A Strategy for Continually Reinventing the Internet

Larry Peterson
Princeton University

Challenges

- Security
  - known vulnerabilities lurking in the Internet
    - DDoS, worms, malware
  - addressing security comes at a significant cost
    - federal government spent $5.4B in 2004
    - estimated $50-100B spent worldwide on security in 2004
- Reliability
  - e-Commerce increasingly depends on fragile Internet
    - much less reliable than the phone network (three vs five 9's)
    - risks in using the Internet for mission-critical operations
    - barrier to ubiquitous VoIP
  - an issue of *ease-of-use* for everyday users

Challenges (cont)

- Scale & Diversity
  - the whole world is becoming networked
    - sensors, consumer electronic devices, embedded processors
  - assumptions about edge devices (hosts) no longer hold
    - connectivity, power, capacity, mobility,…
- Performance
  - scientists have significant bandwidth requirements
    - each e-science community covets its own wavelength(s)
  - purpose-built solutions are not cost-effective
    - being on the “commodity path” makes an effort sustainable

Two Paths

- Incremental
  - apply point-solutions to the current architecture
- Clean-Slate
  - replace the Internet with a new network architecture

- We can't be sure the first path will fail, but…
  - point-solutions result in increased complexity
    - making the network harder to manage
    - making the network more vulnerable to attacks
    - making the network more hostile to new applications
  - architectural limits may lead to a dead-end

Architectural Limits

- Minimize trust assumptions
  - the Internet originally viewed network traffic as fundamentally cooperative, but should view it as adversarial
- Enable competition
  - the Internet was originally developed independent of any commercial considerations, but today the network architecture must take competition and economic incentives into account
- Allow for edge diversity
  - the Internet originally assumed host computers were connected to the edges of the network, but host-centric assumptions are not appropriate in a world with an increasing number of sensors and mobile devices

Limits (cont)

- Design for network transparency
  - the Internet originally did not expose information about its internal configuration, but there is value to both users and network administrators in making the network more transparent
- Enable new network services
  - the Internet originally provided only a best-effort packet delivery service, but there is value in making processing capability and storage capacity available in the middle of the network
- Integrate with optical transport
  - the Internet originally drew a sharp line between the network and the underlying transport facility, but allowing bandwidth aggregation and traffic engineering to be first-class abstractions has the potential to improve efficiency and performance

Barriers to Second Path

- Internet has become ossified
  - no competitive advantage to architectural change
  - no obvious deployment path
- Inadequate validation of potential solutions
  - simulation models too simplistic
  - little or no real-world experimental evaluation
- Testbed dilemma
  - production testbeds: real users but incremental change
  - research testbeds: radical change but no real users

Recommendation

It is time for the research community, federal government, and commercial sector to jointly pursue the second path. This involves experimentally validating new network architecture(s), and doing so in a sustainable way that fosters wide-spread deployment.

Why Now?

- Active research community
  - scores of architectural proposals
  - ready to step up to the challenge of making it real
- Enabling technologies
  - OS virtualization and interposition mechanisms
  - overlay networks are maturing
  - high-speed data pipes in the core
  - fast network processors and FPGAs
- Infrastructure exists
  - PlanetLab
  - National Lambda Rail (NLR)

PlanetLab

- 580 machines spanning 275 sites and 30 countries
  - nodes within a LAN-hop of > 2M users
- Supports *distributed virtualization*
  - each of 425 network services running in their own *slice*

Examples Services

- Content Distribution Networks
  - CoDeeN (Princeton), Coral (NYU), Coweb (Cornell)
- Distributed Hash Tables
  - OpenDHT (Berkeley), Chord (MIT)
- Large File Transfer
  - CoBlitz (Princeton), SplitStream (Rice), Bullet (UCSD)
- Routing Overlays
  - i3 (Berkeley), Pluto (Princeton)
- Network Measurement
  - ScriptRoute (Maryland, Washington)
- Anomaly Detection & Fault Diagnosis
  - NetBait (Intel), PlanetSeer (Princeton)
- Multicast, Mobility, Network Games, DNS,…

National LambdaRail

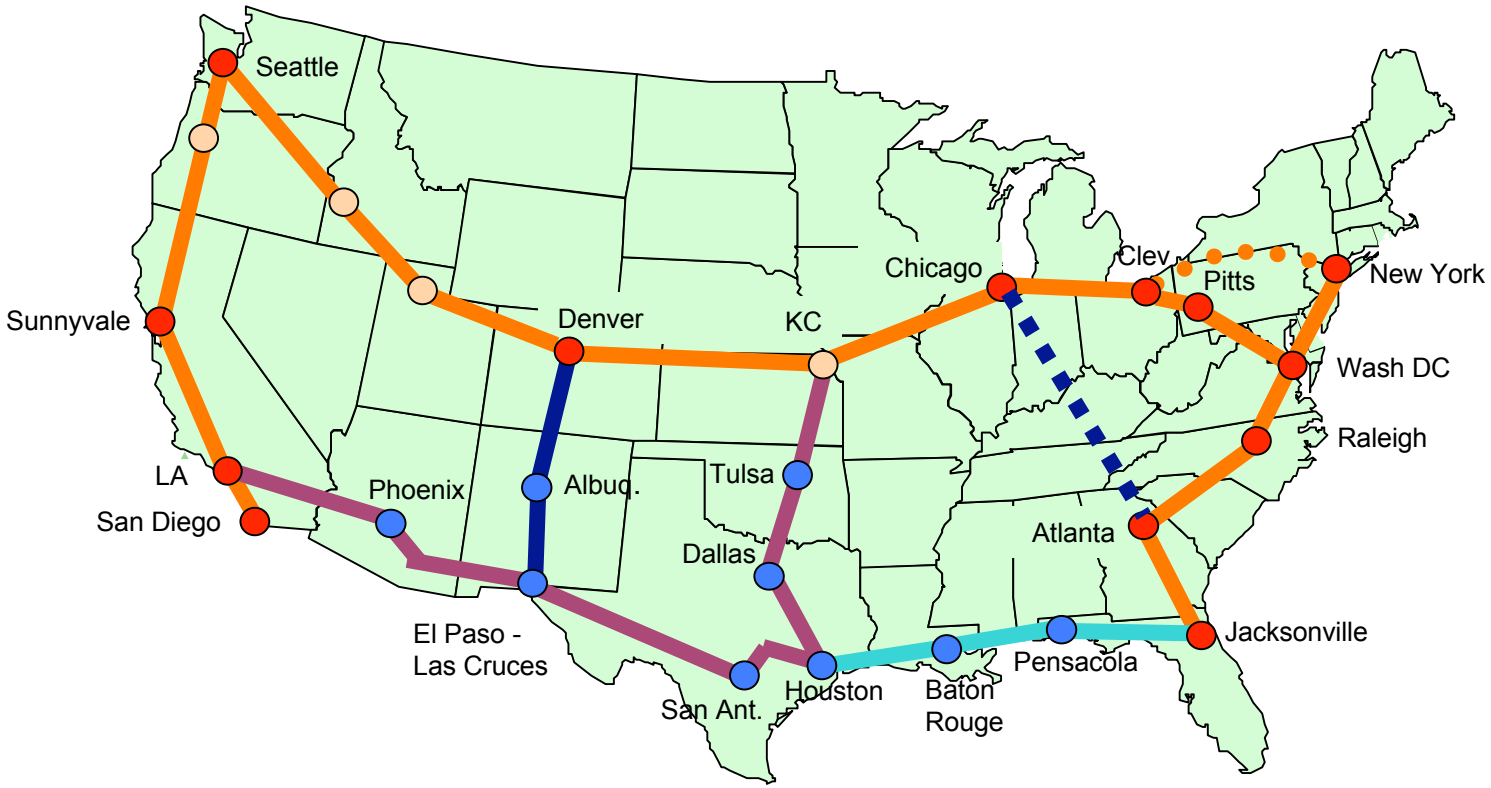

- 10Gbps per-lambda
- Lambdas set aside for network research

Next Step: Meta Testbed

- Goals
  - support experimental validation of new architectures
    - simultaneously support real users and clean slate designs
    - allow a thousand flowers to bloom
  - provide plausible deployment path
- Key ideas
  - virtualization
    - multiple architectures on a shared infrastructure
    - shared management costs
  - opt-in on a per-user / per-application basis
    - attract real users
    - demand drives deployment / adoption

Meta Testbed

- Infrastructure
  - PlanetLab provides “access network” with global reach
    - user desktops run *proxy* that allows them to opt-in
    - treat nearby PlanetLab node as *ingress router*
  - NLR provides high-speed backbone
    - populate with programmable routers
    - extend slice abstraction to these routers
- Usage model
  - each architecture (service) runs in its own slice
  - two modes of use
    - short-term experiments
    - long-running stable architectures and services

Slices

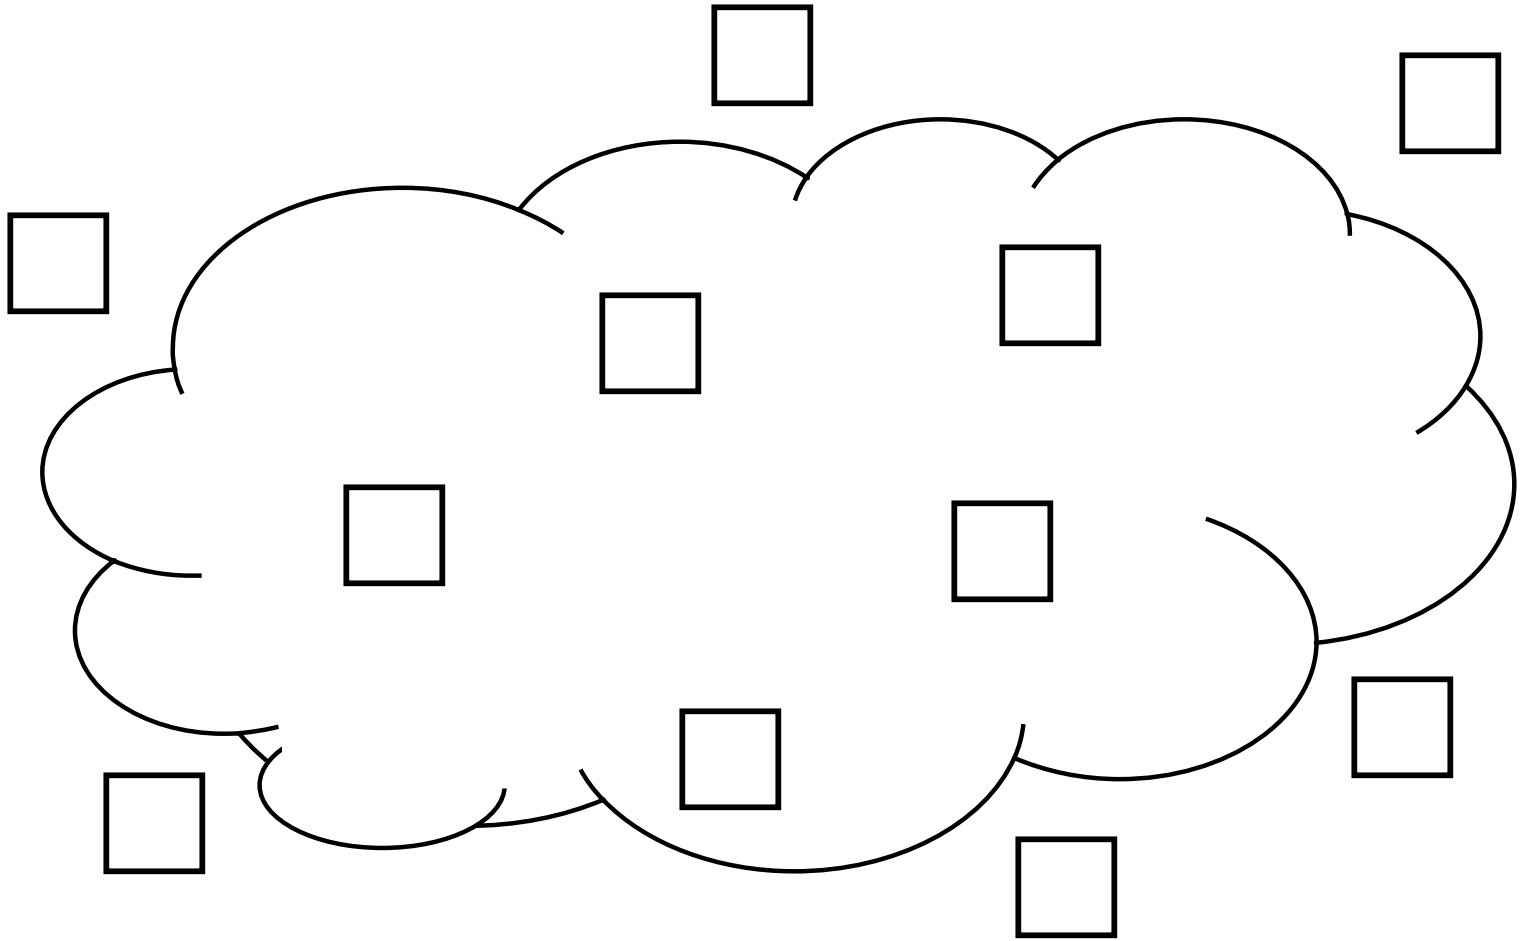

Slices

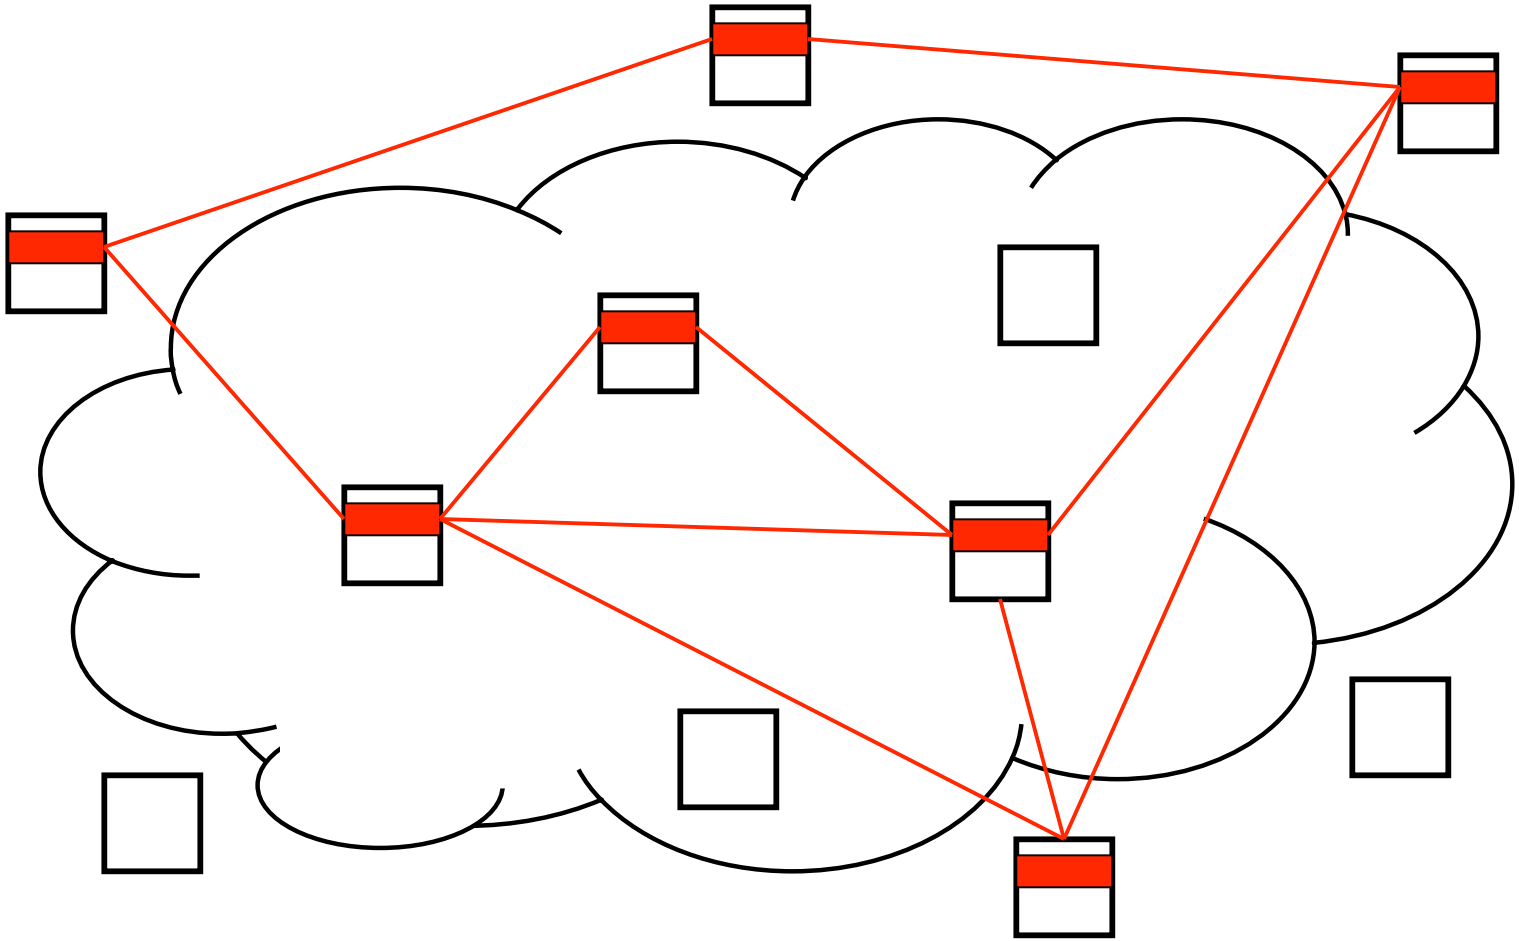

Slices

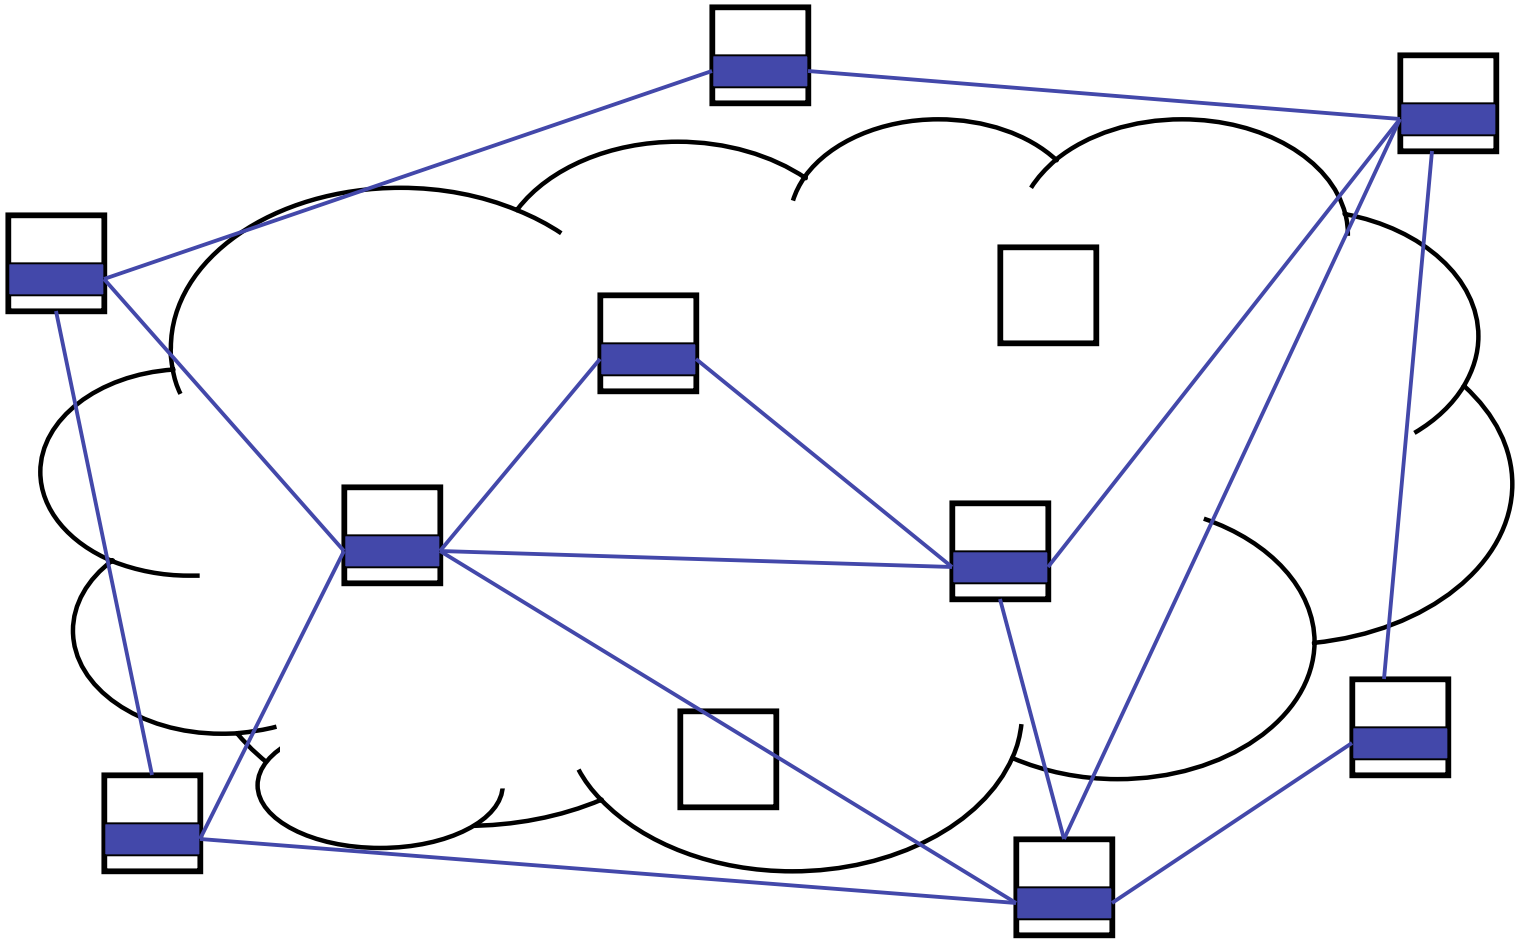

Per-Node View

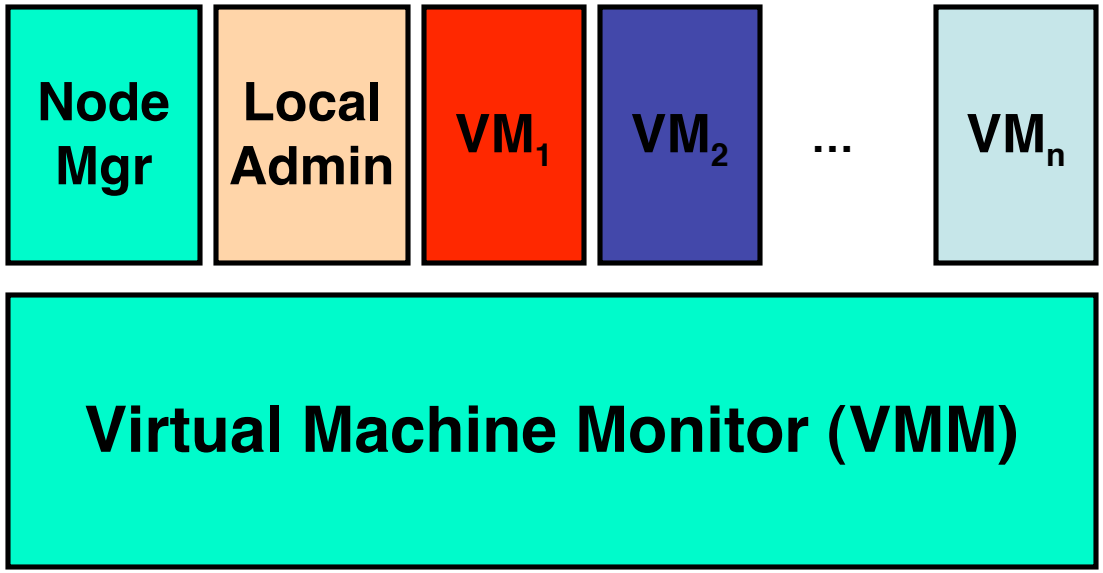

Extending Slices to NLR

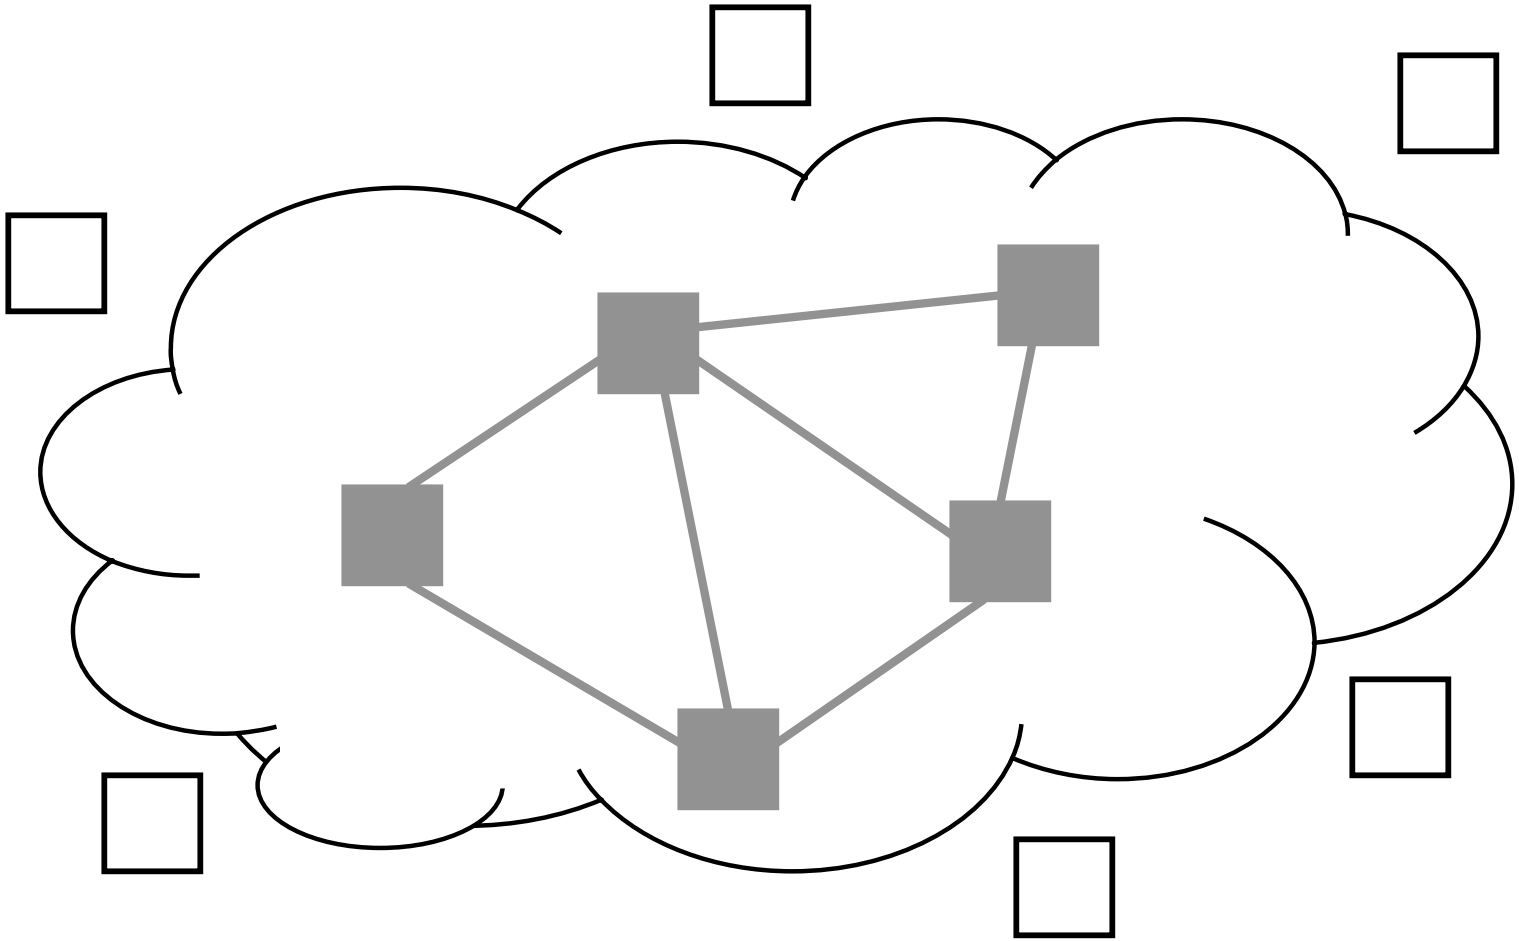

Extending Slices to NLR

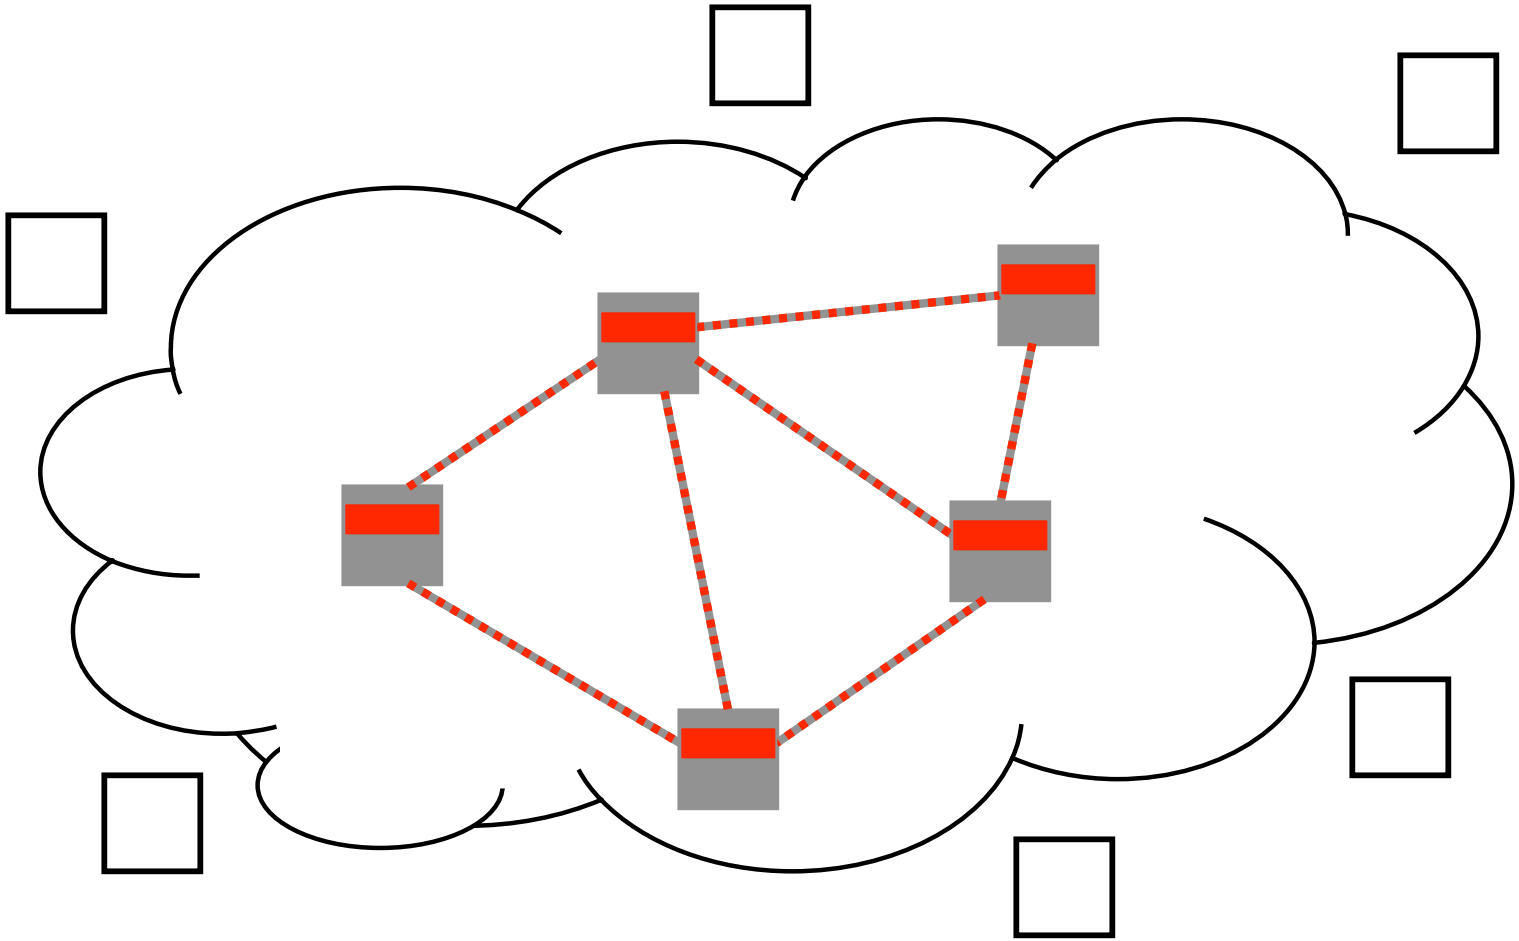

NLR + PlanetLab

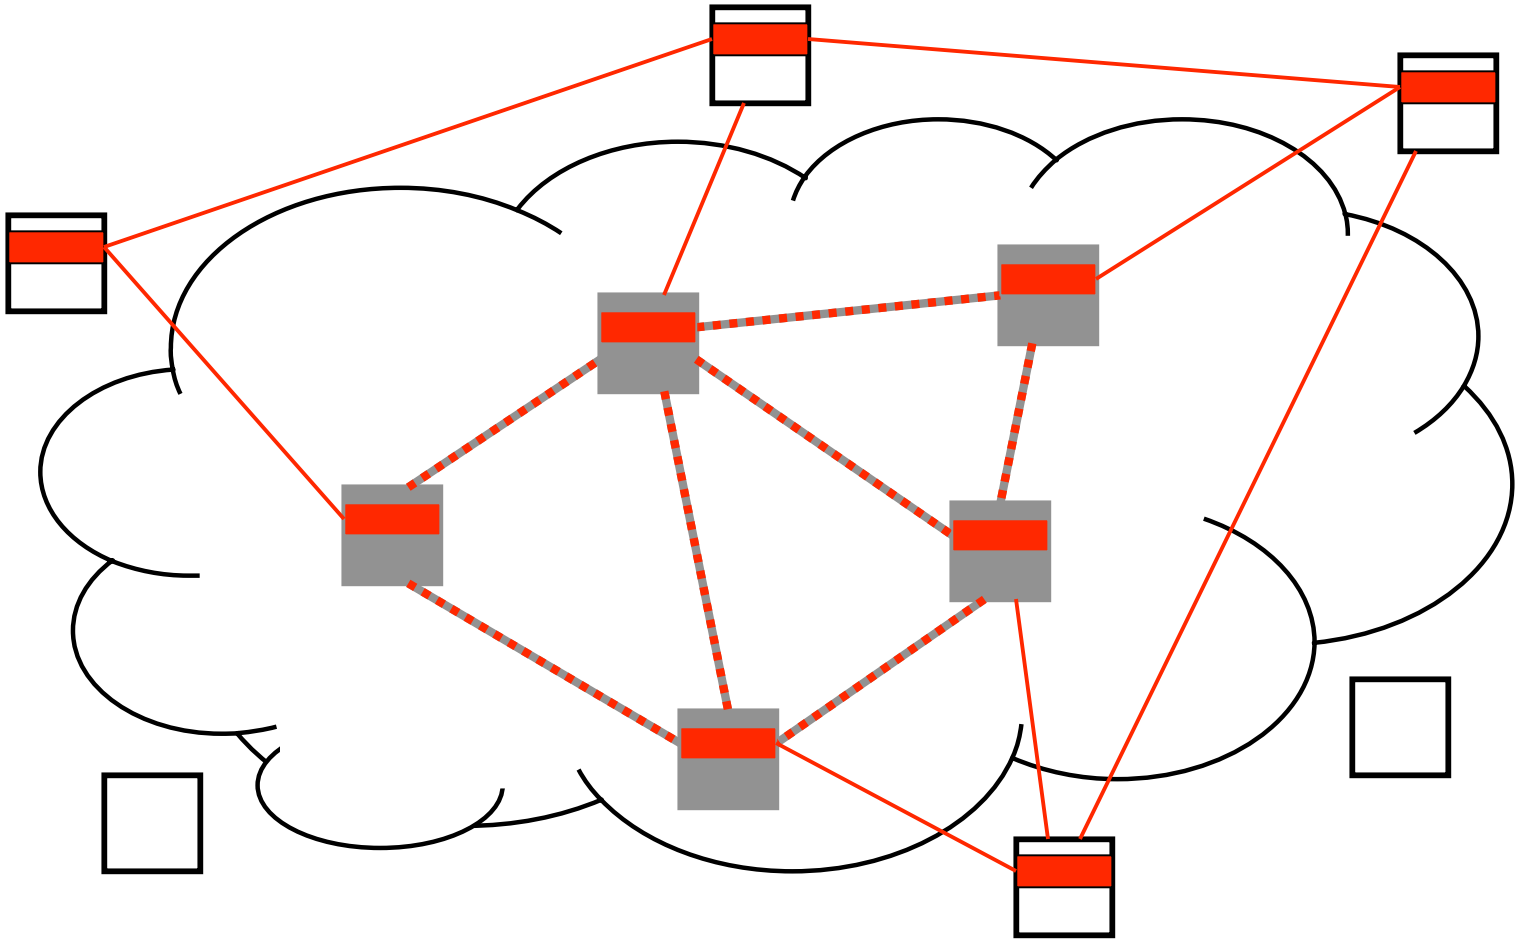

User Opt-in

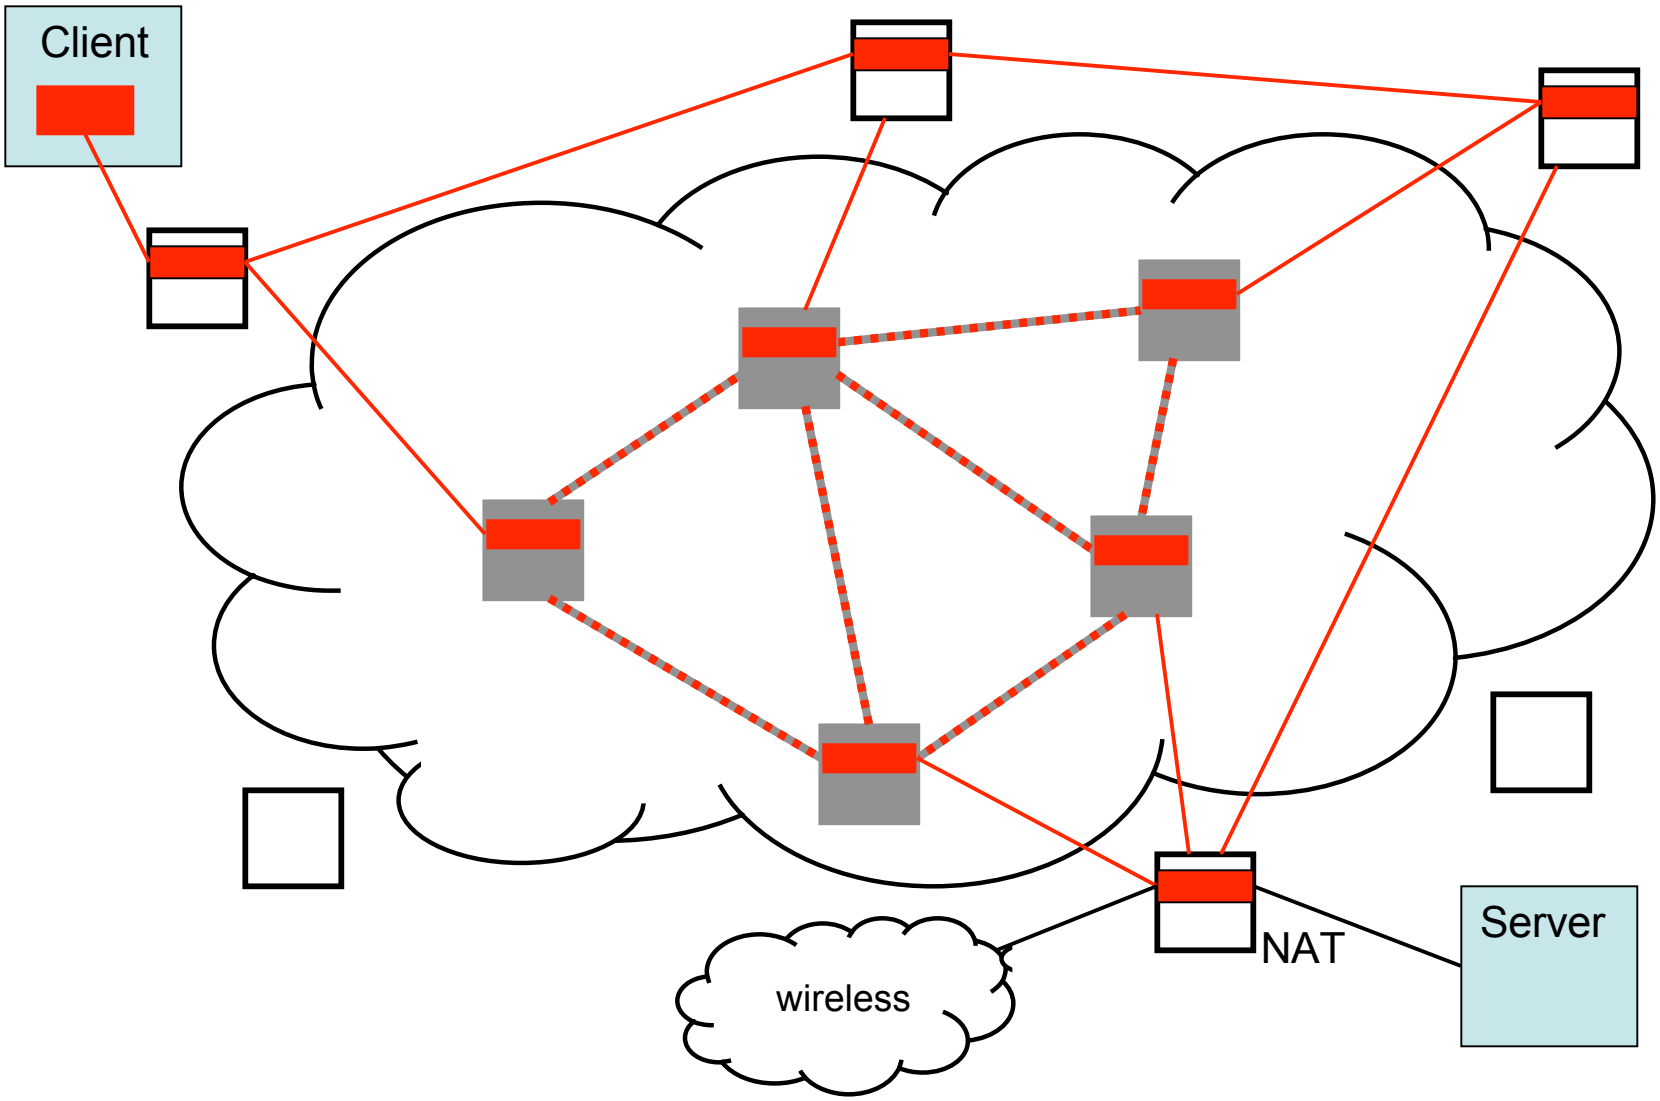

Another View

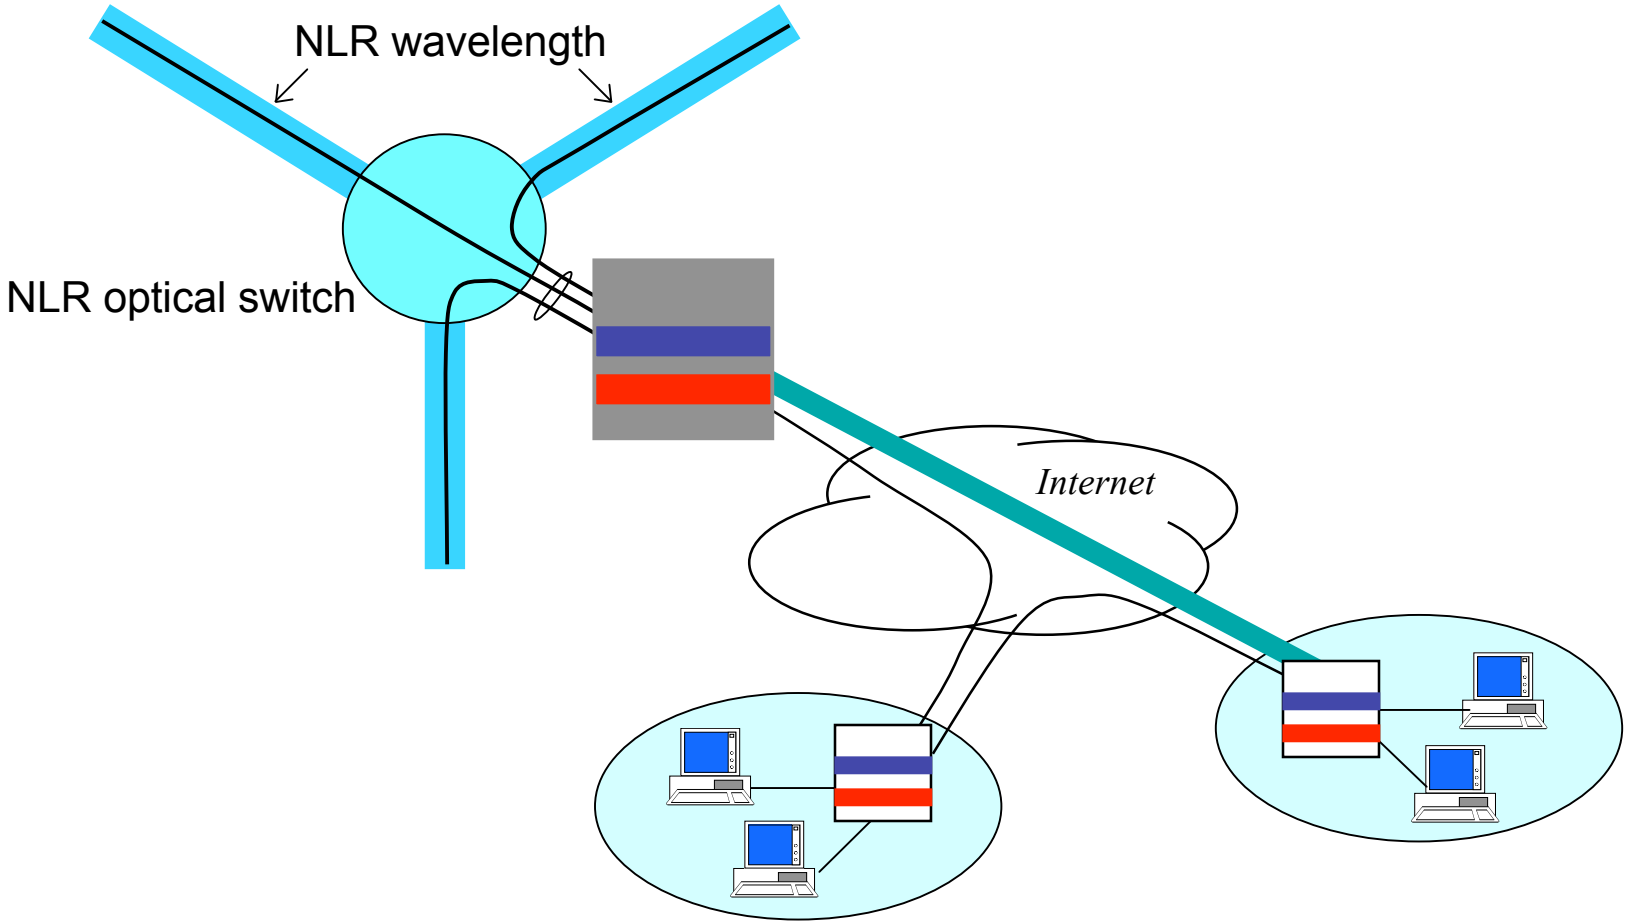

Per-Node View (NLR)

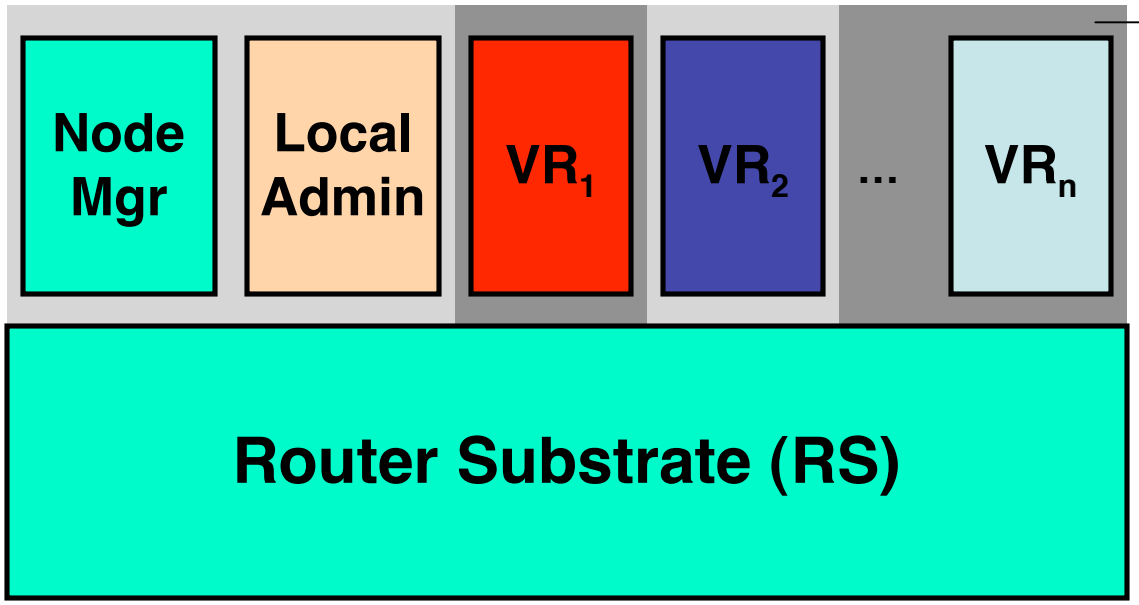

Processing Engine(s)

- COTS PC
- Network Processor
- FPGA

Deployment Story

- Old model
  - global up-take of new technology
  - does not work due to ossification
- New model
  - incremental deployment via user opt-in
  - lowering the barrier-to-entry makes deployment plausible
- Process by which we define the new architecture
  - *purists:* settle on a single common architecture
    - virtualization is a *means*
  - *pluralists:* multiplicity of continually evolving elements
    - virtualization is an *ends*
- What architecture do we deploy?
  - research happens…

Empirical Research Process

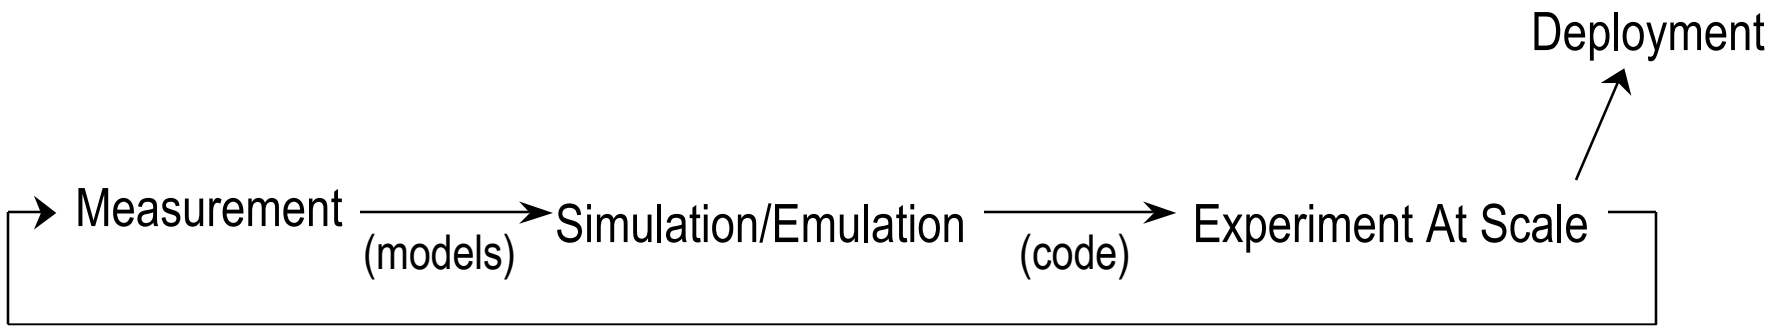

Architectural Thrusts

- Built-in security
  - worm and virus containment, DDoS prevention,…
- Knowledge/Information/Decision Plane
  - managability, fault & anomaly diagnosis, reliability,…
- Network service infrastructure
  - functionality, evolvability, reliability, heterogeneity,…
- Naming and Addressing
  - mobility, ease-of-use, reliability, evolvability,…
- Global sensor network
  - scalability, heterogeneity, mobility,…
- e-Science infrastructure
  - performance, managability, ease-of-use,…
- Optical integration
  - performance, evolvability,…

Success Scenarios

- Create a new network architecture
  - convergence of multiple architectural visions
  - approach to deployment succeeds
  - ready for commercialization
- Meta testbed becomes the new architecture
  - multiple architectures co-exist
  - create a climate of continual re-invention
- Gain new insights and architectural clarity
  - ideas retro-fitted into today's architecture
  - pursuing second path improves the odds of first path succeeding

Acknowledgements

- Tom Anderson, *University of Washington*
- Dan Blumenthal, *UC Santa Barbara*
- David Clark, *Massachusetts Institute of Technology*
- David Culler, *UC Berkeley*
- Guru Parulkar, *National Science Foundation*
- Jennifer Rexford, *Princeton University*
- Scott Shenker, *UC Berkeley*
- David Tennenhouse, *Intel Corporation*
- Jonathan Turner, *Washington University, St. Louis*
- John Wroclawski, *Massachusetts Institute of Technology*

- NSF Workshop on *Overcoming Barriers to Disruptive Innovation in Networking*

www.planet-lab.org/doc/barriers.pdf