GENI Engineering Conference
March 4, 2008

A couple of platforms

(Or: *"Why can't I innovate in my wiring closet?"*)

Nick McKeown
nickm@stanford.edu

The Stanford Clean Slate Program
http://cleanslate.stanford.edu

GENI: How to try new network architectures in my local network?

PlanetLab
- Suitable for distributed systems using overlays
- Not suitable for radical network architectures
- Performance is limited

XORP
- Open-source router code for Linux, FreeBSD, etc.
- Performance is limited
- Point solution

OpenWRT
- Linux for wireless access points and routers

VINI, WUSTL SPP, …
- Designed for large nationwide backbone infrastructure

Outline

NetFPGA

- Open source networking hardware
- 4-ports of 1GE
- Operates at line-rate
- Available for $500 (university)

OpenFlow

- Enabling innovation on campus
- Standard way to control flow-tables in commercial switches and routers
- Being deployed at Stanford
- Consider deploying it at your campus too

NetFPGA

Open source networking hardware

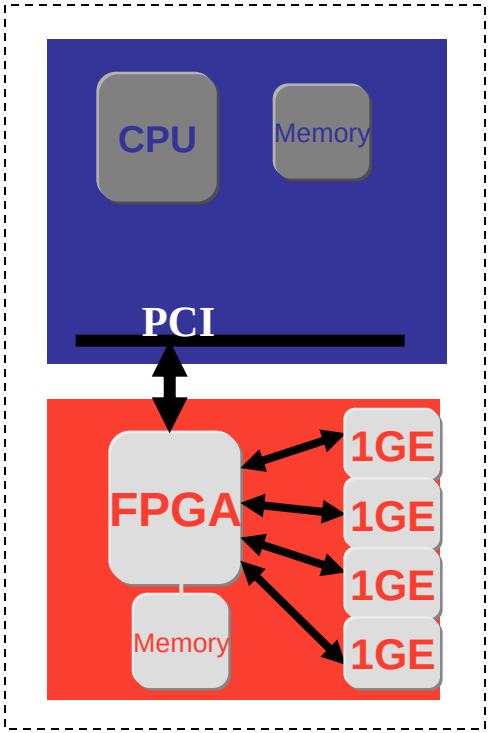

PC with NetFPGA

NetFPGA Board

Rack of 1U servers
96 x 1GE ports

- ✓ Hardware available from 3rd party ~$500 (universities)
- ✓ PCI board, or pre-built system (desktop or rack-mount)

NetFPGA: Usage Models

Types of user

1. Teachers
2. Students
3. Researchers
4. Commercial development

Supported Use Models

1. Router Kit
2. Supported Modular Reference Designs
3. Free Design (or 3rd party design)

All software and designs available under BSD-like license

Usage #1

NetFPGA "Router Kit"

User-space development, 4x1GE line-rate forwarding

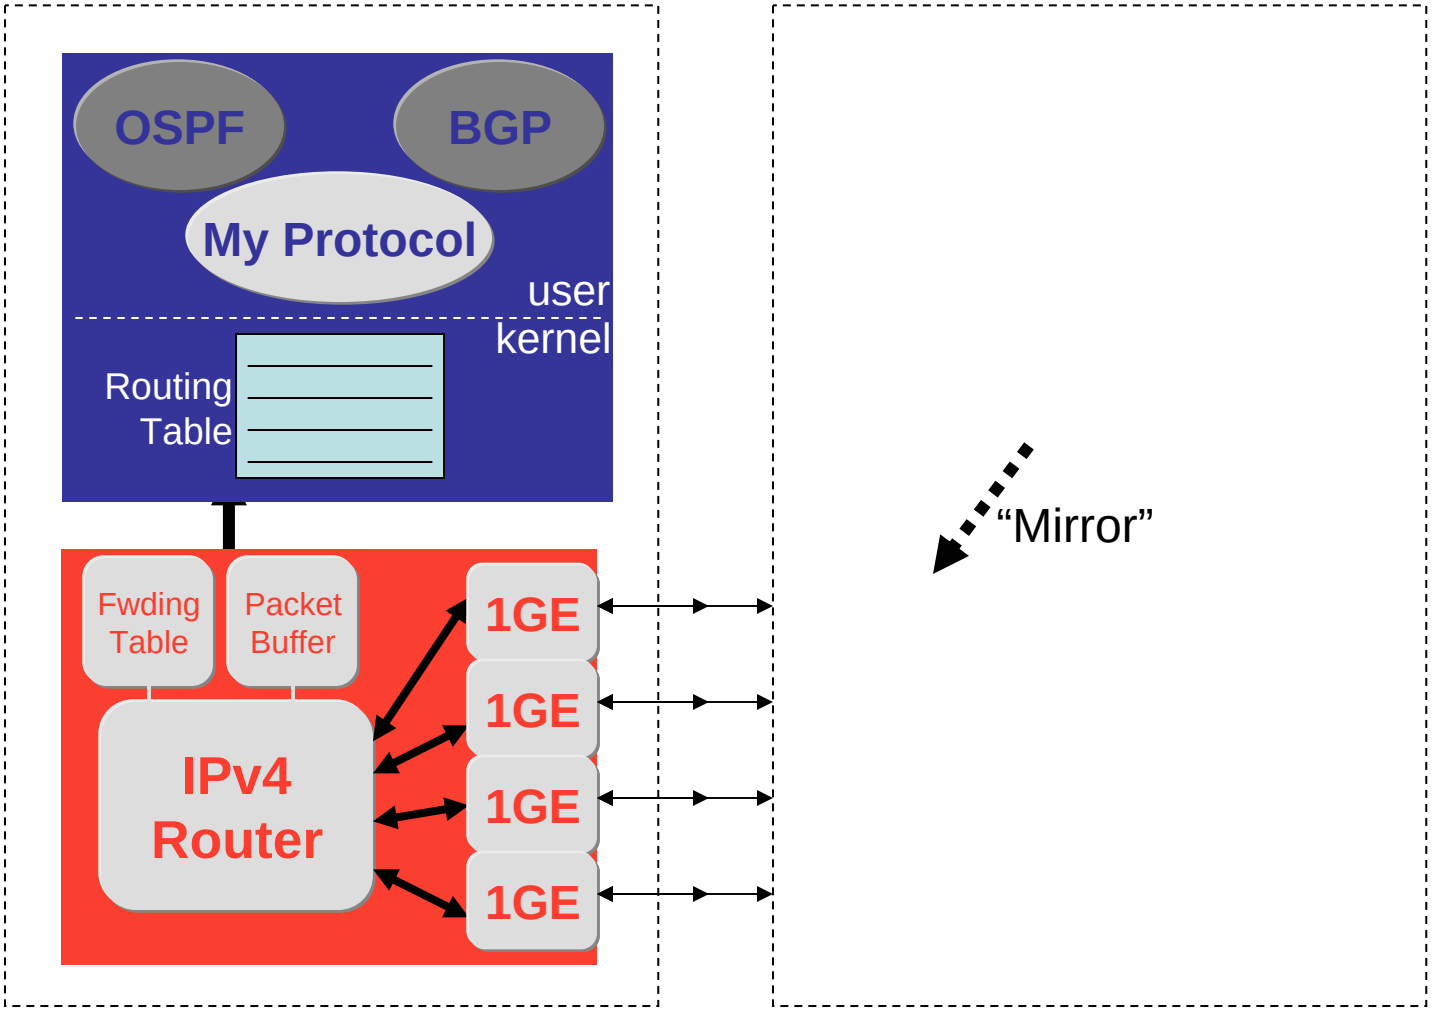

Usage #2

NetFPGA "Reference Router"

Modules in hardware

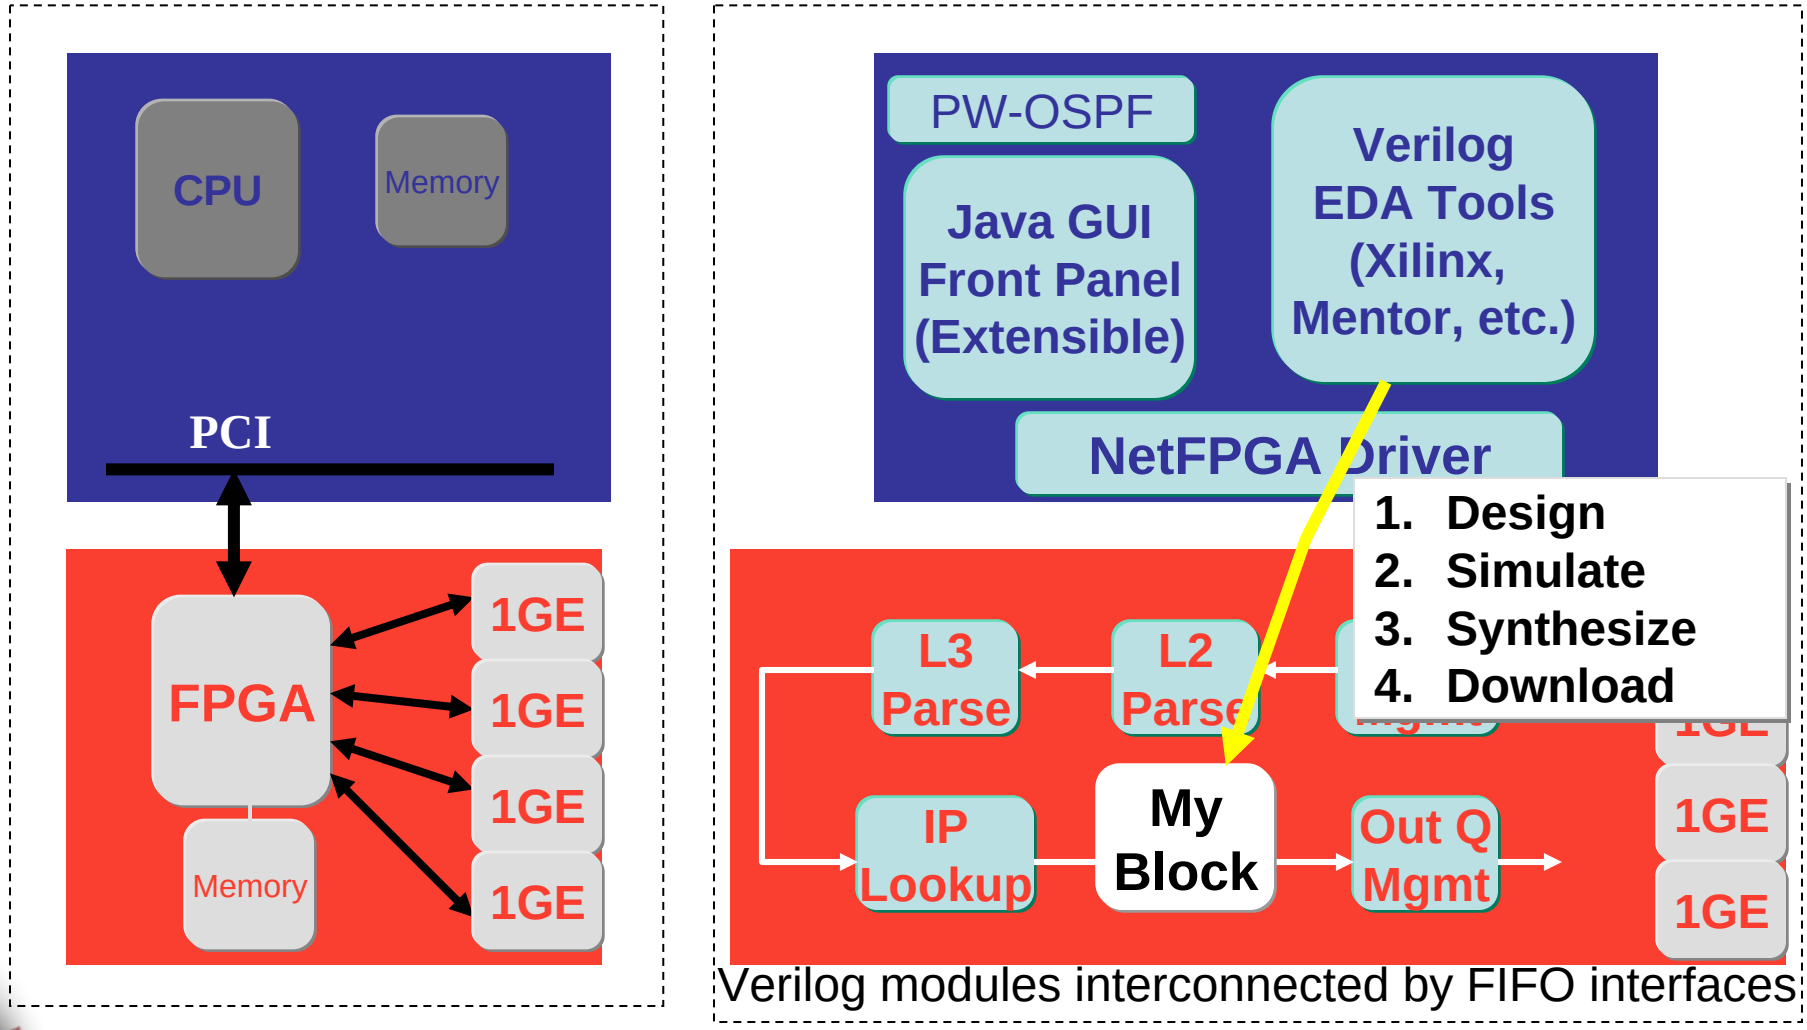

Usage #3

NetFPGA

Free Design

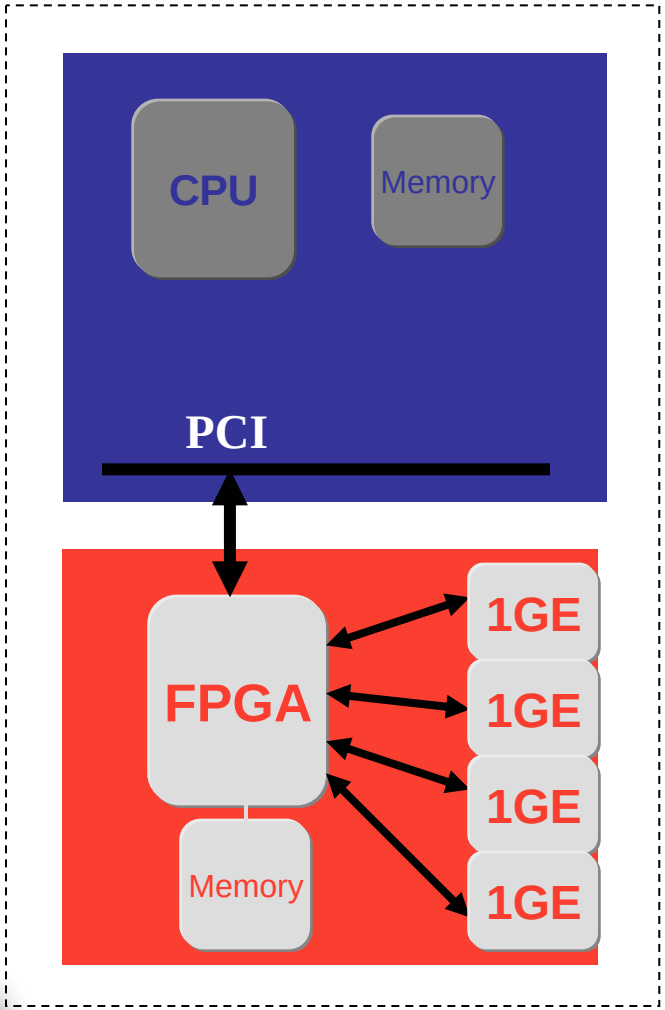

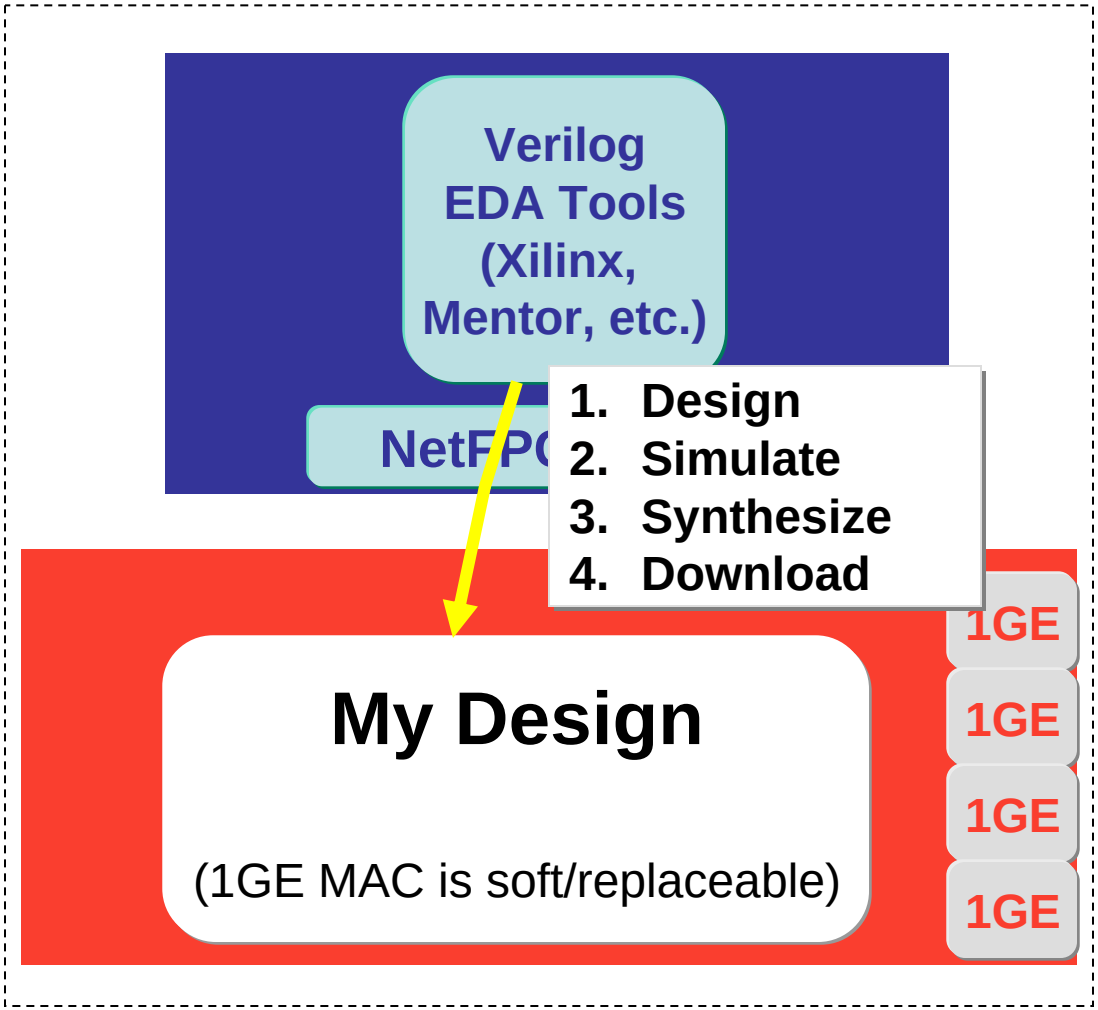

NetFPGA Resources

http://NetFPGA.org

Available

- NetFPGA hardware
- Linux Driver
- Java GUI
- Test & regression software

Courseware

- Undergraduate EE: Ethernet switch
- Masters EE/CS: IPv4 router hardware and software

Tutorials

- 6 tutorials worldwide in 2008
- 1 week Stanford summer camp

NetFPGA Resources (2)

http://NetFPGA.org

Reference Designs

- 4x1GE NIC
- 4x1GE IPv4 Router
- 4x1GE OpenFlow (soon)

User-contributed

- NAT
- IEEE 1588 Timing Sync
- Real-time buffer monitoring
- RCP (Rate Control Protocol)

NetFPGA Funding

http://NetFPGA.org

Launched with funds from NSF EIA Program

Outline

NetFPGA

- Open source networking hardware
- 4-ports of 1GE
- Operates at line-rate
- Available for $500 (university)

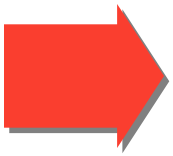

OpenFlow

- Enabling innovation on campus
- Standard way to control flow-tables in commercial switches and routers
- Being deployed at Stanford
- Consider deploying it at your campus too

Innovations in campus wiring closets

Experiments we'd like to do

- Mobility management
- Network-wide energy management
- New naming/addressing schemes
- Network access control

Problem with our network

- Paths are fixed (by the network)
- IP-only
- Addresses dictated by DNS, DHCP, etc
- No means to add our own processing

OpenFlow Switching

1. A way to run experiments in the networks we use everyday.
2. Bring GENI to college campuses.

A “pragmatic” compromise

Allow researchers to run experiments in their network…

…without requiring vendors to expose internal workings.

Basics

An Ethernet switch (e.g. 128-ports of 1GE)

An open protocol to remotely add/remove flow entries

Experimenter's Dream
(Vendor's Nightmare)

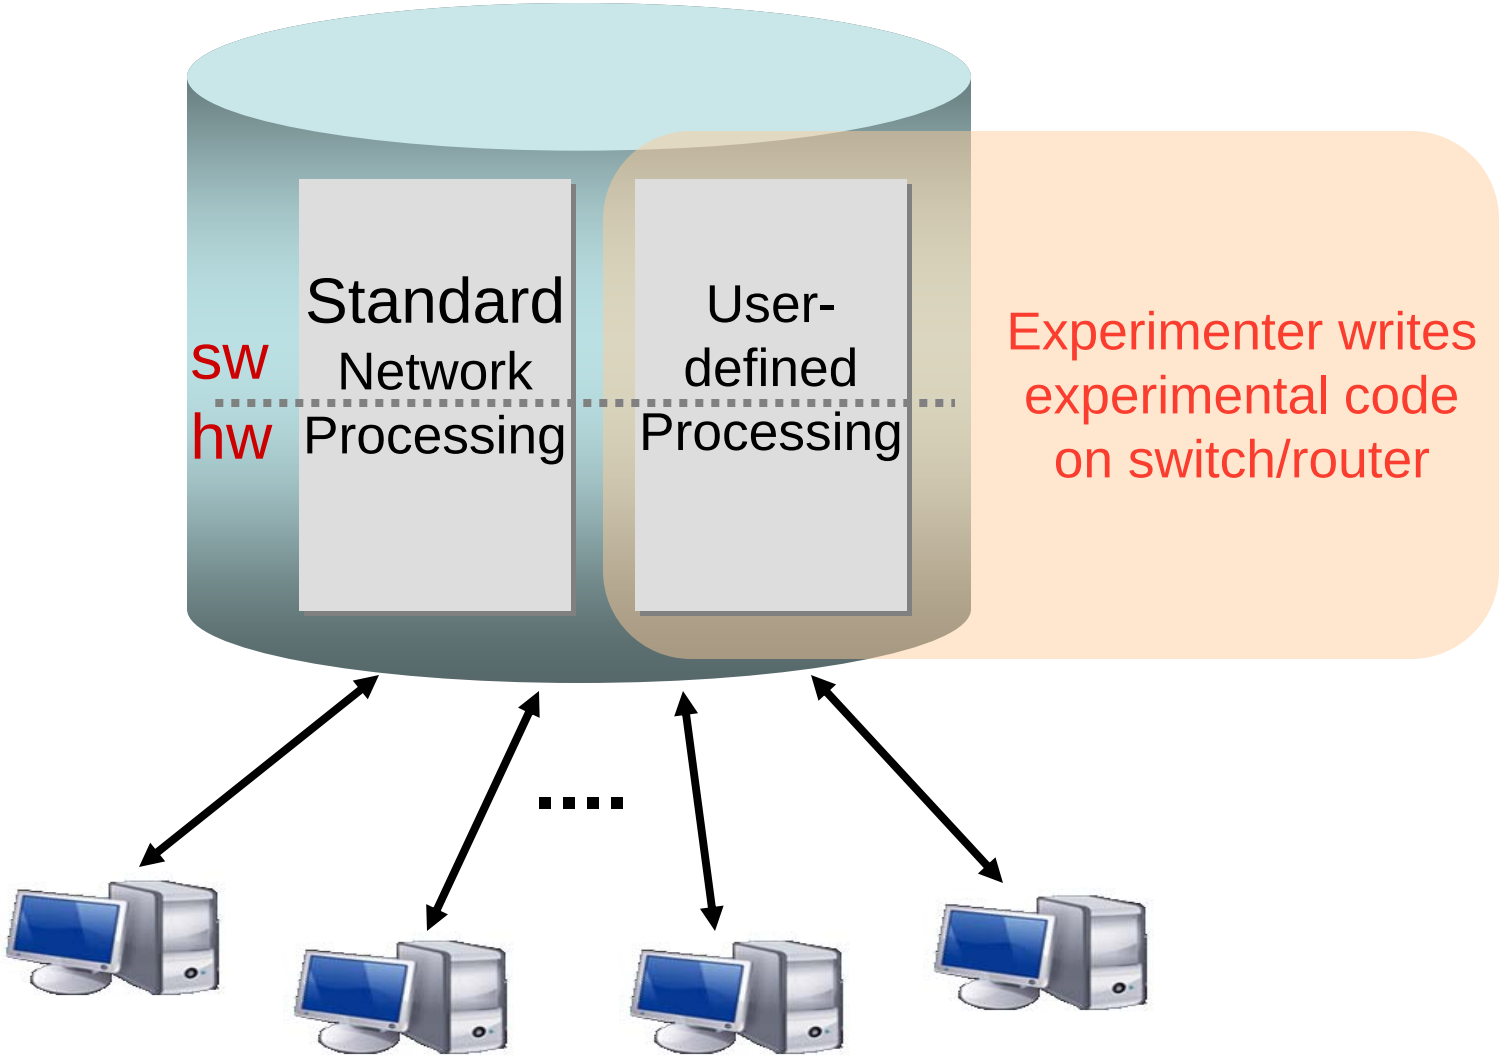

No obvious way

Commercial vendor won't open software and hardware development environment

- Complexity of support
- Market protection and barrier to entry

Hard to build my own

- Prototypes are flakey
- Software only: Too slow
- Hardware/software: Fanout too small (need >100 ports for wiring closet)

Furthermore, we want…

- Isolation: Regular production traffic untouched
- Virtualized and programmable: Different flows processed in different ways
- Equipment we can trust in our wiring closet
- Open development environment for all researchers (e.g. Linux, Verilog, etc).
- Flexible definitions of a flow
  - Individual application traffic
  - Aggregated flows
  - Alternatives to IP running side-by-side
  - …

OpenFlow Switching

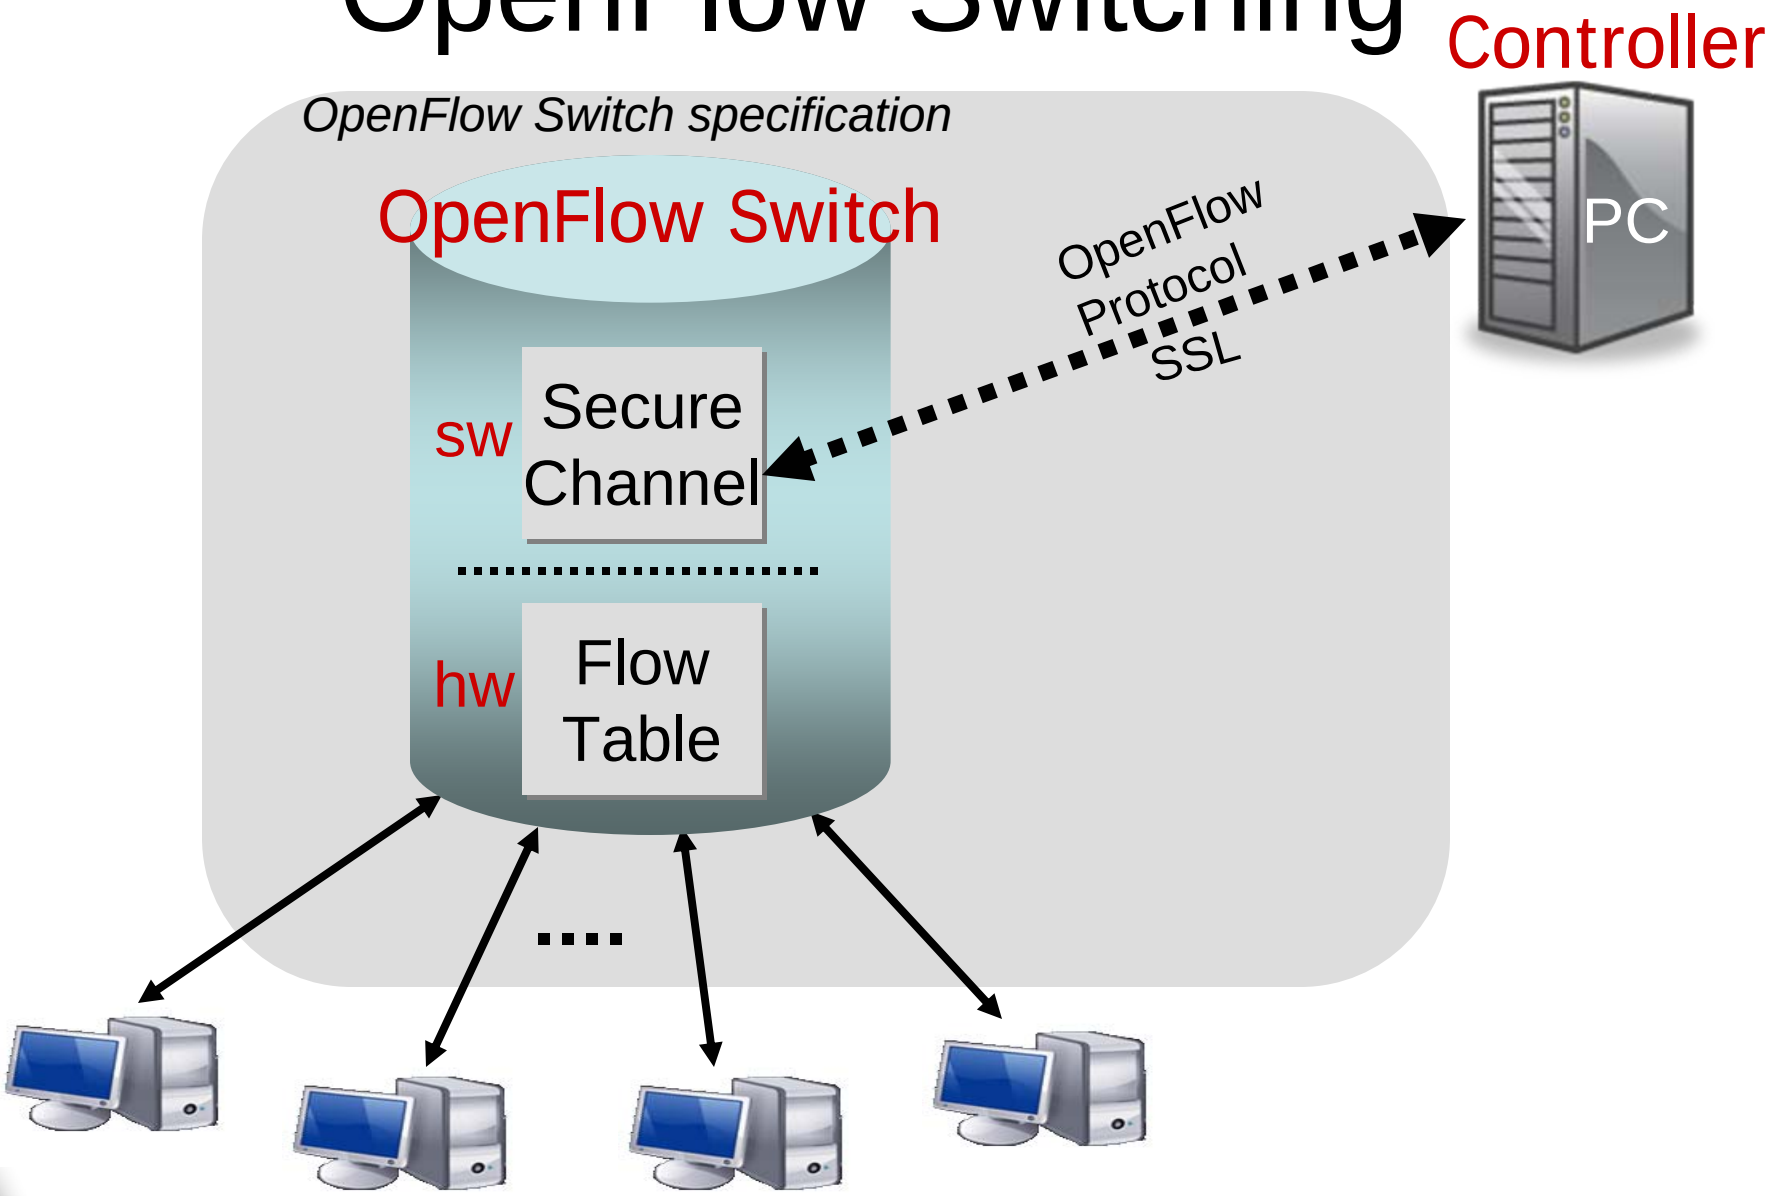

Flow Table Entry

"Type 0" OpenFlow Switch

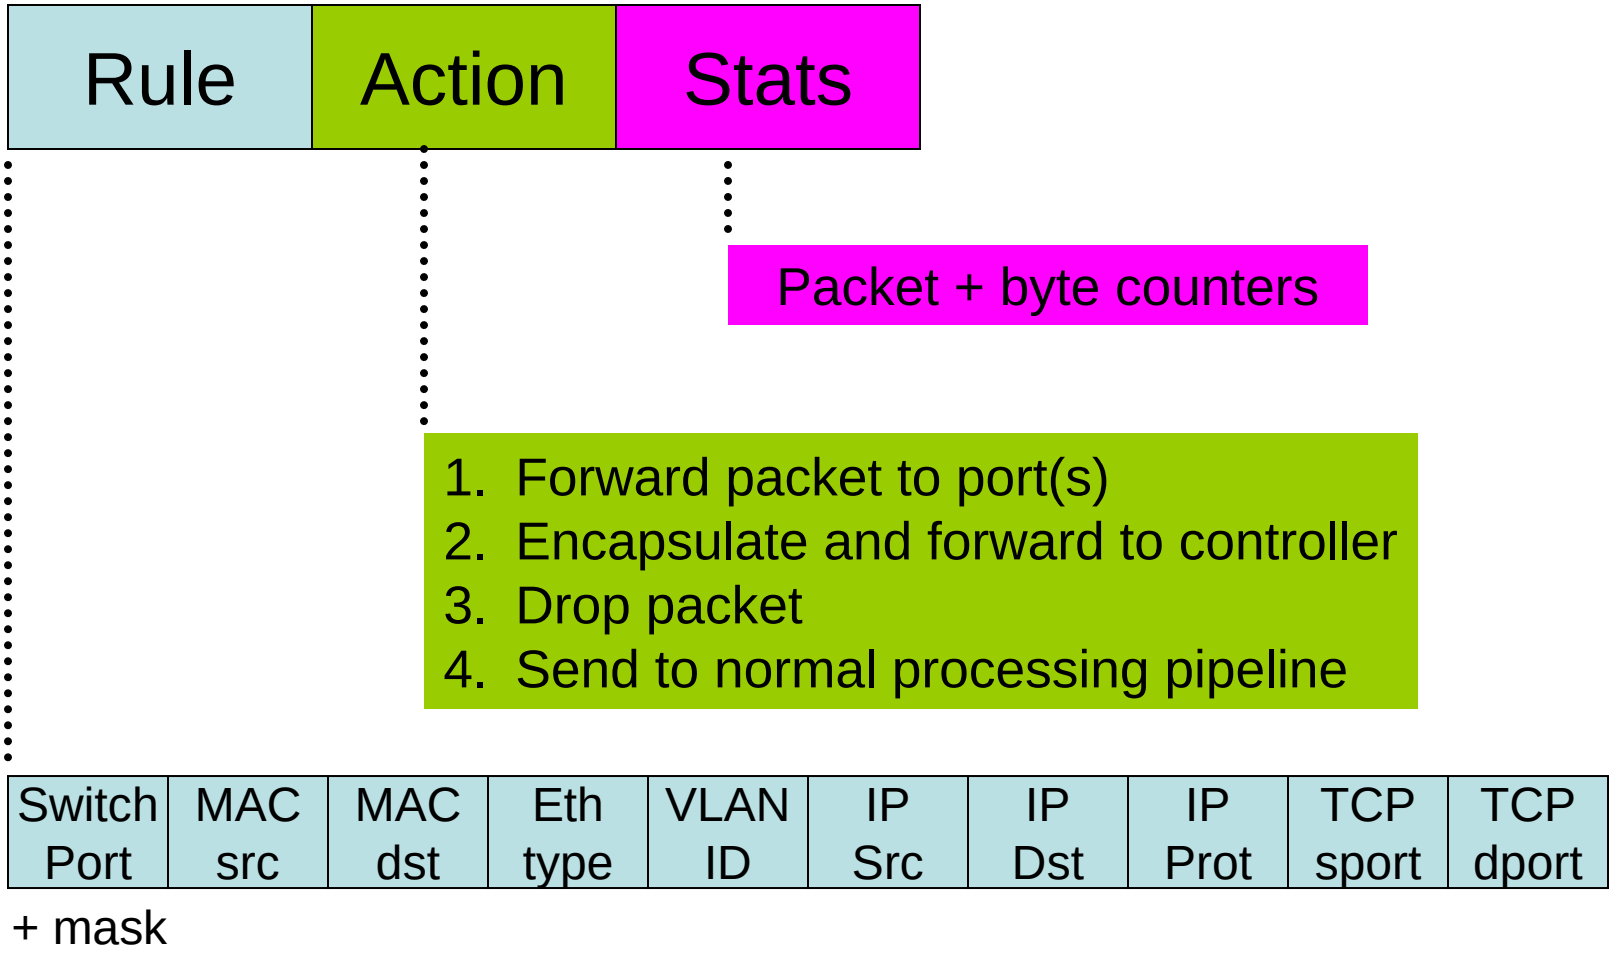

OpenFlow “Type 1”

- Definition in progress
- Additional actions
  - Rewrite headers
  - Map to queue/class
  - Encrypt
- More flexible header
  - Allow arbitrary matching of first few bytes
- Support multiple controllers
  - Load-balancing and reliability

Secure Channel

- SSL Connection, site-specific key
- Controller discovery protocol
- Encapsulate packets for controller
- Send link/port state to controller

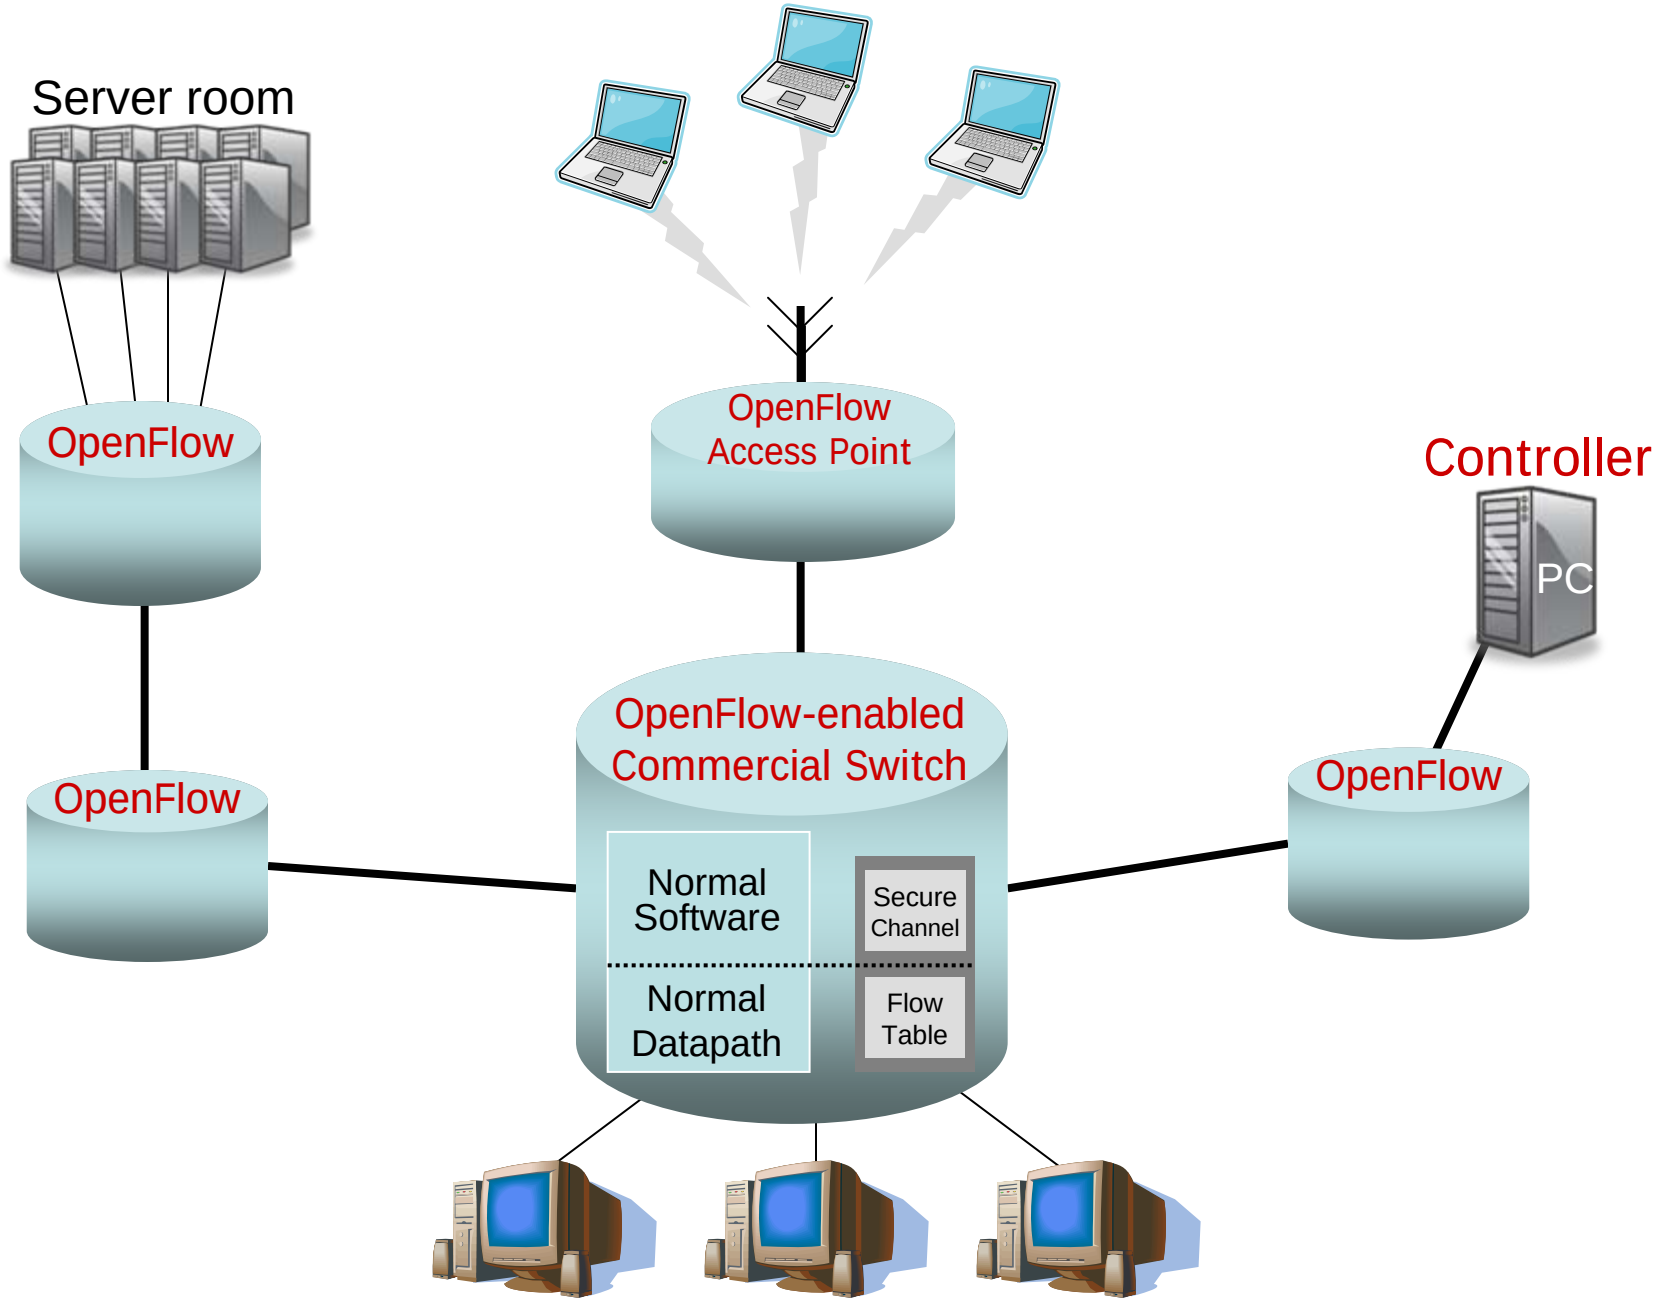
Server room
OpenFlow
OpenFlow Access Point
Controller
PC
OpenFlow
OpenFlow-enabled Commercial Switch
OpenFlow
Normal Software
Secure Channel
Normal Datapath
Flow Table

OpenFlow Usage Models

1. Experiments at the flow level
   - User-defined routing protocols
   - Admission control
   - Network access control
   - Network management
   - Energy management
   - VOIP mobility and handoff
   - …

   - Experiment-specific controllers
   - Static or dynamic flow-entries

2. Experiments at the packet level
   - Slow: Controller handles packet processing
   - Fast: Redirect flows through programmable hardware
   - Modified routers, firewalls, NAT, congestion control…

3. Alternatives to IP

Example Experiment at the flow level

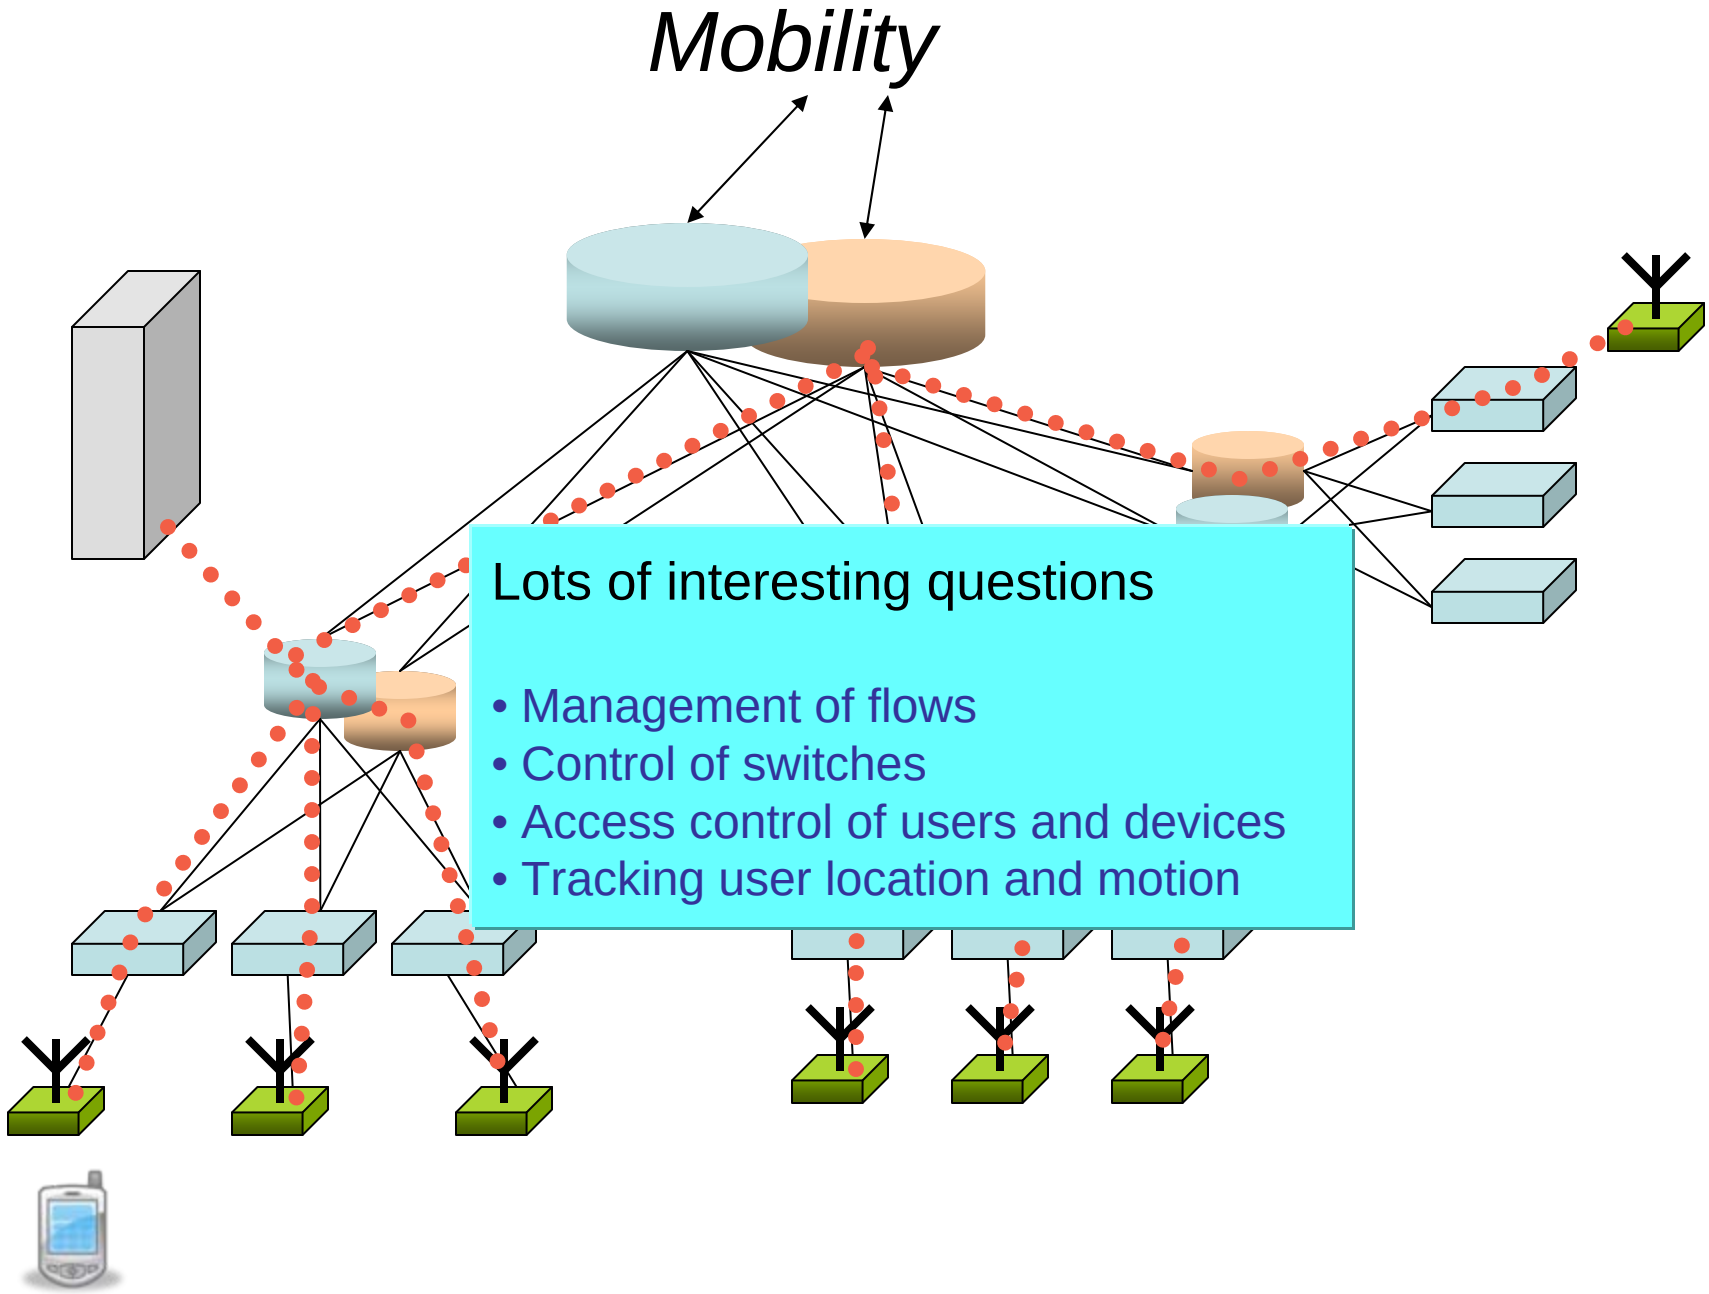

Experiments at the packet level

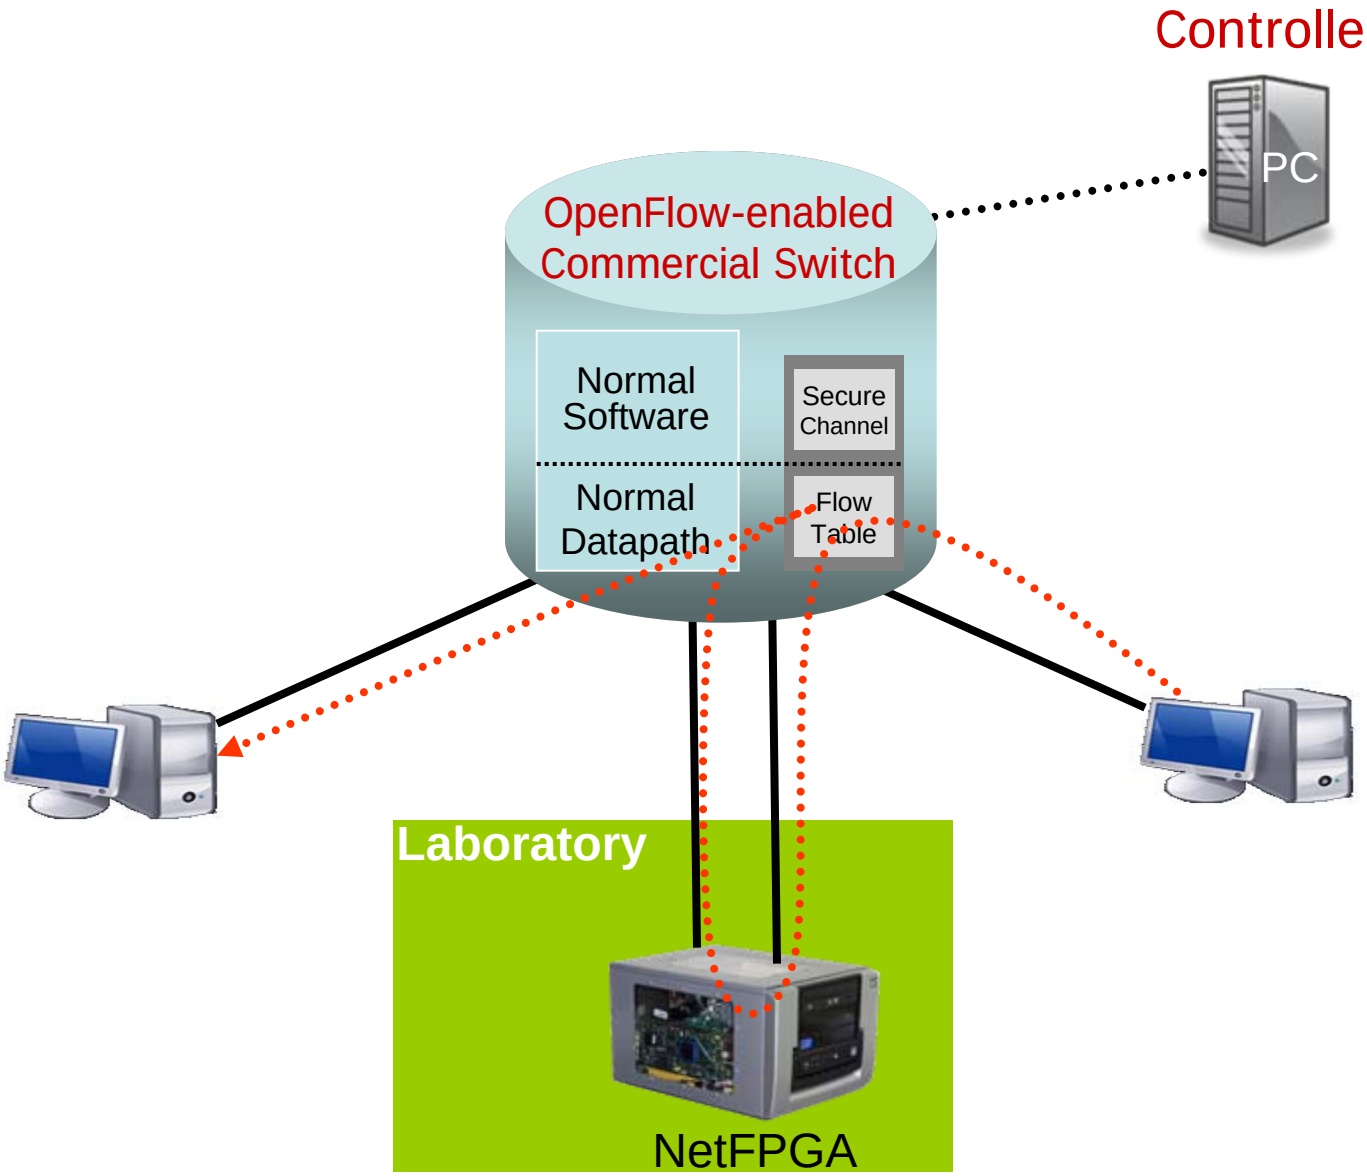

OpenFlow Usage Models

1. Experiments at the flow level
2. Experiments at the packet level

3. Alternatives to IP
   - Flow-table is Layer-2 based
   - e.g. new naming and addressing schemes
   - …

OpenFlow Consortium

http://OpenFlowSwitch.org

Goal: Evangelize OpenFlow to vendors

Free membership for all researchers

Whitepaper, OpenFlow Switch Specification, Reference Designs

Licensing: Free for research and commercial use

OpenFlow: Status

Commercial Ethernet switches and routers

- Working with six vendors to add to existing products
- Expect OpenFlow “Type 0” to be available in 2008-09

Reference switches

- Software: Linux and OpenWRT (for access points)
- Hardware: NetFPGA (line-rate 1GE; available soon)
- Working on low-cost 48-port 1GE switch based on Broadcom reference design

Reference controller

- Simple test controller
- NOX controller (Martin Casado; available soon)

Deployment at Stanford

Stanford Computer Science Department

Gates Building

~1,000 network users

23 wiring closets

Stanford Center for Integrated Systems (EE)

Paul Allen Building

~200 network users

6 wiring closets

Working with HP Labs and Cisco on deployment

If you are interested in deploying OpenFlow on your campus…

Please contact me!

nickm@stanford.edu

http://OpenFlowSwitch.org