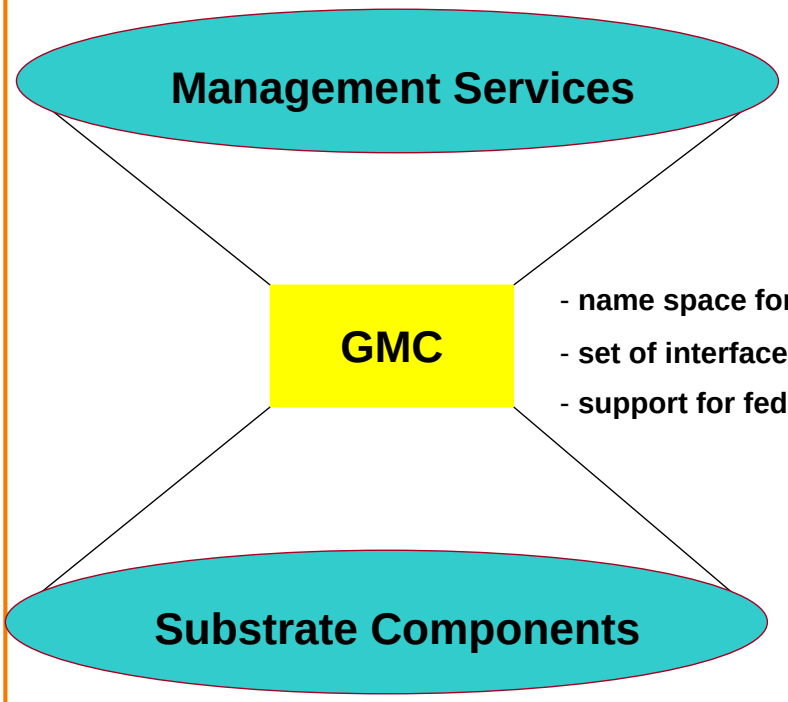

- name space for users, slices, & components
- set of interfaces (“plug in” new components)
- support for federation (“plug in” new partners)

Hardware Components

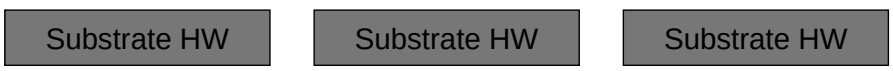

Virtualization Software

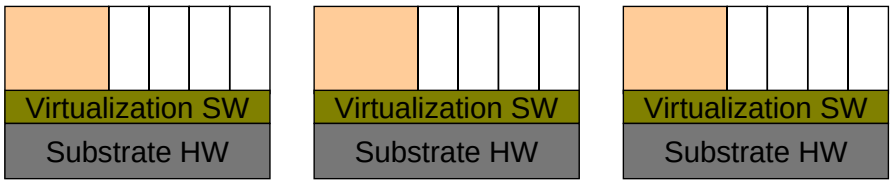

Component Manager

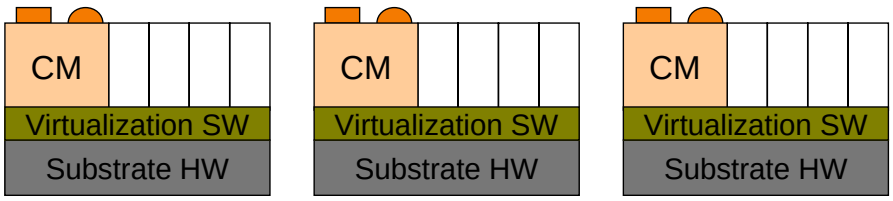

GENI Management Core (GMC)

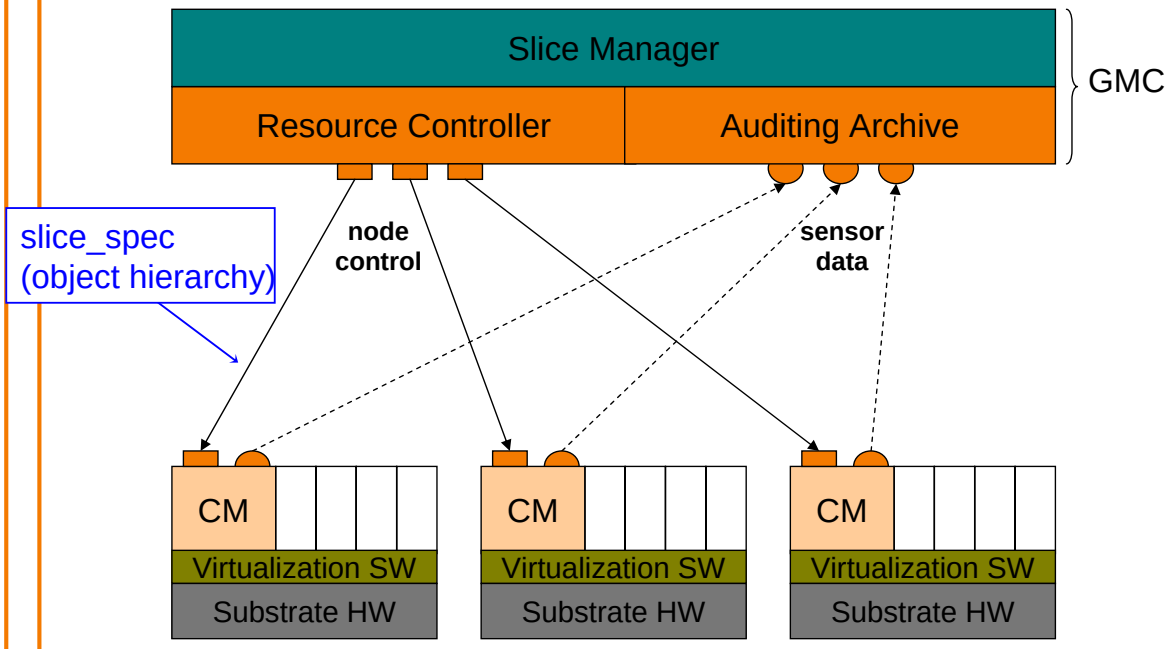

Federation

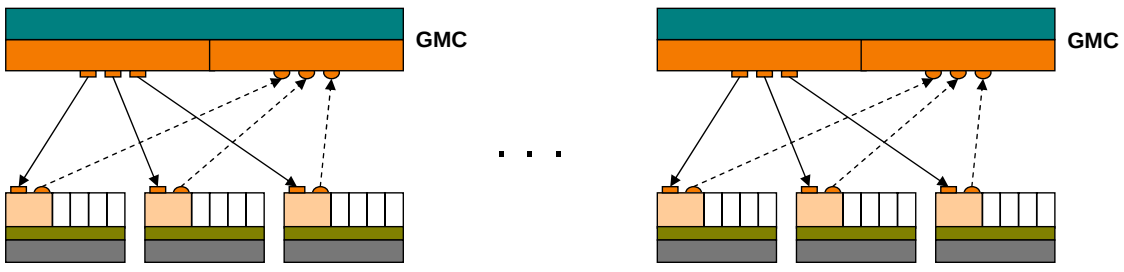

User Front-End(s)

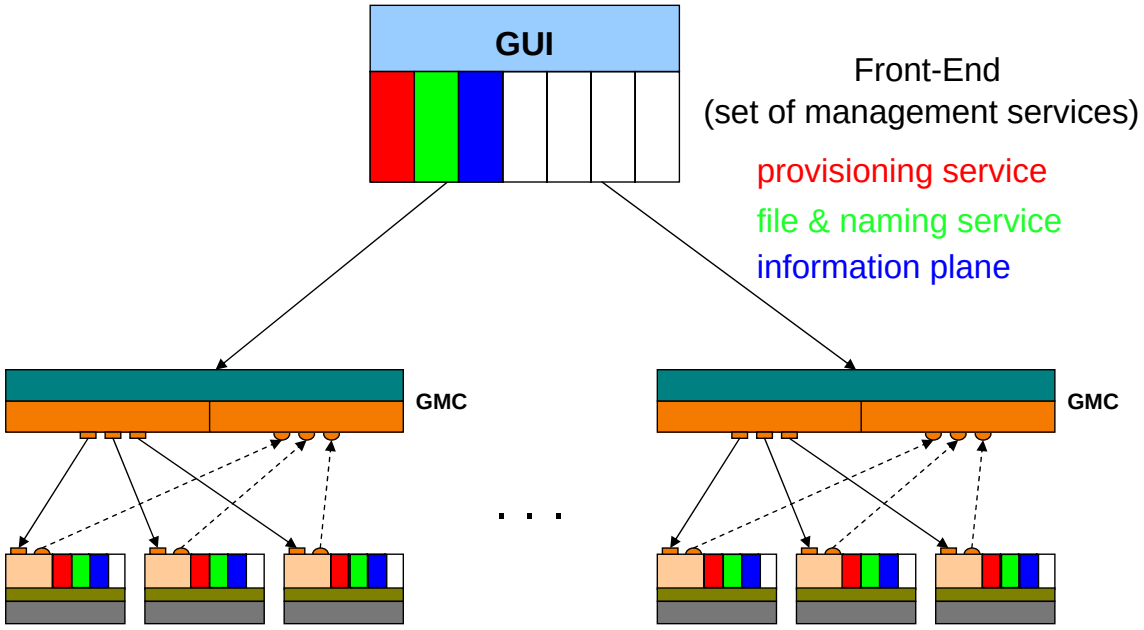

Virtualization in GENI

- Multiple levels possible
  - Different level required by different experiments
  - Different level depending on the technology
- Example “base cases”
  - Virtual server (socket interface / overlay tunnels)
  - Virtual router (virtual line card / static circuits)
  - Virtual switch (virtual control interface / dynamic circuits)
  - Virtual AP (virtual MAC / fixed spectrum allocation)
- Specialization
  - The ability to install software in your own virtual-*

Distributed Services in GENI

- Goals
  - Complete the GENI management story
  - Lower the barrier-to-entry for researchers (students)
- Example focus areas
  - Provisioning (slice embedder)
  - Security
  - Information plane
  - Resource allocation
  - Files and naming
  - Topology discovery
  - Development tools
  - Interfacing with the Internet, and IP

GENI Security

- Limits placed on a slice's "reach"
  - Restricted to slice and GENI components
  - Restricted to GENI sites
  - Allowed to compose with other slices
  - Allowed to interoperate with legacy Internet
- Limits on resources consumed by slices
  - Cycles, bandwidth, disk, memory
  - Rate of particular packet types, unique addrs per second
- Mistakes (and abuse) will still happen
  - Auditing will be essential
  - Network activity ⟶ slice ⟶ responsible user(s)

Success Scenarios

- Change the research process
  - Sound foundation for future network architectures
  - Experimental evaluation, rather than paper designs
- Create new services
  - Demonstrate new services at scale
  - Attract real users
- Aid the evolution of the Internet
  - Demonstrate ideas that ultimately see real deployment
  - Provide architectural clarity for evolutionary path
- Lead to a future global network
  - Purist: converge on a single new architecture
  - Pluralist: virtualization supporting many architectures

Working Groups to Flesh Out Design

- Research (Dave Clark and Scott Shenker)
  - Usage policy / requirements / instrumentation
- Architecture (Larry Peterson and John Wroclawski)
  - Define core modules and interfaces
- **Backbone (Jen Rexford and Dan Blumenthal)**
  - **Fiber facility / routers & switches / tail circuits / peering**
- Wireless (Dipankar Raychaudhuri and Deborah Estrin)
  - RF technologies / deployment
- Services (Tom Anderson, Reiter)
  - Edge sites / infrastructure and underlay services
- Education
  - Training / outreach / course development

GENI Backbone Requirements

- Programmability
  - Flexible routing, forwarding, addressing, circuit set-up, …
- Isolation
  - Dedicated bandwidth, circuits, CPU, memory, disk
- Realism
  - User traffic, ISP connection, propagation delays, failures
- Control
  - Inject failures, create circuits, exchange routing updates
- Performance
  - High-speed packet forwarding and low delays
- Security
  - Preventing attacks on the Internet, and on GENI itself

A Researcher's View of GENI Backbone

- Virtual network topology
  - Nodes and links in a particular topology
  - Resources and capabilities per node/link
  - Embedded in the GENI backbone
- Virtual router and virtual switch
  - Abstraction of a router and switch per node
  - To evaluate new architectures (routing, switching, addressing, framing, grooming, layering, …)
- GENI backbone capabilities evolve over time
  - To realize the abstractions at finer detail
  - To scale to a larger number of experiments

Creating a Virtual Topology

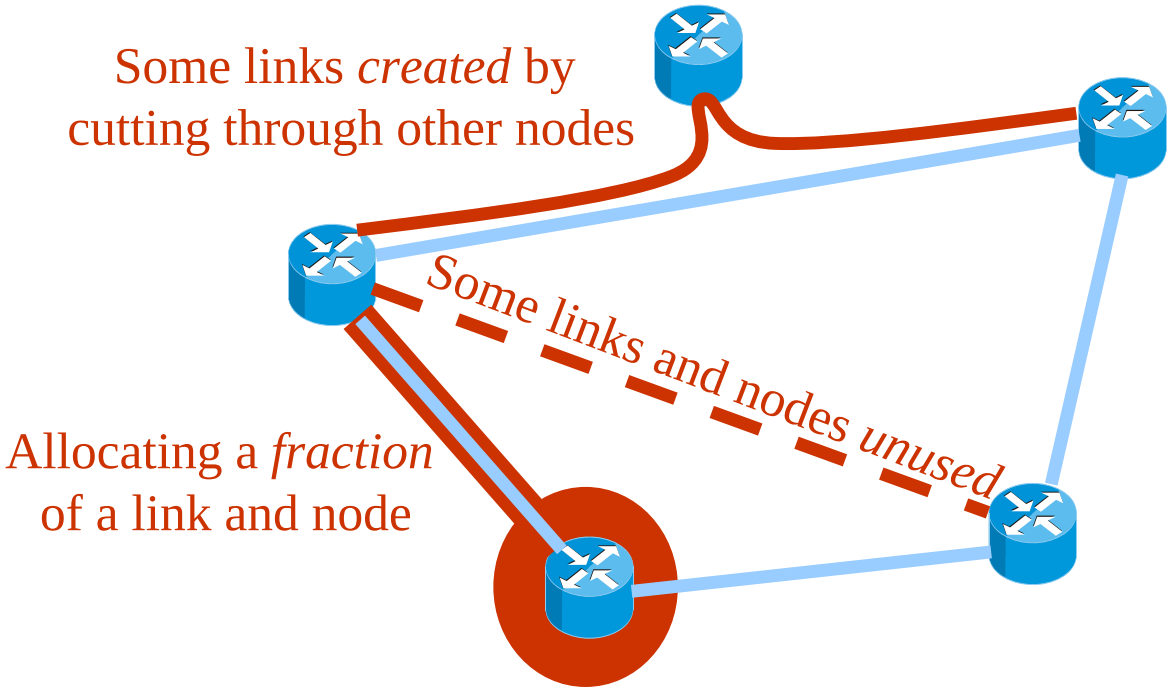

GENI Backbone

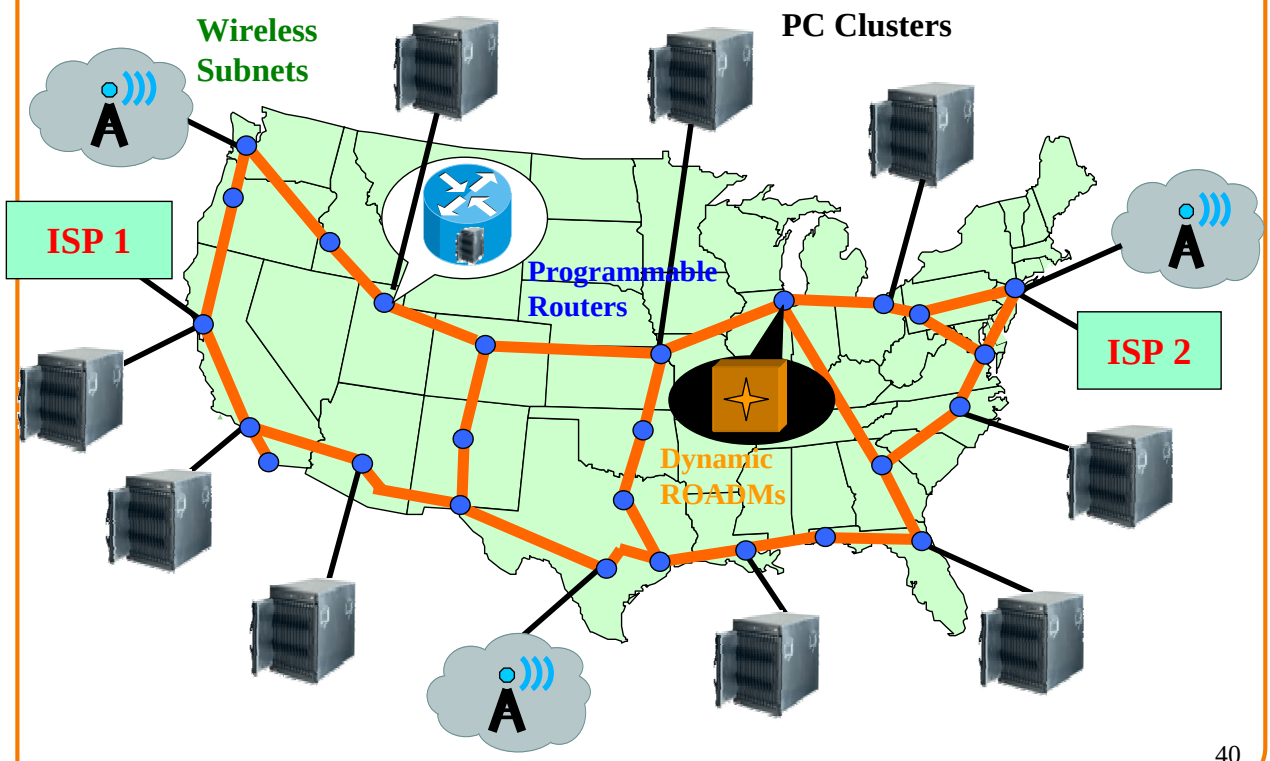

GENI Backbone Node Components

- Phase 0 – General purpose blade server
  - Single node with collection of assignable resources
  - Virtual Router may be assigned VM, blade or >1 blades
- Phase 1 – Adding higher performance components
  - Assignable Network Processor blades and FPGA blades
  - NPs also used for I/O for better control of I/O bandwidth
- Phase 2 – Adding reconfigurable cross-connect
  - Enable experiments with configurable transport layer
  - Provide “true circuits” between backbone virtual routers
- Phase 3 – Adding dynamic optical switch
  - Dynamic optical switch with programmable groomer and framer, and reconfigurable add/drop multiplexers

GENI Backbone Node Components

- Phase 0 – General purpose blade server
  - Node with collection of assignable resources
  - Virtual Router may be assigned a virtual machine, blade, or multiple blades

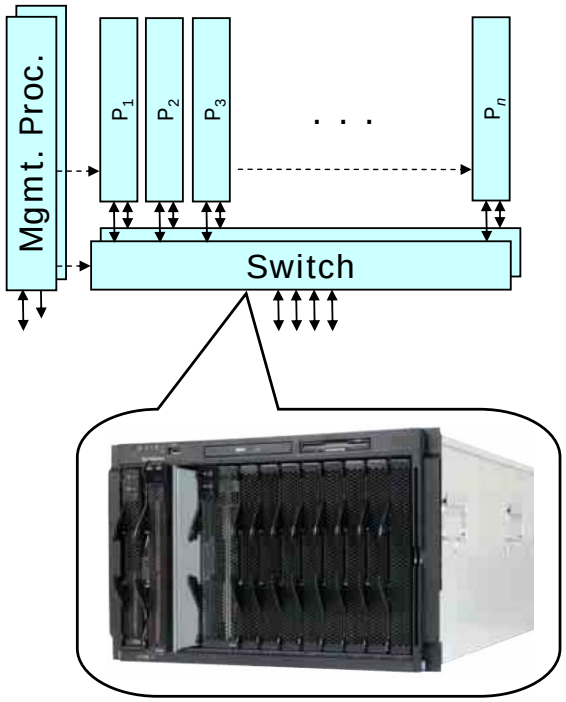

GENI Backbone Node Components

- Phase 1 – Adding higher performance components
  - Assignable Network Processor blades and FPGA blades
  - NPs also used for I/O for better control of bandwidth
  - ATCA chassis and blades

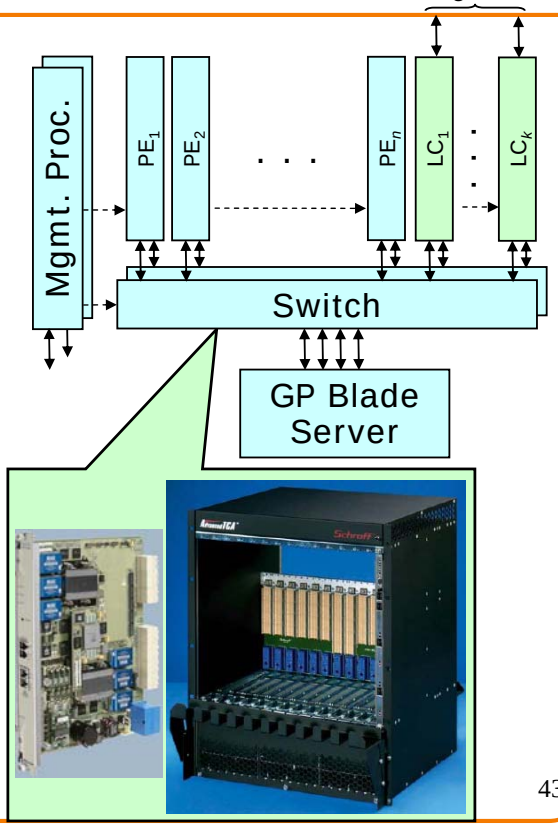

GENI Backbone Node Components

- Phase 2 – Reconfigurable cross-connect
  - Enable experiments with configurable transport layer
  - Provide "true circuits" between backbone virtual routers
  - Cut-through traffic circumvents the router

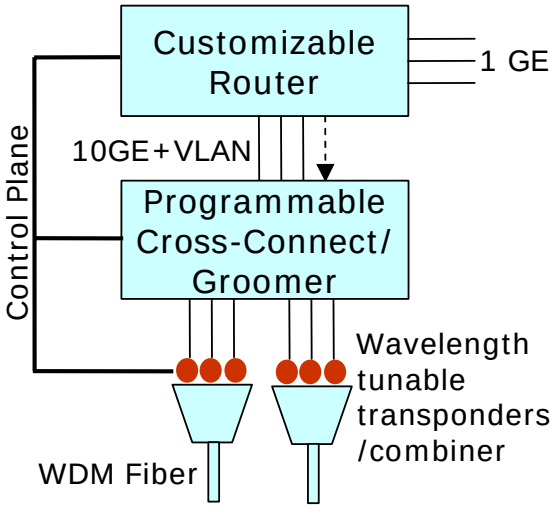

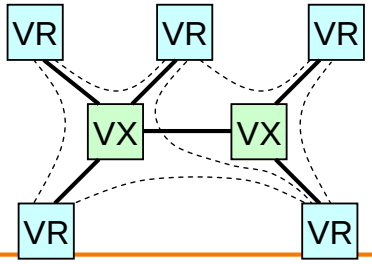

GENI Backbone Node Components

- Phase 3 – Adding dynamic optical switch
  - Dynamic optical switch with programmable groomer and framer, and reconfigurable add/drop multiplexers
  - Maleable bandwidth
  - Arbitrary framing

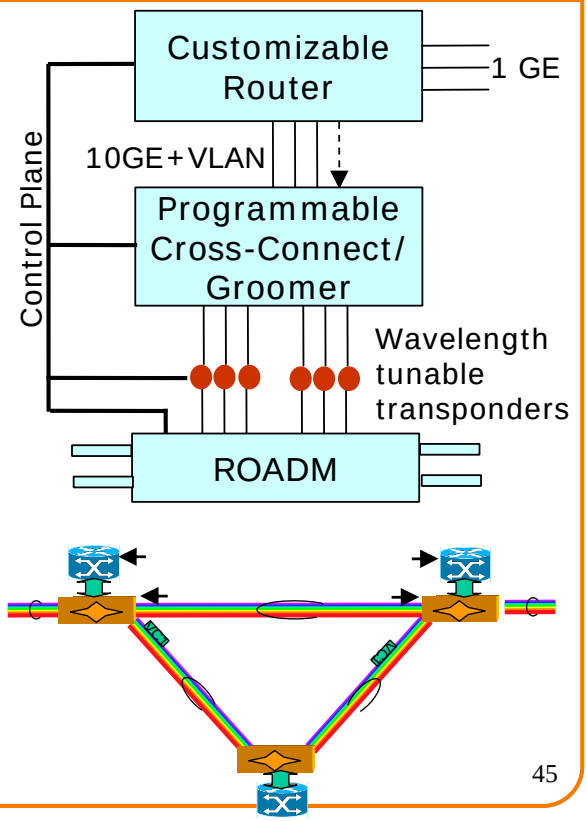

GENI Backbone Software

- Component manager and virtualization layer
  - Abstraction of virtual router and virtual switch
  - Setting scheduling parameters for subdividing resources
- Multiplexers for resources hard to share
  - Single BGP session with the outside world
  - Single interface to an element-management system
- Exchanging traffic with the outside world
  - Routing and forwarding software to evaluate & extend
  - VPN servers and NATs at the GENI/Internet boundary
- Libraries to support experimentation
  - Specifying, controlling, and measuring experiments
  - Auditing and accounting to detect misbehavior

Feasibility

- Industrial trends and standards
  - Advanced Telecom Computing Architecture (ATCA)
  - Network processors and FPGAs
  - SONET cross connects and ROADMs
- Open-source networking software
  - Routing protocols, packet forwarding, network address translation, diverting traffic to an overlay
- Existing infrastructure
  - PlanetLab nodes, software, and experiences
  - National Lambda Rail and Abilene backbones
  - New Virtualized Network Infrastructure (VINI)

Conclusions

- Future Internet poses many research challenges
  - Security, network management, economics, layer-2, …
- Research community should rise to the challenge
  - Conceive of future network architectures
  - Prototype and evaluate architectures in realistic settings
- Global Environment for Network Innovations
  - Facility for evaluating new network architectures
  - Virtualization, programmability, and user opt-in
- GENI backbone design
  - Fiber facility, tail circuits, and upstream connectivity
  - Programmable router and dynamical optical switch