





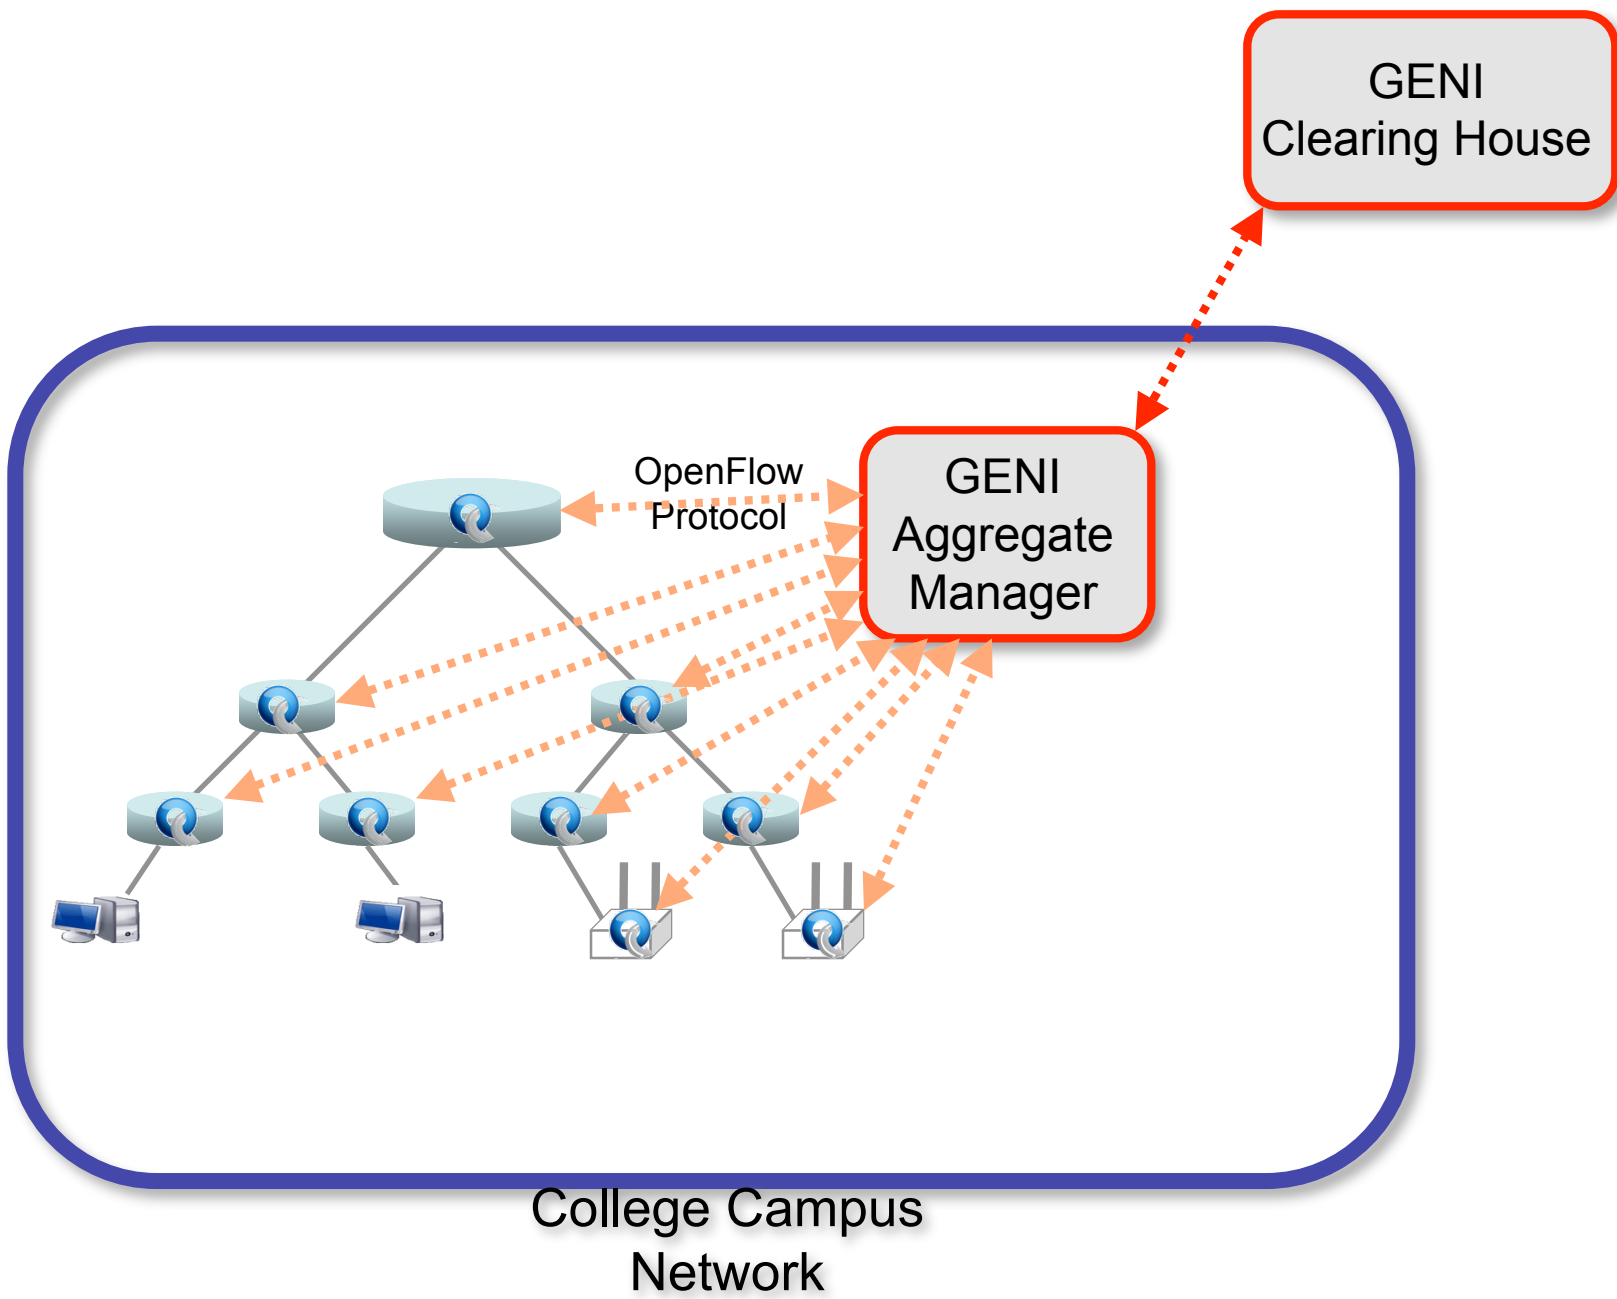
Enterprise GENI

Nick McKeown
nickm@stanford.edu

Stanford OpenFlow team: Guido Appenzeller, Glen Gibb, David Underhill,
David Erickson, Adam Covington, Brandon Heller, Rob Sherwood,
Masayoshi Kobayashi, Guru Parulkar, Srinivasan Seetharaman, Yiannis Yiakoumis







Staged Approach

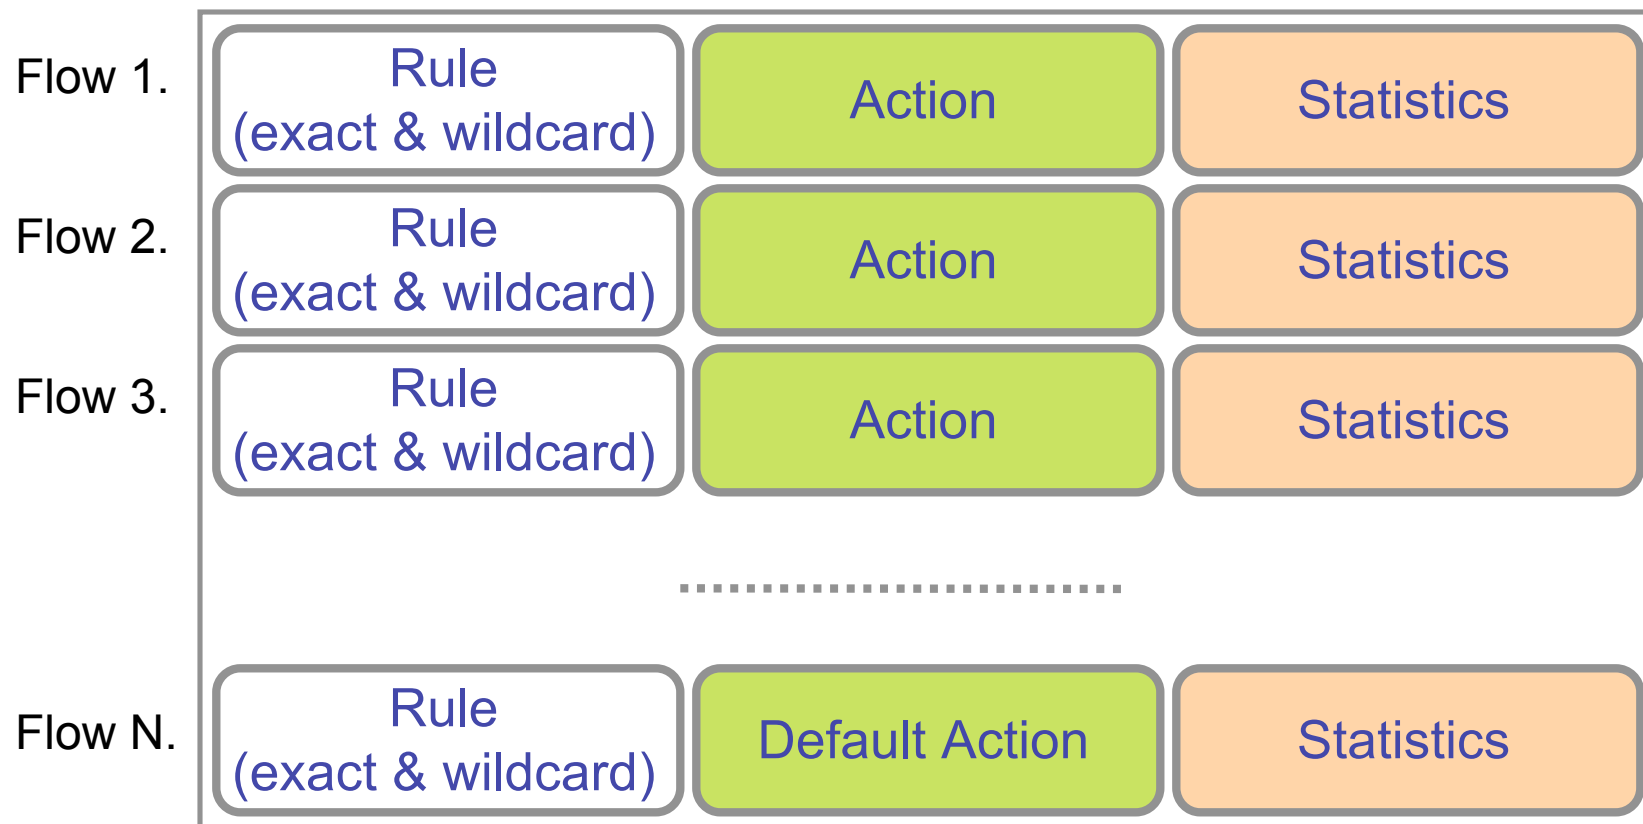
1. Define OpenFlow feature
2. Add OpenFlow to commercial switches and APs
3. Deploy at Stanford
4. Integrate with GENI Clearing House
5. Deploy on many college campus networks
6. We all create lots of open-source software so researchers can build on each other's work

(We're part-way into Stage 2)


OpenFlow Basics

OpenFlow Basics (1)

Exploit the flow table in switches, routers, and chipsets




OpenFlow Basics (2)



OpenFlow Basics (3)

Controller



OpenFlow Switch specification



sw


Secure Channel

hw

Flow Table

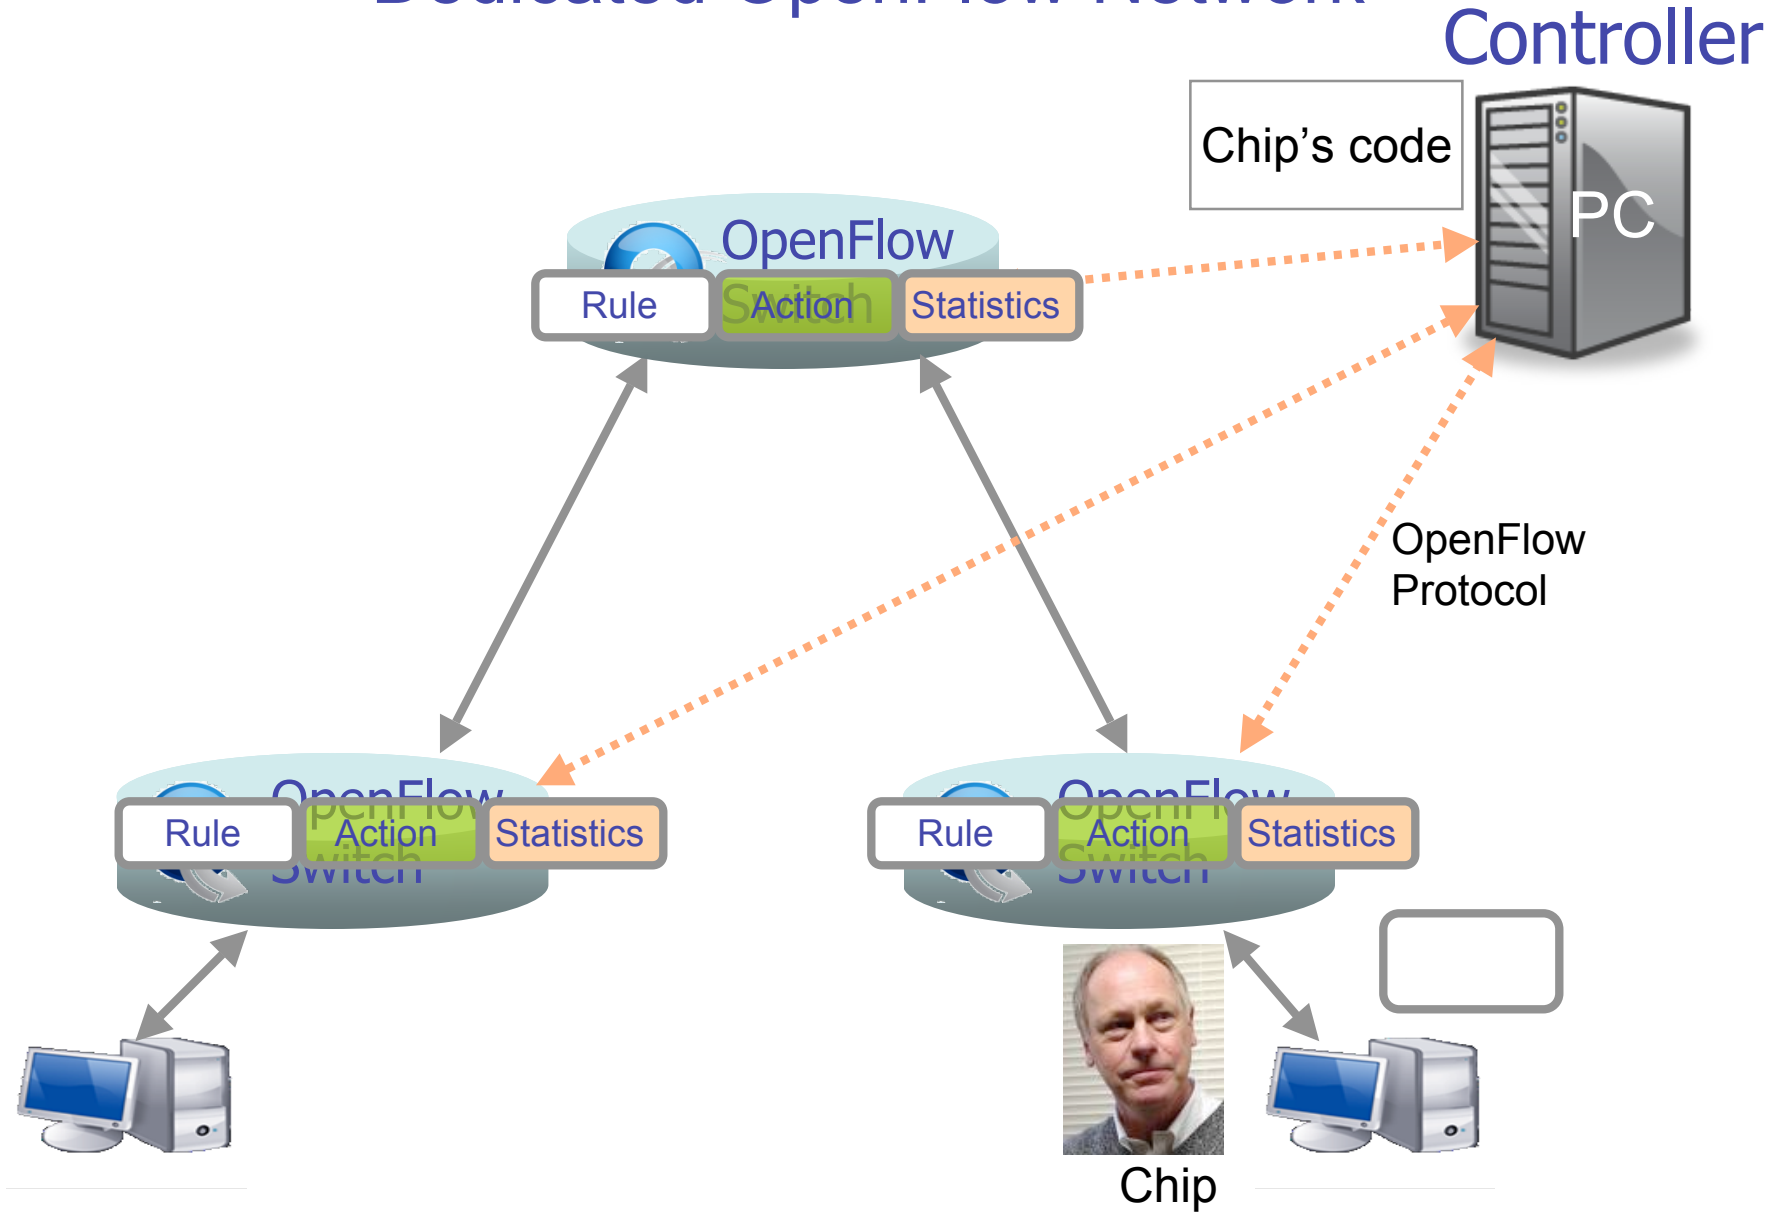
OpenFlow Protocol
SSL

Add/delete flow entries
Encapsulated packets
Controller discovery



OpenFlow Usage

Dedicated OpenFlow Network



Usage examples

Chip's code:

- Static “VLANs”
- His own new routing protocol: unicast, multicast, multipath, load-balancing
- Network access control
- Home network manager
- Mobility manager
- Energy manager
- Packet processor (in controller)
- IPvChip
- Network measurement and visualization
- ...

<http://OpenFlowSwitch.org>



[Documents](#) [Downloads](#) [Wiki](#) [Community](#) [Blog](#) [About](#)



Enable Innovation in your Network.

OpenFlow is an open standard that allows you to run experimental protocols in production networks. It is in the process of being implemented by major switch vendors and used today by universities to deploy innovative networking technology in their campus networks.



Learn More

white paper, video,
sample deployments



Get Started

downloads, tutorial
quickstart guide



Develop

reference implementations
specs and test suites

Let's see an example...

HP



- Experimental feature on ProCurve 5400-series
- 144-ports of 1GE, hardware forwarding
- OpenFlow added by HP Labs and ProCurve group
- In 23 wiring closets in CS Building at Stanford



Praveen
Yalagandula



Jean
Tourrilhes



Sujata
Banerjee



Rick
McGeer



Charles
Clark

NEC



- Experimental feature on IP8800 series router
- 24-ports of 1GE, 2-ports of 10GE, hardware forwarding
- OpenFlow added by NEC team in Japan
- NEC announced plans for OpenFlow products
- Deployed at Stanford and in JGN2plus in Tokyo



Atsushi
Iwata



Hideyuki
Shimonishi



Jun
Suzuki



Masanori
Takashima



Nobuyuki
Enomoto



Philavong
Minaxay



Shuichi
Saito
NEC/NICT



Tatsuya
Yabe



Yoshihiko
Kanaumi
NEC/NICT

Juniper



- OpenFlow added to Junos SDK
- First platform: MX-480 carrier class Ethernet
- 24-ports 10GE or 240-ports 1GE
- Hardware forwarding
- Deployed in Internet2 in NY and at Stanford



Umesh
Krishnaswamy



Michaela
Mezo



Parag
Bajaria



James
Kelly



Bobby
Vandalore

Cisco



- Experimental feature on Catalyst 6509
- Software forwarding
- Deployed at Stanford



Pere
Monclus



Sailesh
Kumar




Flavio
Bonomi

Nicira

Controller



- Created NOX controller
- Available at <http://NOXrepo.org> (GPL)
- Deployed at Stanford



Martin
Casado



Scott
Shenker



Teemu
Koponen




Natasha
Gude




Justin
Pettit

Using Enterprise GENI for research

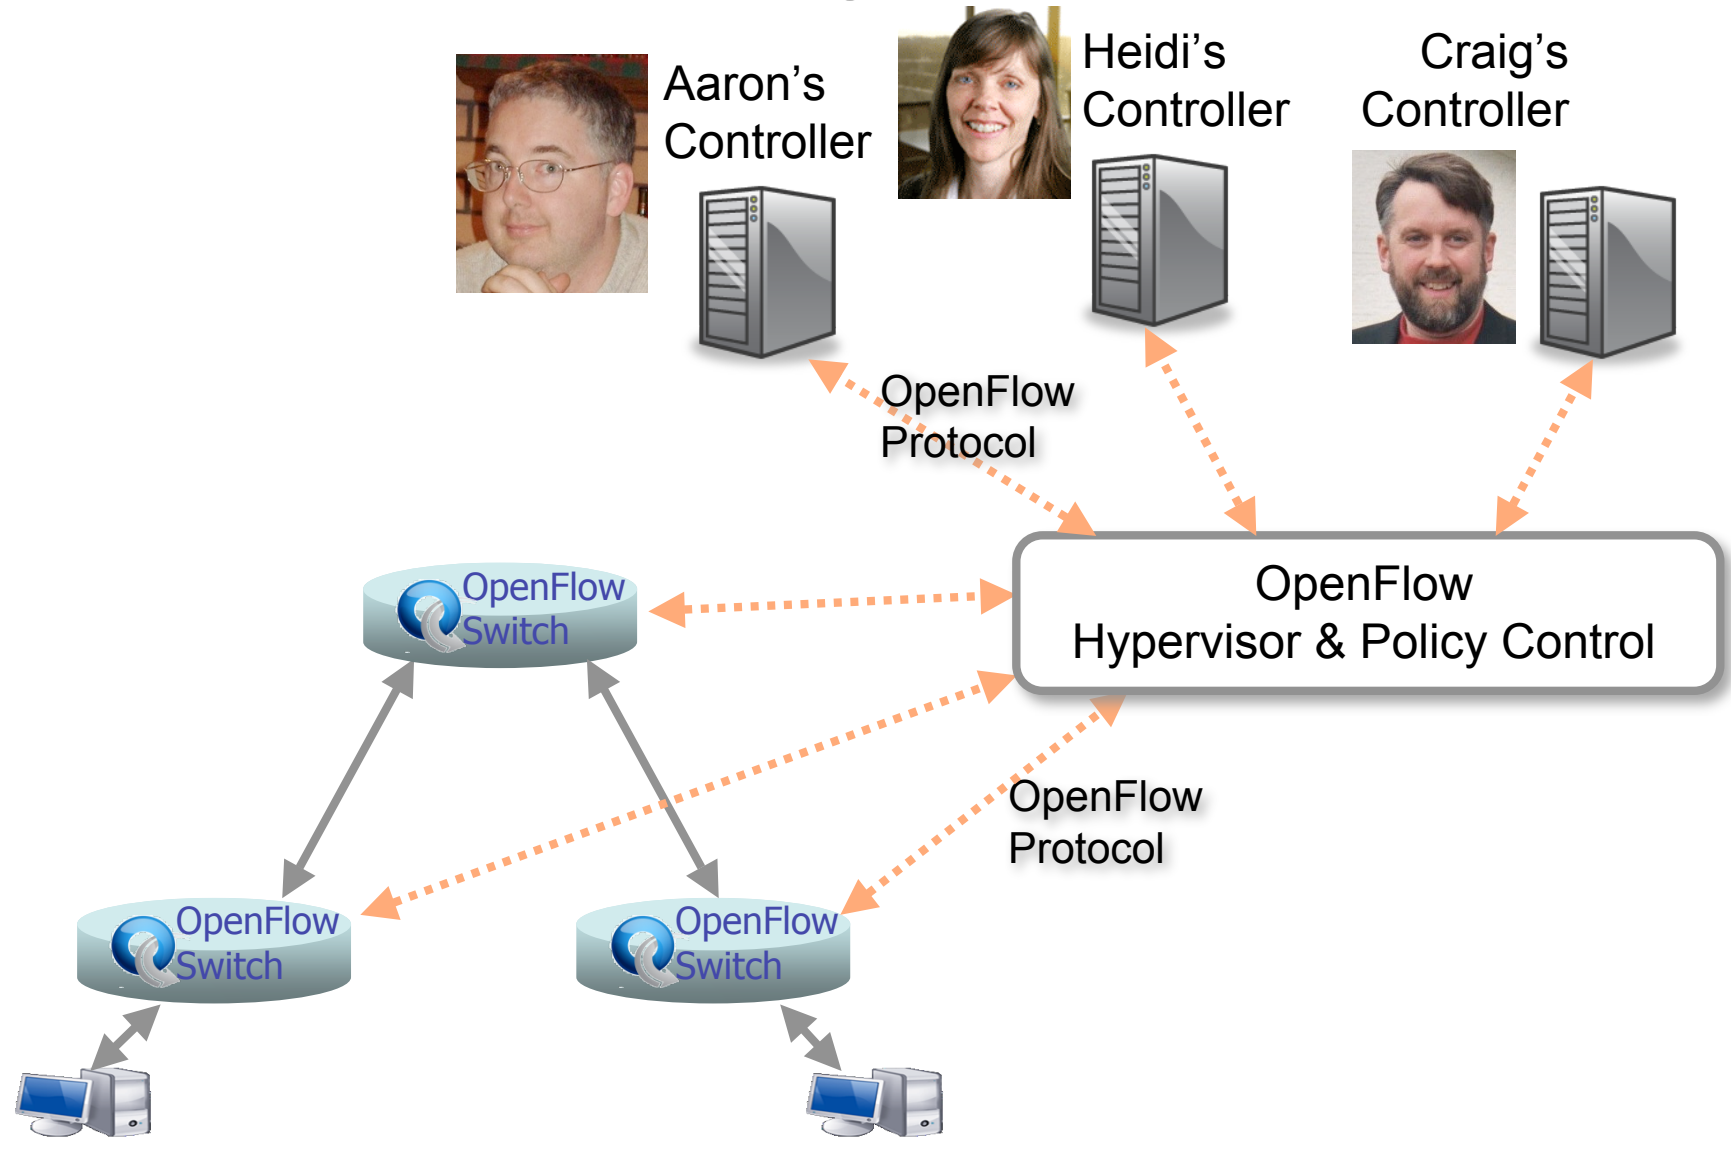
Step 1: Separate VLANs for Production and Research Traffic



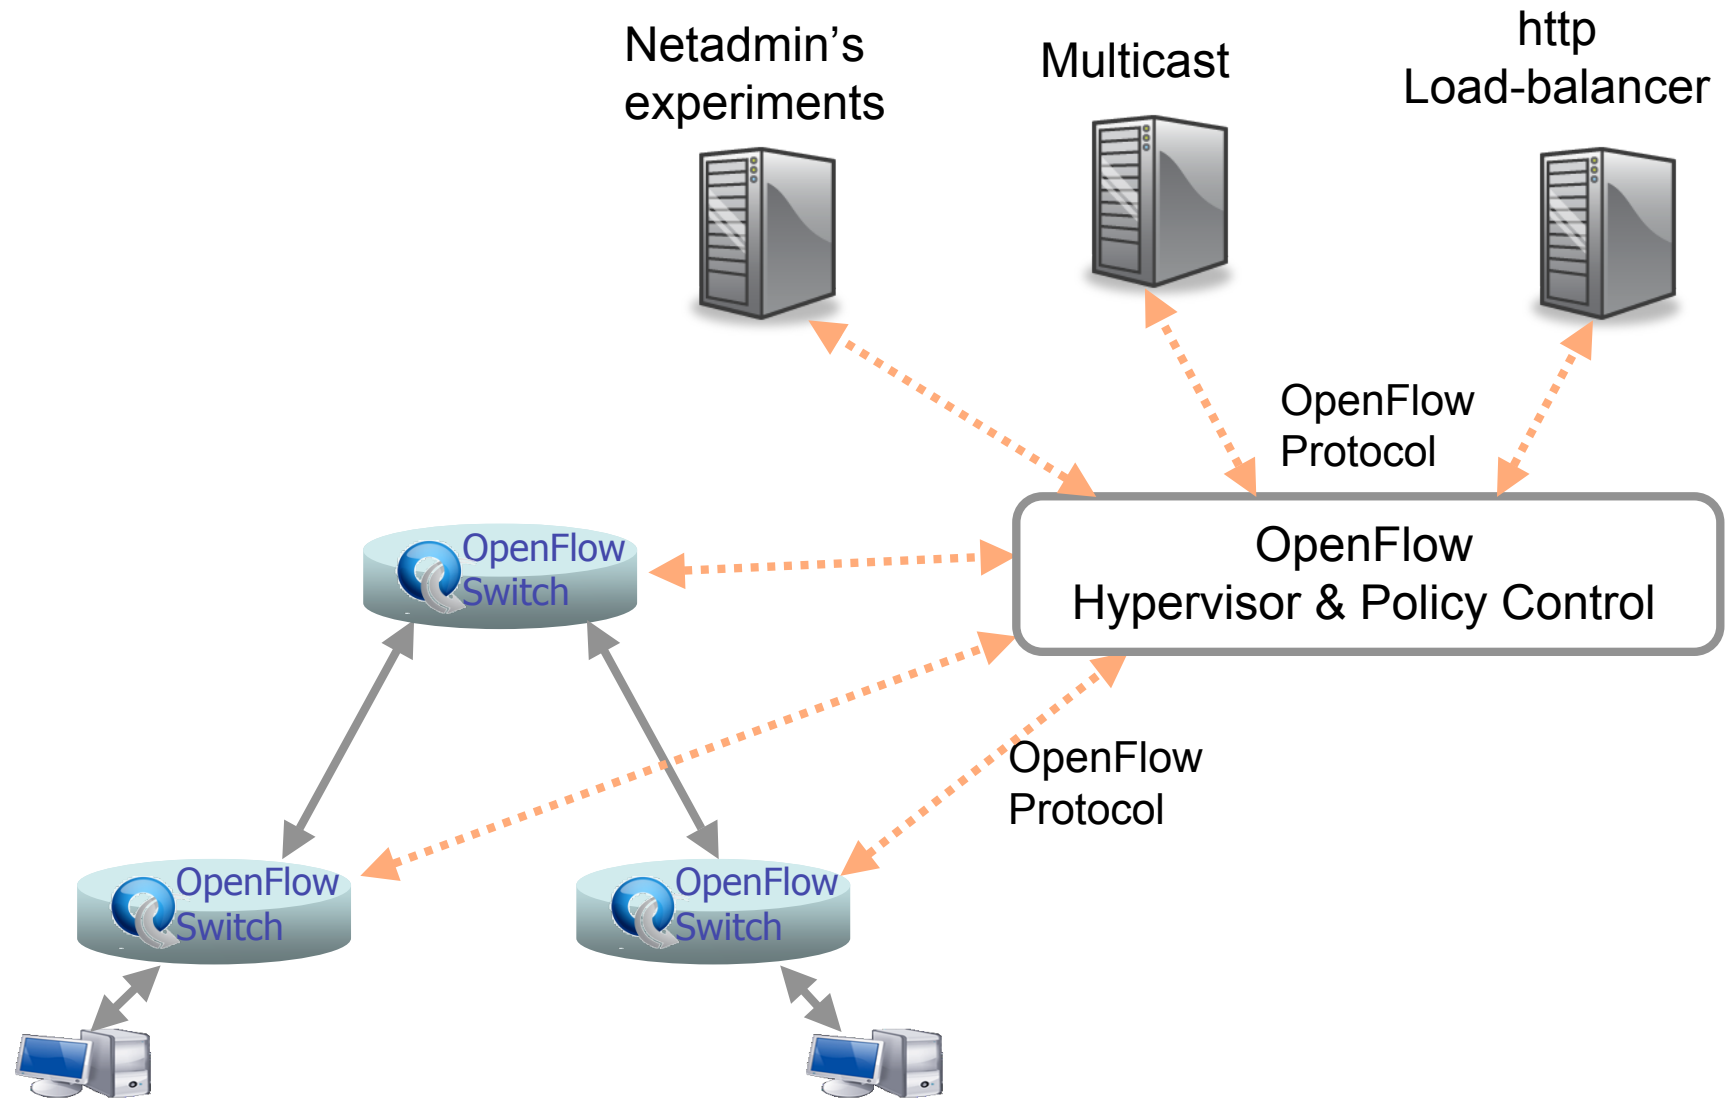
Step 2: Virtualize OpenFlow Switch



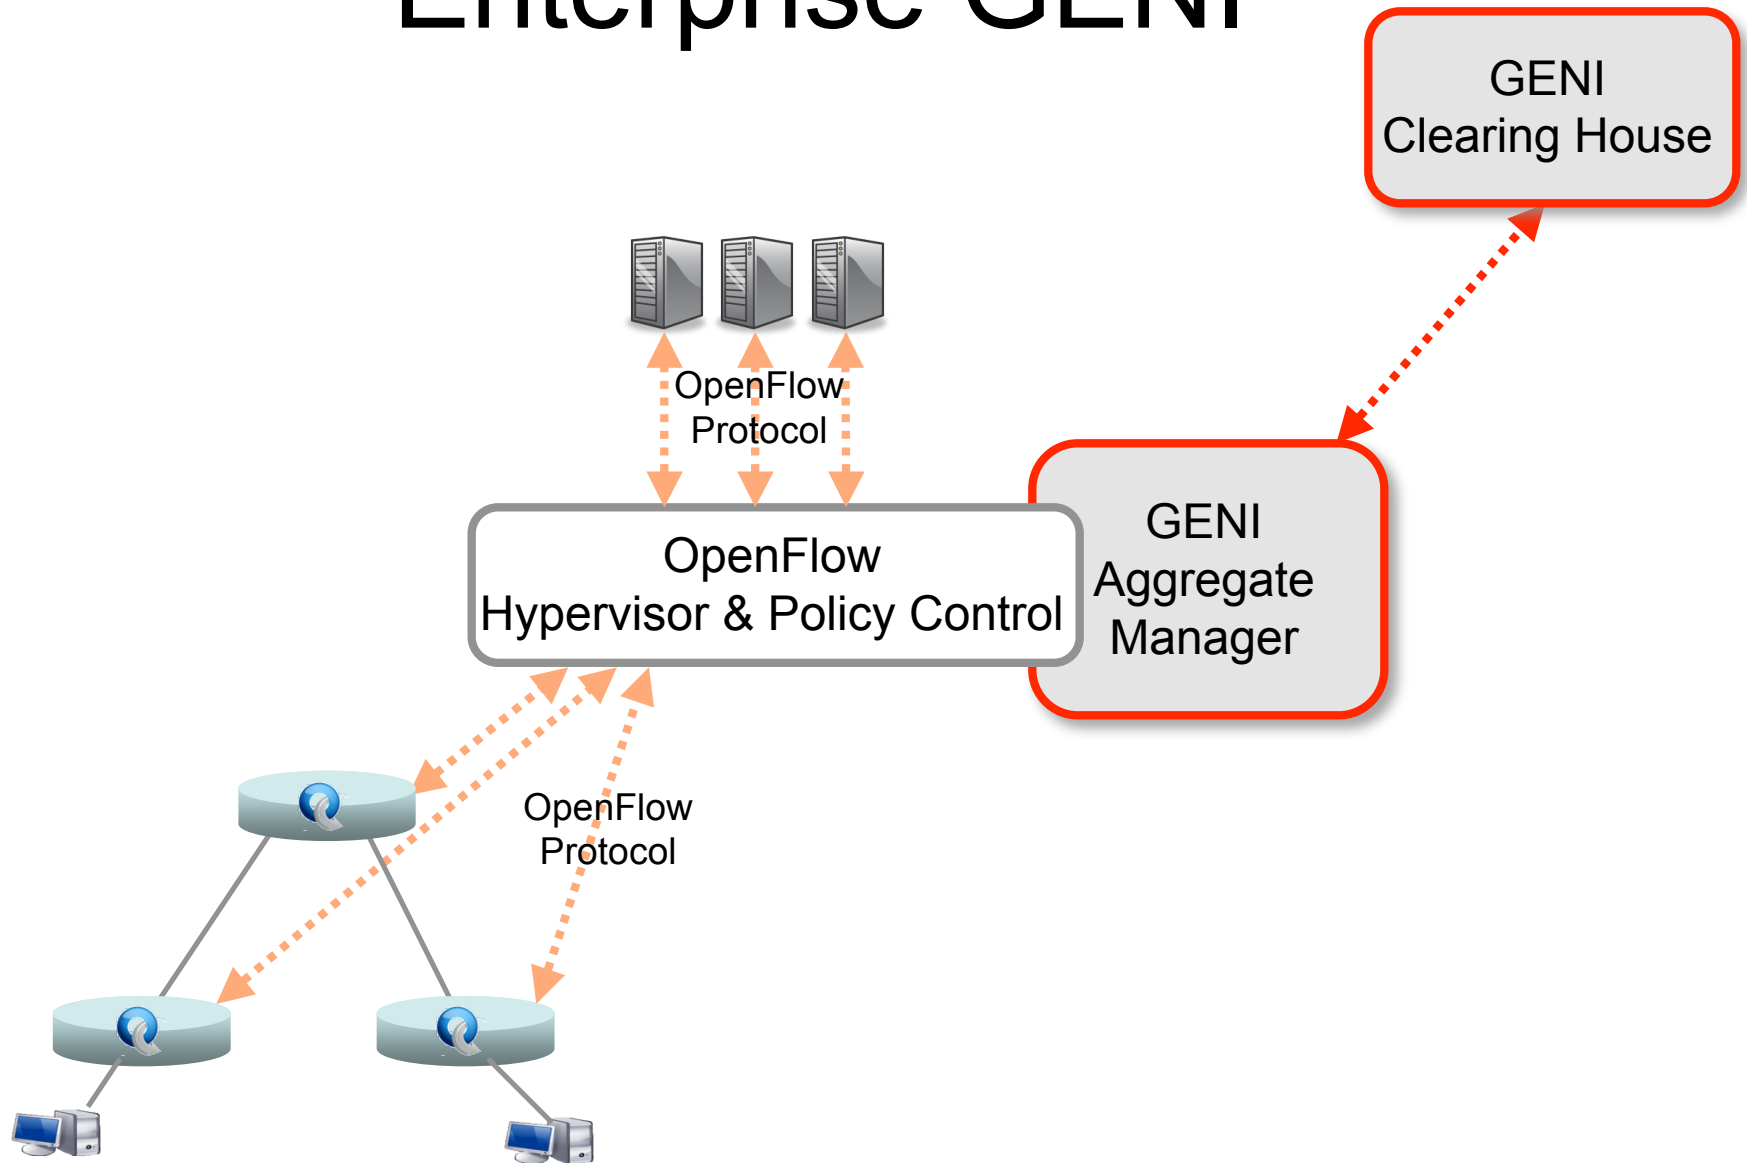
Virtualizing Control



Virtualizing Control



Enterprise GENI



Many Open Questions!


- ❖ Scalability of a controller
- ❖ Load-balancing over redundant controllers
- ❖ Federation, hierarchy and aggregation
- ❖ Protecting the controller against DDOS

Our goal is to enable the research community to explore all these questions

Thanks...

(It takes a village)

Internet2 Team



Chris Small



Matt Zekauskas

Installing Juniper MX-480 in NY



Stanford Team

**We plan trials in early 2009
5-6 college campuses**

**Contact us if you would like to take part
nickm@stanford.edu**