OpenFlow Workshop

APAN FIT Workshop - Hong Kong

Chris Small – Indiana University

Feb 22 2011

GlobalNOC
Global Research Network Operations Center

Sections

- OpenFlow concepts, hardware and software
- OpenFlow use cases
  - Network Operators View
- Demos
- Discussion

Operations

- Focus on why and how to deploy a OpenFlow network
  - Someone deploying OpenFlow Apps not necessarily building them
  - Concepts
  - Nuts and Bolts – What software is availible

- Resources for OpenFlow http://www.openflowswitch.org/wk/index.php/HOTITutorial2010

Keys to Openflow/Software-Defined Networking

- Separation of Control Plane & Data Plane with Open API Between the Two
- Logically Centralized Control-Plane with Open API to Applications
- Network Slicing/Virtualization

- Creates Open Interfaces between Hardware, OS and Applications Similar to Computer Industry
- Increases Competition, Enables Innovation

So why interesting to operations?

- Researchers can use to OpenFlow to explore new network ideas
  - Quick turn around from idea to deployment
- Operators also can use OpenFlow to build (or eventually purchase) interesting apps
  - "À la carte" networking
  - Inexpensive hardware
  - Provide an infrastructure

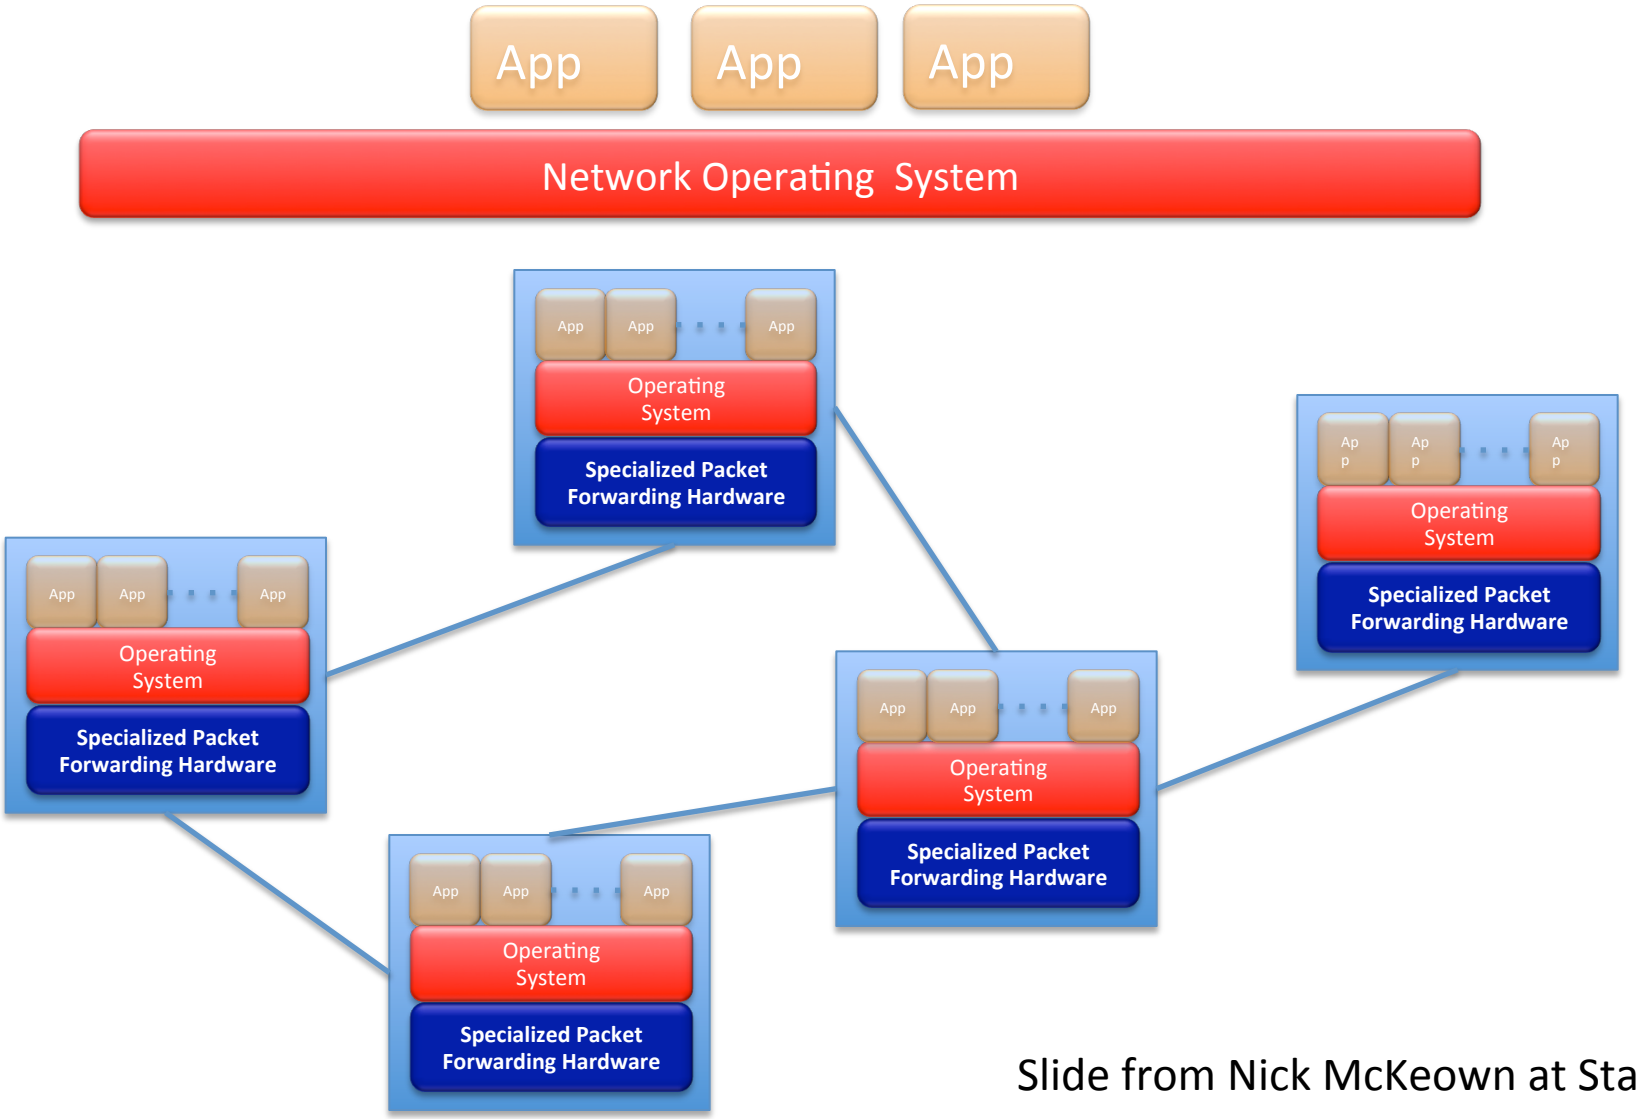

Slide from Nick McKeown at Stanford

The "Software-defined Network"

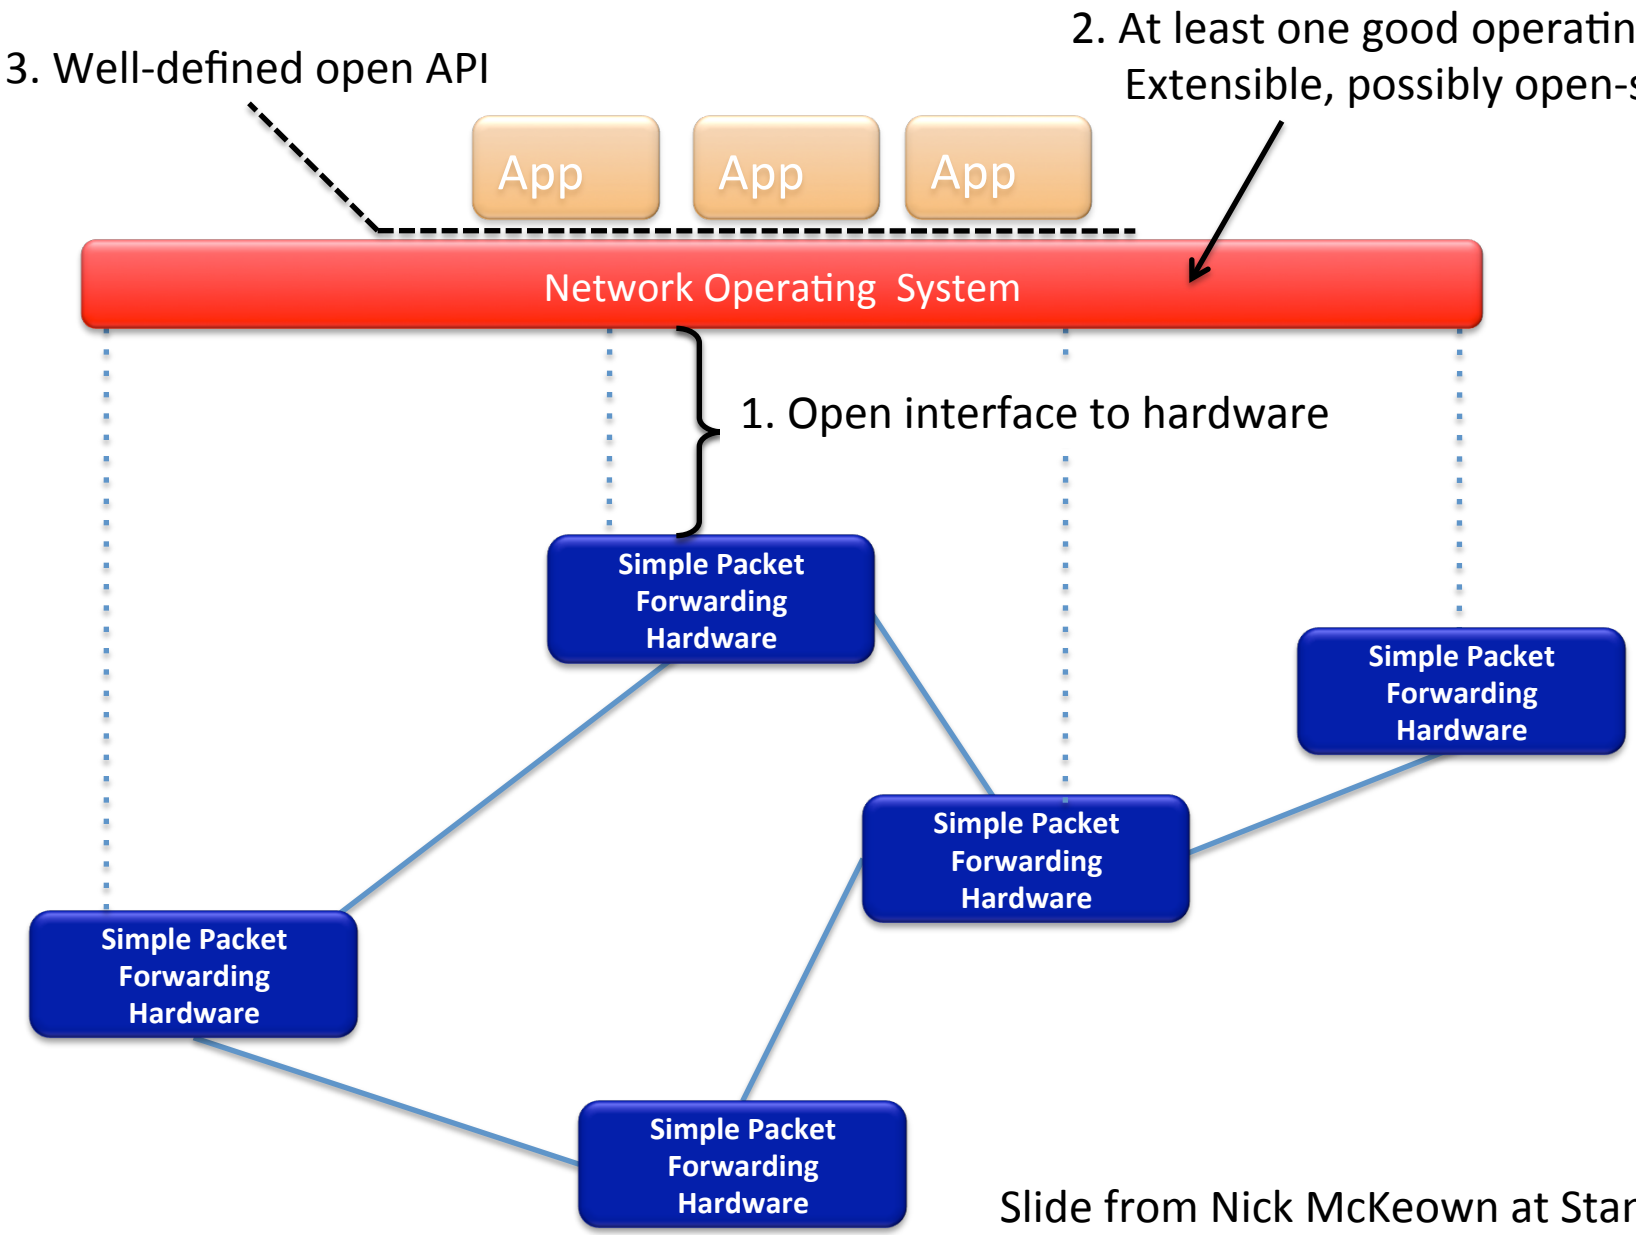

Slide from Nick McKeown at Stanford

Trend

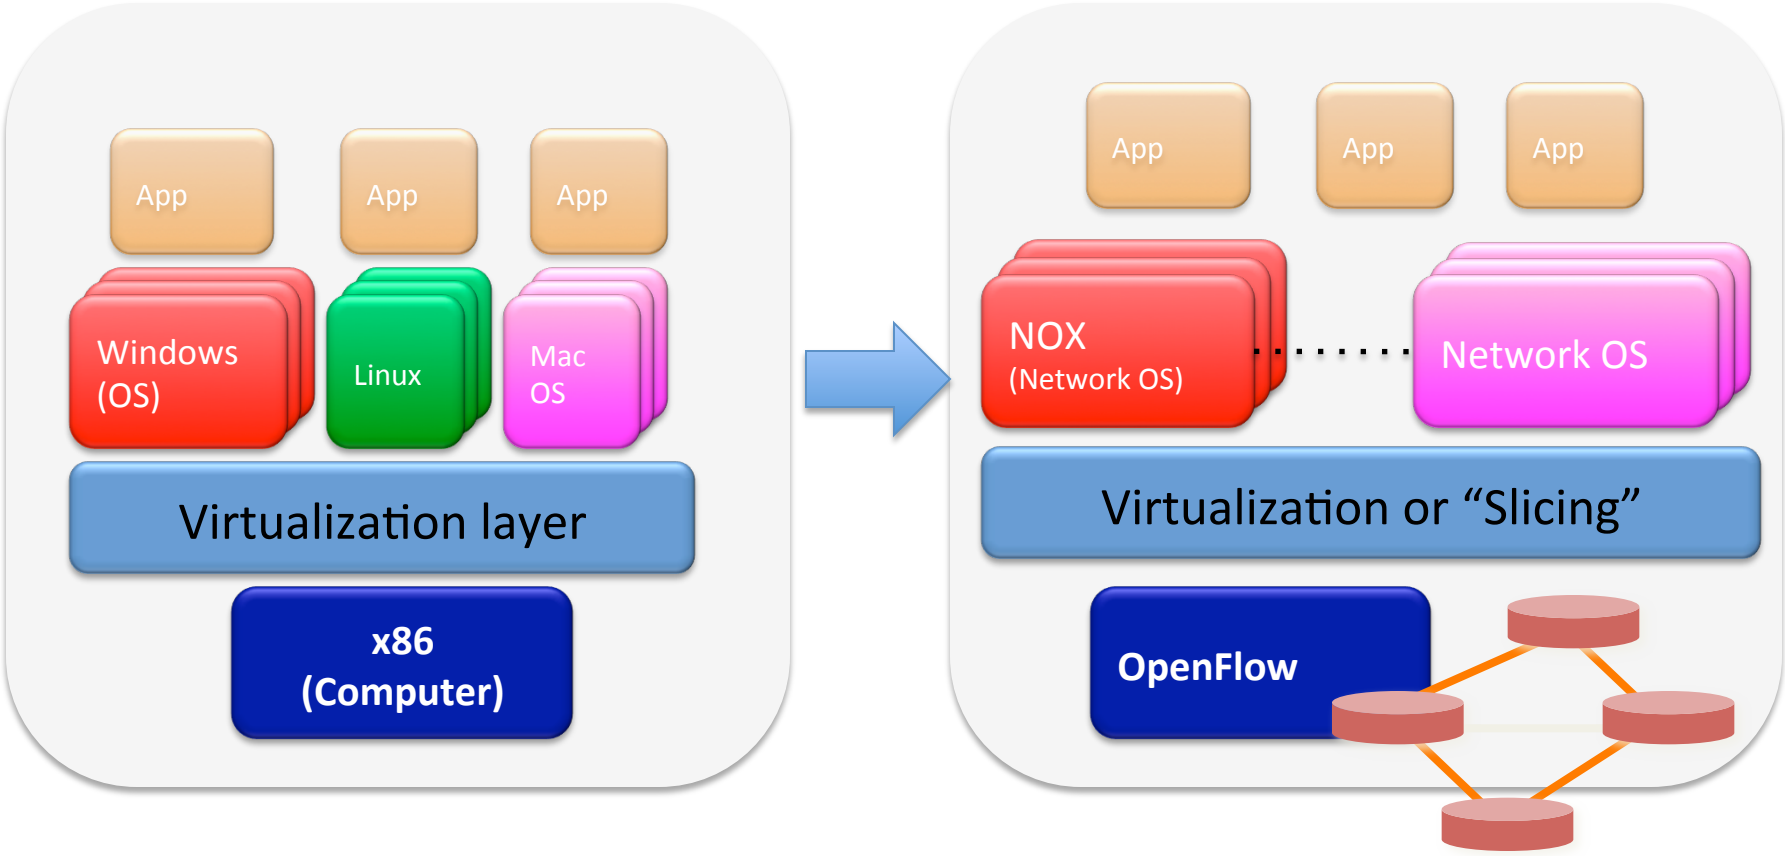

Computer Industry

Network Industry

Slide from Nick McKeown at Stanford

OpenFlow Basics

OpenFlow Basics (1)

Exploit the flow table in switches, routers, and chipsets

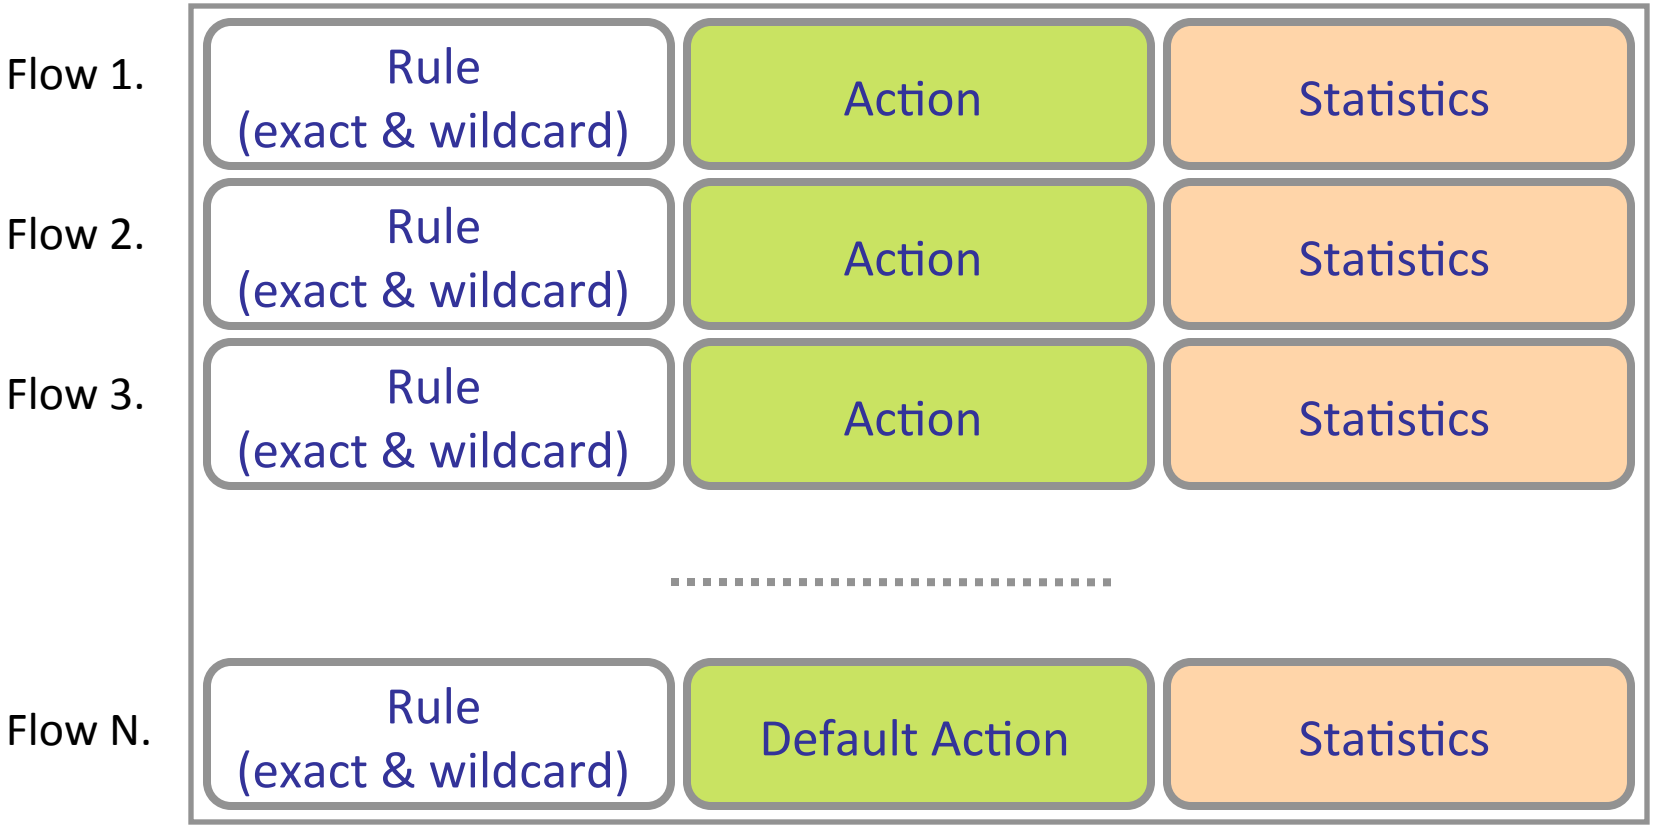

OpenFlow Basics (2)

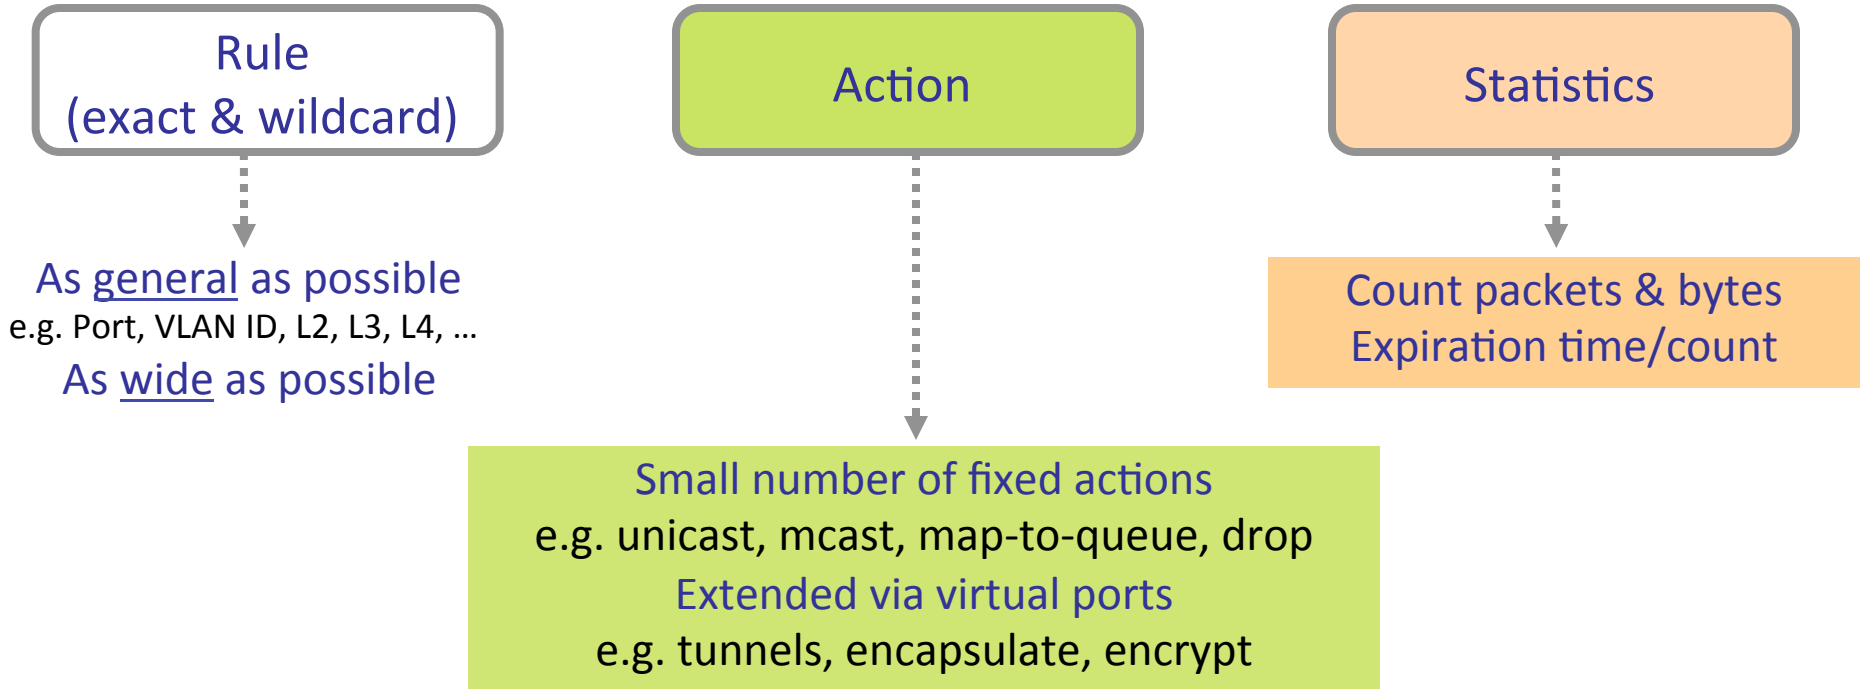

Flow Table Entry

OpenFlow 1.0 Switch

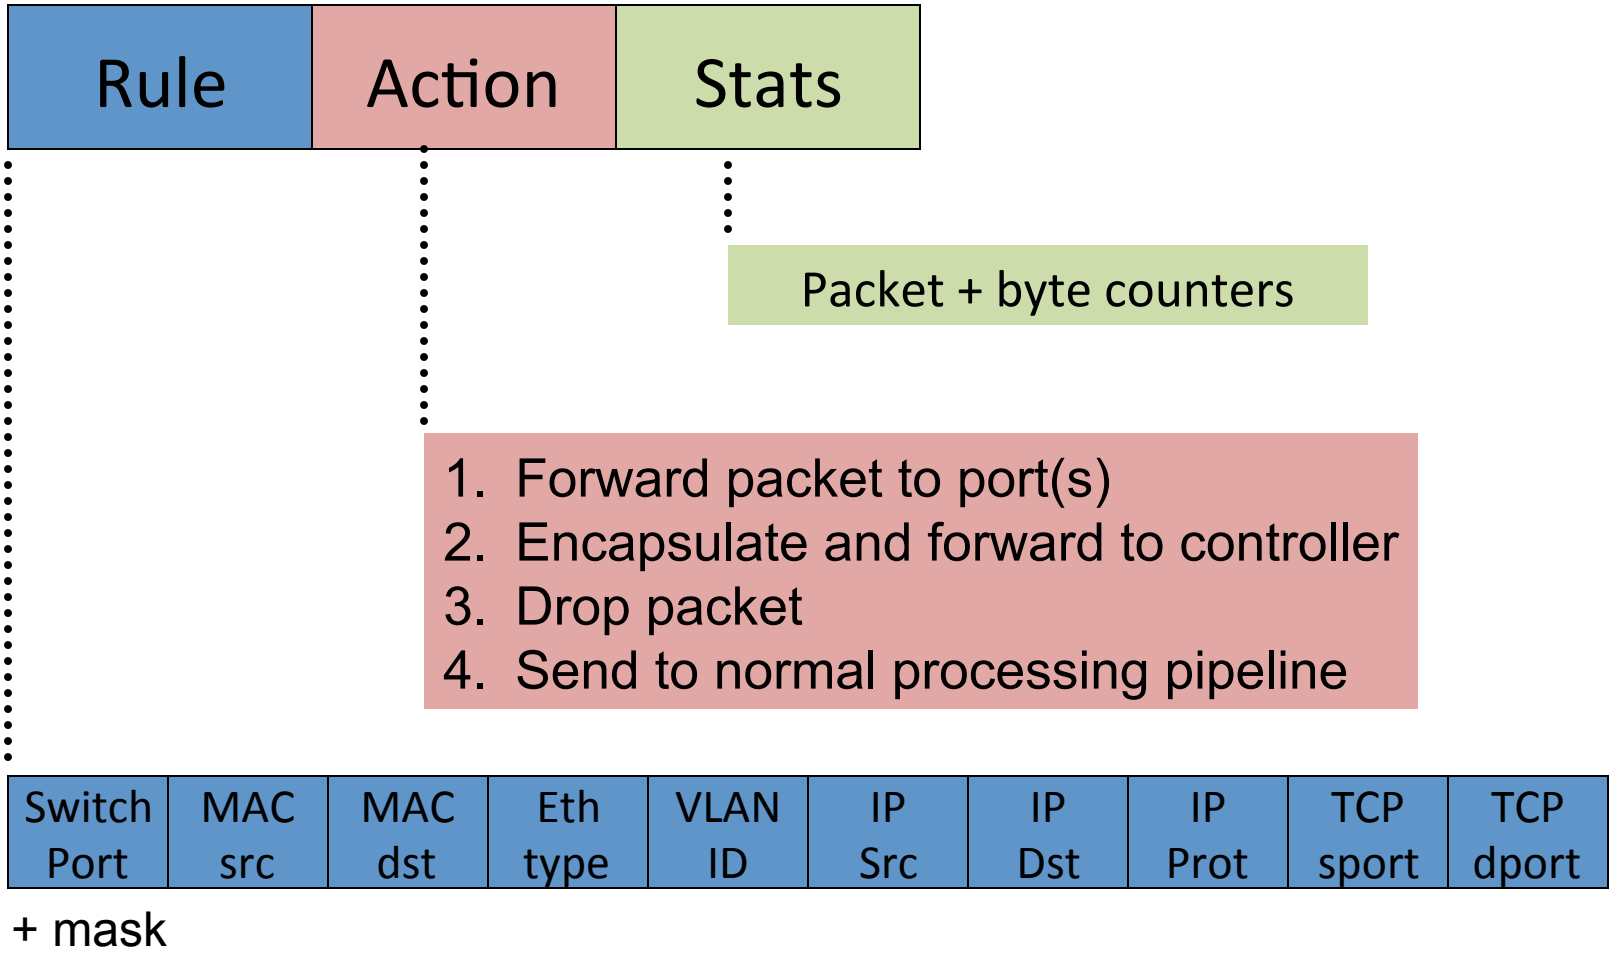

OpenFlow Basics (3)

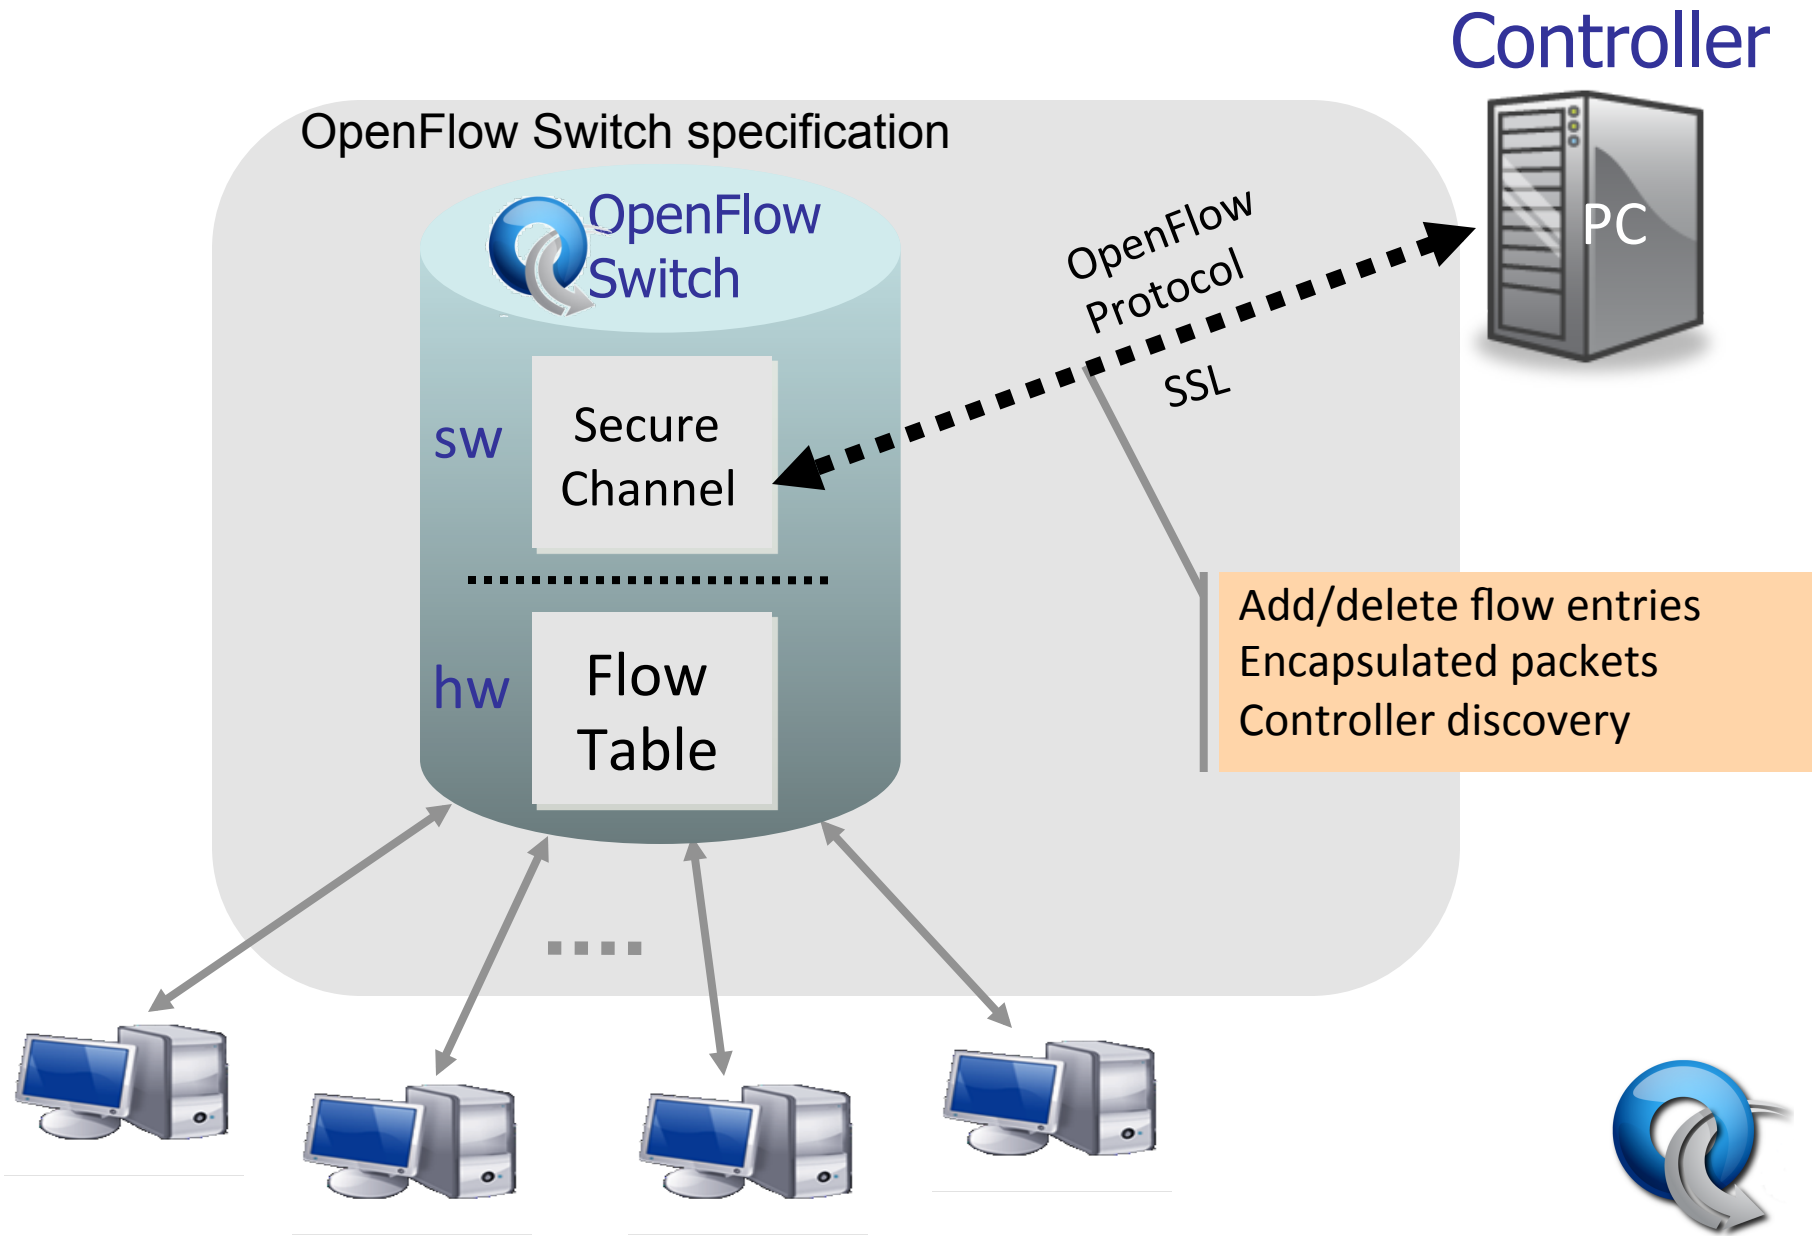

OpenFlow Usage

Dedicated OpenFlow Network

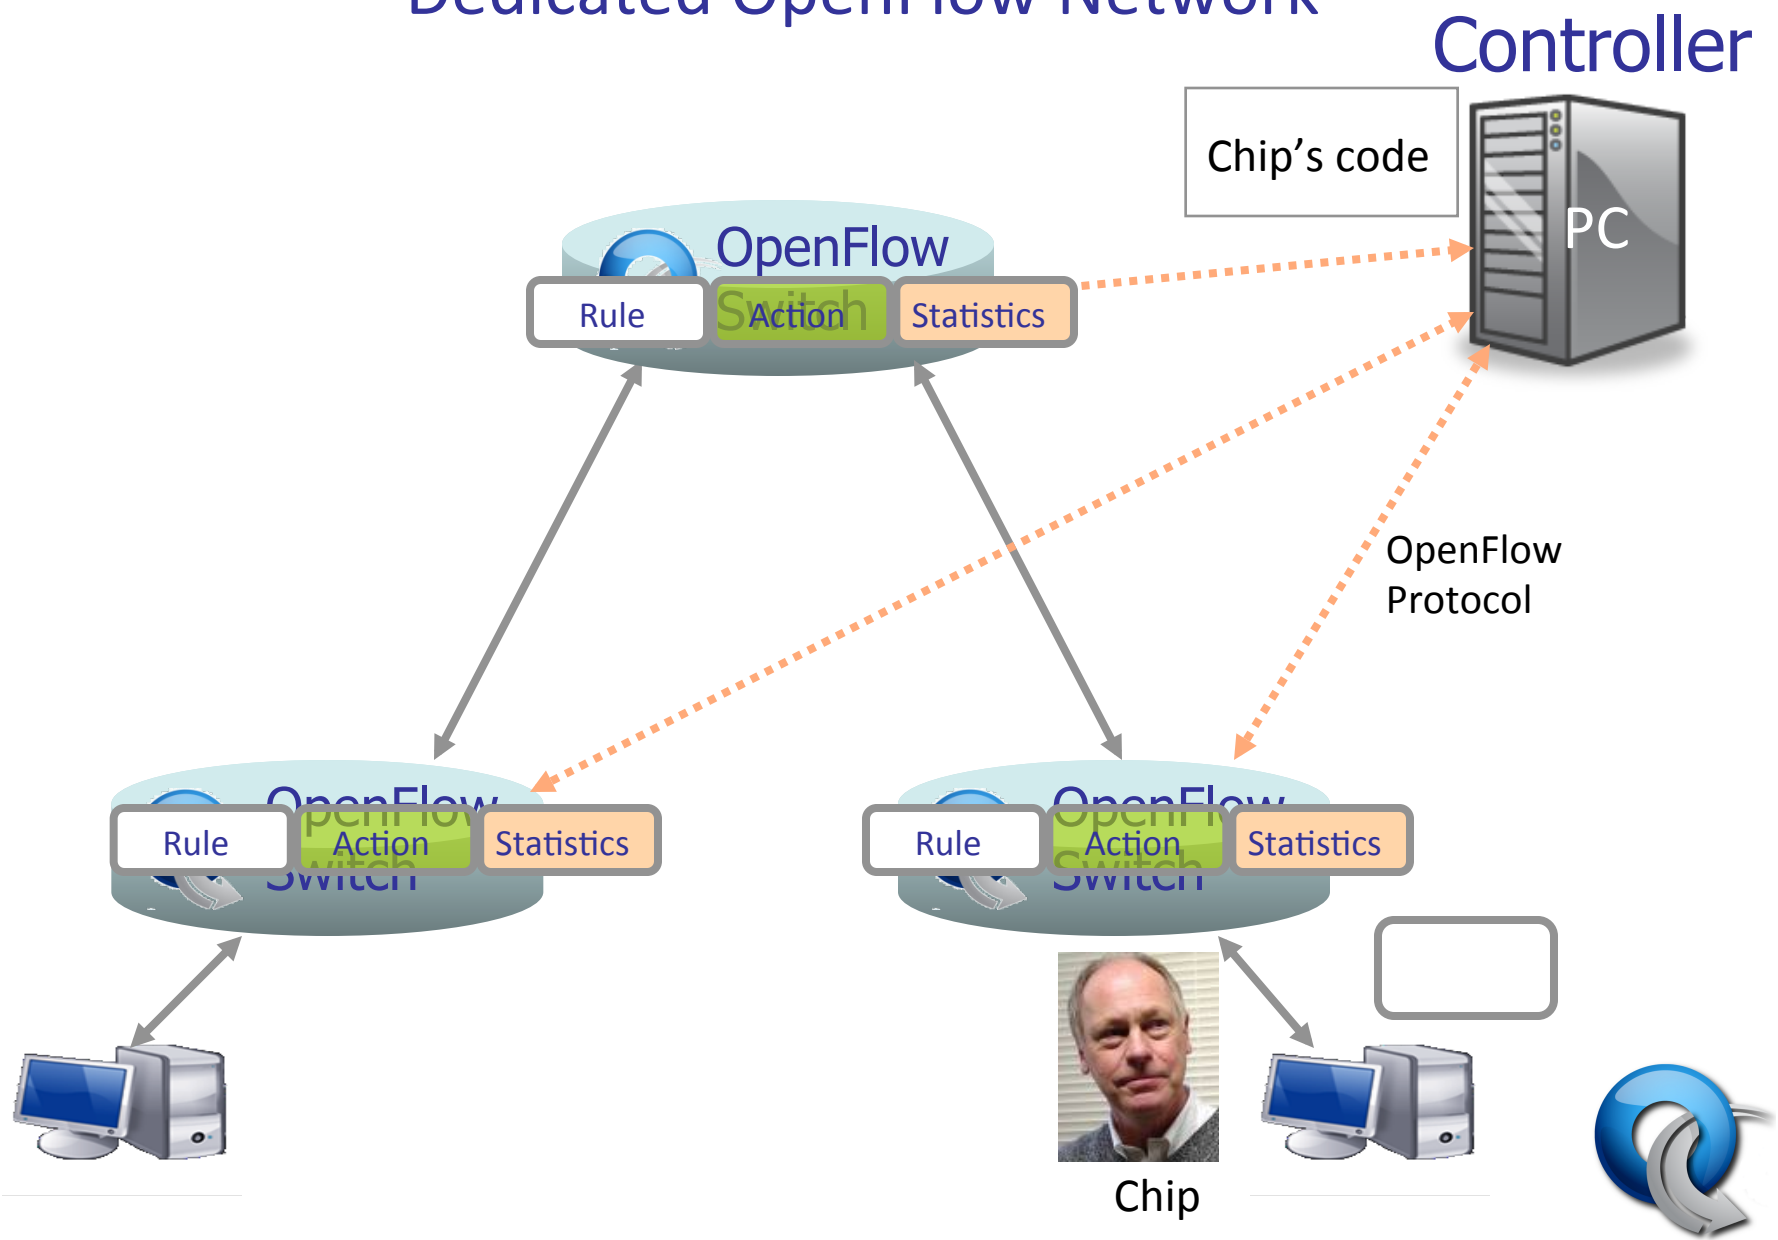

What to do with OpenFlow ?

- 1k-3k TCAM Entries in Typical Edge Switch
- Difficult to take advantage of:
  - Manual Config, SNMP Writes, RADIUS
  - Limited Actions (allow/deny)
  - Vendor Specific
- But what if you could program these through a standard API ?

Possible Uses of Openflow (Quick Wins)

- Security Applications
  - NAC
  - IDS/IPS
  - Remote Packet Capture & Injection
- VM Mobility
  - Redirect specific application traffic to remote site
  - Flow-based forwarding – no need to extend entire broadcast domain – no STP issues

Other Applications

- Load Balancing
- n-cast
  - multiple streams over lossy networks
- Policy (Firewall)
  - SNAC
- Flow based network provisioning

Intercontinental VM Migration

Moved a VM from Stanford to Japan without changing its IP.
VM hosted a video game server with active network connections.

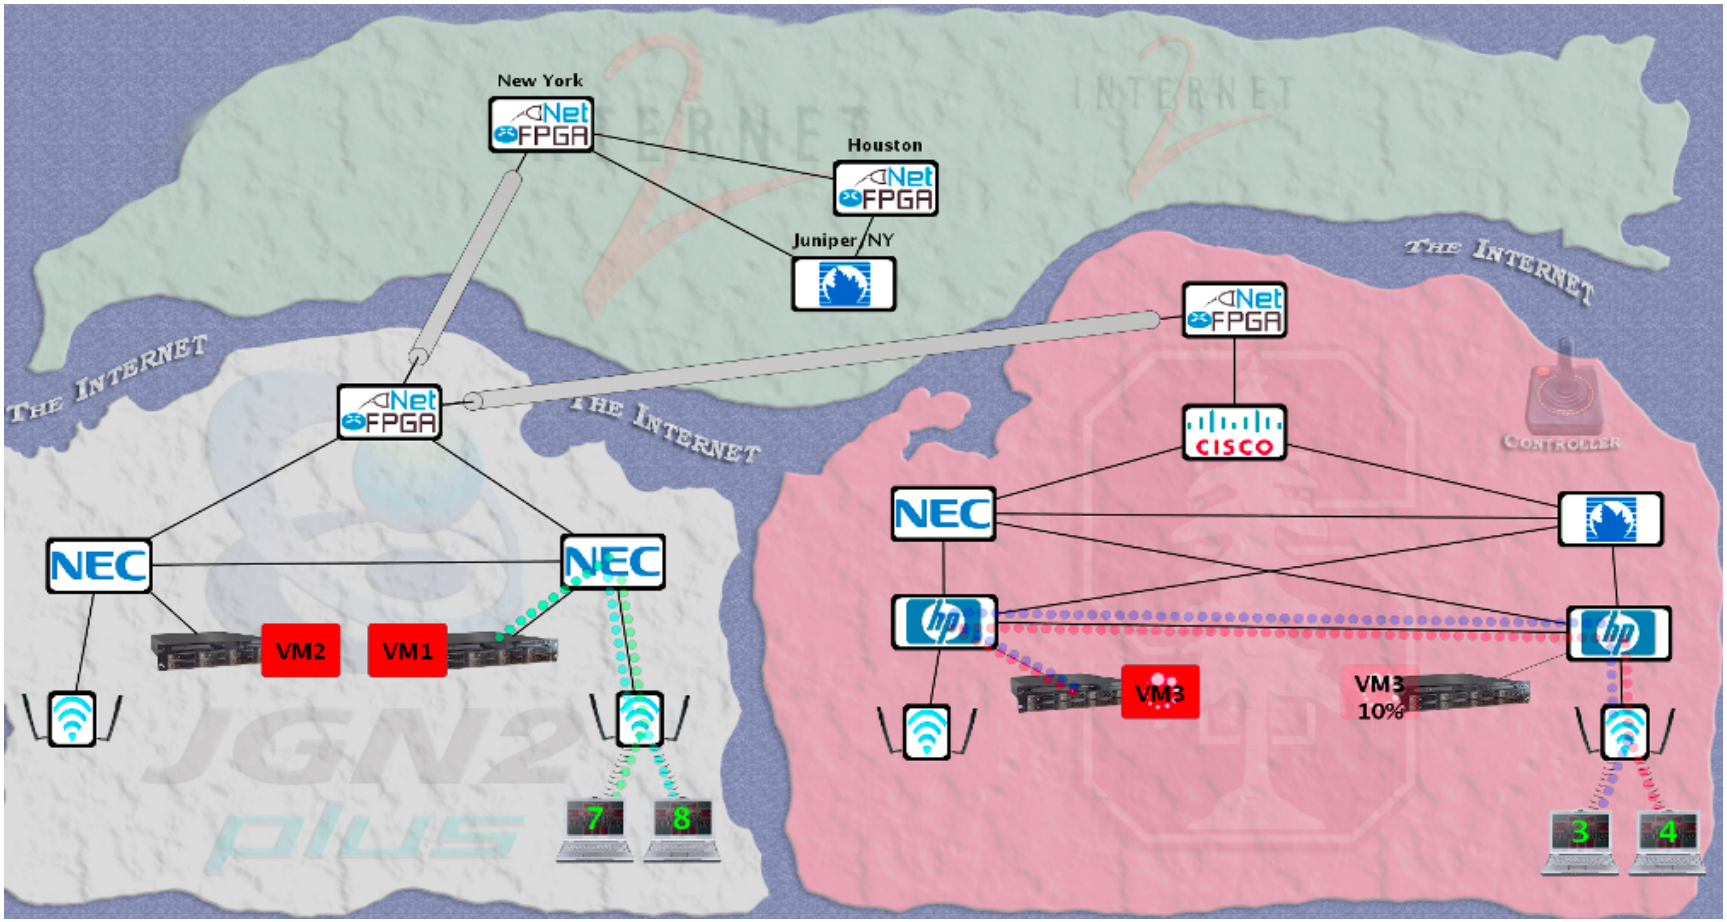

GlobalNOC
Global Research Network Operations Center

Possible Uses of Openflow (Quick Wins)

- Dynamic Circuit Provisioning
  - Don't need to extend layer-2 end-to-end
  - Simply direct specific flows down a engineered path with guaranteed priority
  - Don't have to rely on scripted SSH sessions, SNMP or other sub-optimal ways to programmatically configure switches/routers.

Possible Uses of Openflow (Grand Challenges)

- Distributed Control-Plane Architecture Requires a Lot of State to be Synchronized Across Many Devices
- Many Protocols Needed for Synchronization Internally to Networks (OSPF, RSVP, STP, etc)
- Can these “internal” protocols eventually be removed entirely with only BGP for inter-domain route advertisements ?

OpenFlow Paradigm shifts

- "Wireless like" management of wired switches
- Manipulate virtual switches over many physical devices
  - VM Migration demo
- OSI model breakdown
- Control at the flow level

Deployments

GENI

- GENI OpenFlow deployment on 8 campuses
- Internet2 and NLR backbones
- Integrated with Production hardware on campuses
- Backbone, Regionals (funded in GENI Solicitation 3) and Campuses interconnected
- Outreach to more campuses in future?

Internet2 and NLR

- Internet2
  - Backbone of 5 NEC IP8800
  - Multiple 1G connections (in each direction)
  - L2circuits between sites
- NLR
  - Backbone of 5 HP 6600-24XG
  - 10 G wave between sites

NLR – I2 OpenFlow Core

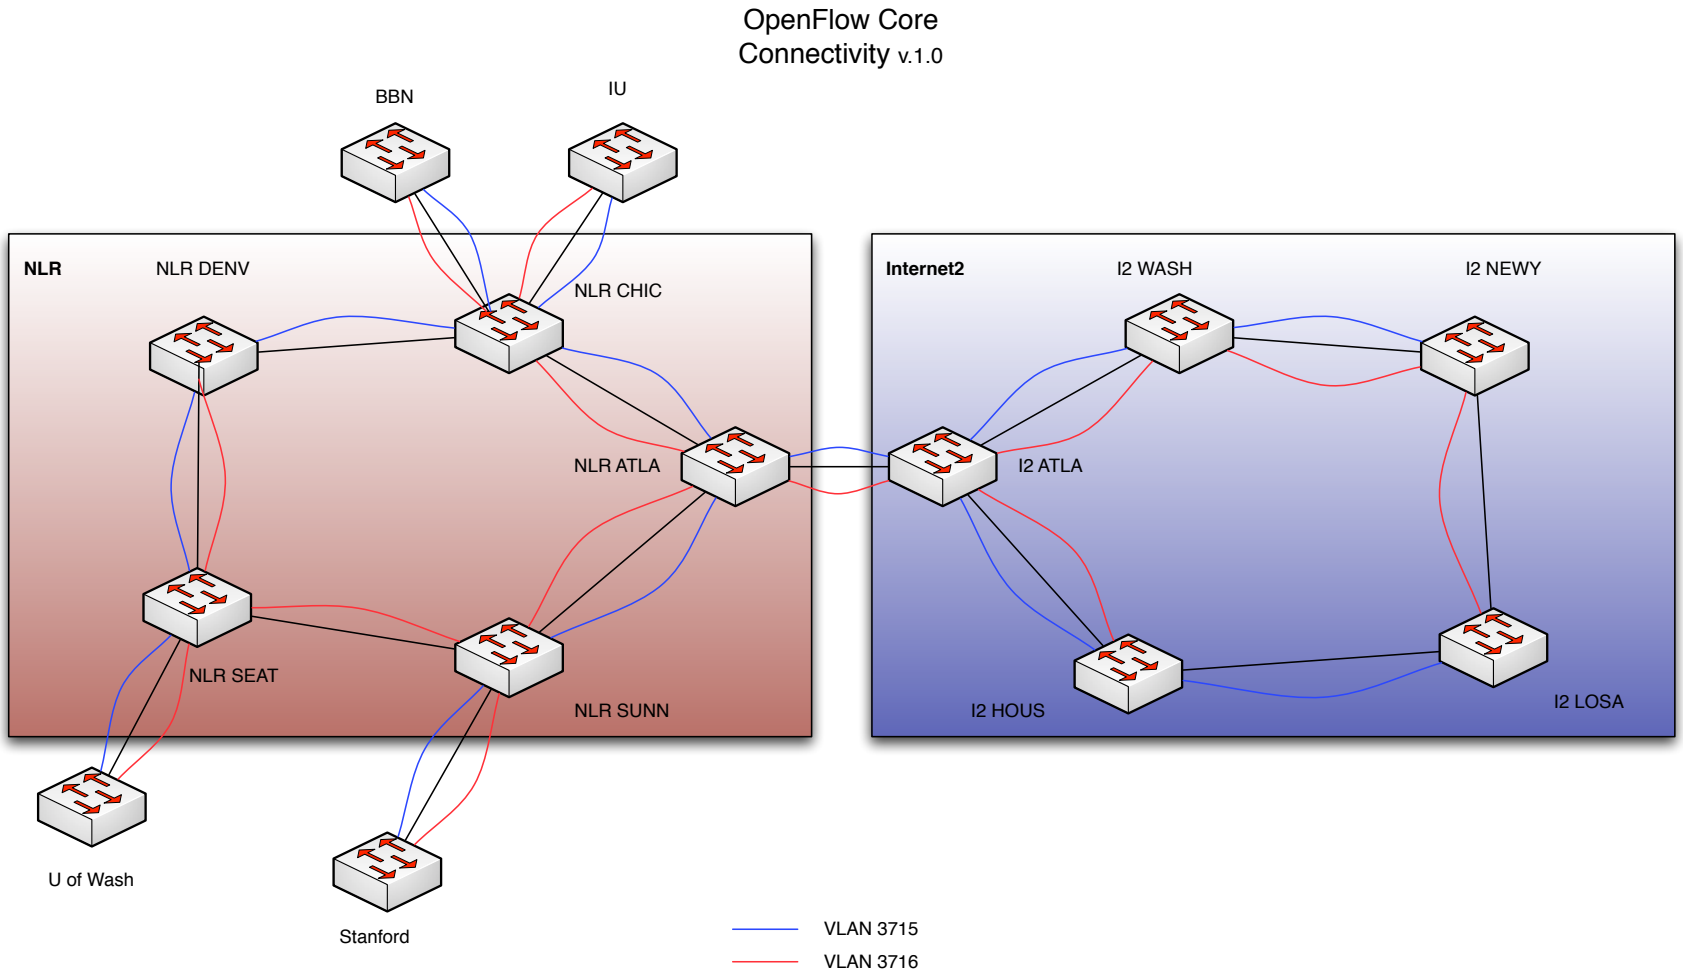

GlobalNOC
Global Research Network Operations Center

IU Campus Deployment

- Focused on Edge (Closet) Deployment
- Goals:
  - Stress-Test Current Implementations
  - Verify “Sandboxing” of Openflow
  - Develop Monitoring Tools
  - Prepare for Production Deployments

IU Deployment

- HP switches in Testlab and Production
  - 4 6600s in Bloomington testlab
  - 1 5406in Testlab/Wireless
  - 2 5406 used by Engineering
  - 3500 in Gigapop
- Pronto switches (w/ Purdue Calumet)
- NetGear switches
- NetFPGA 10G and 1G?

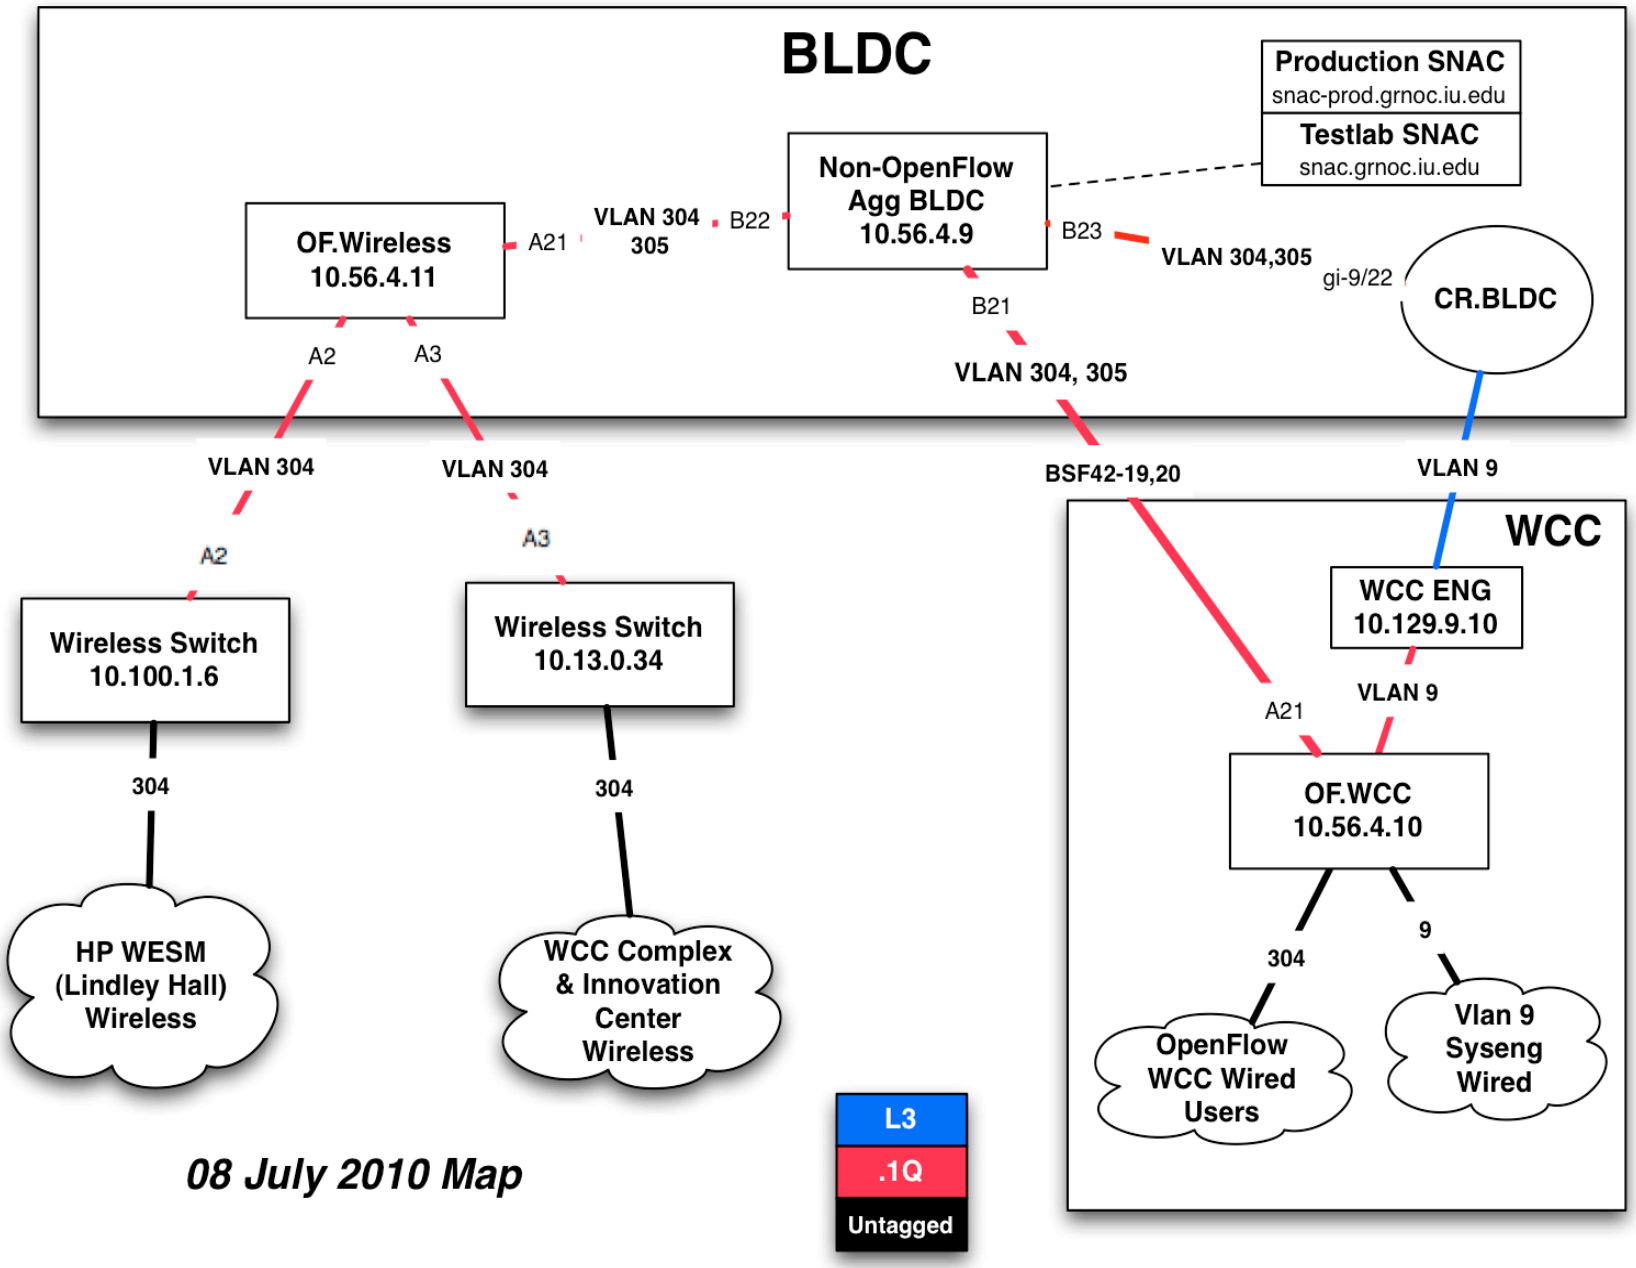

08 July 2010 Map

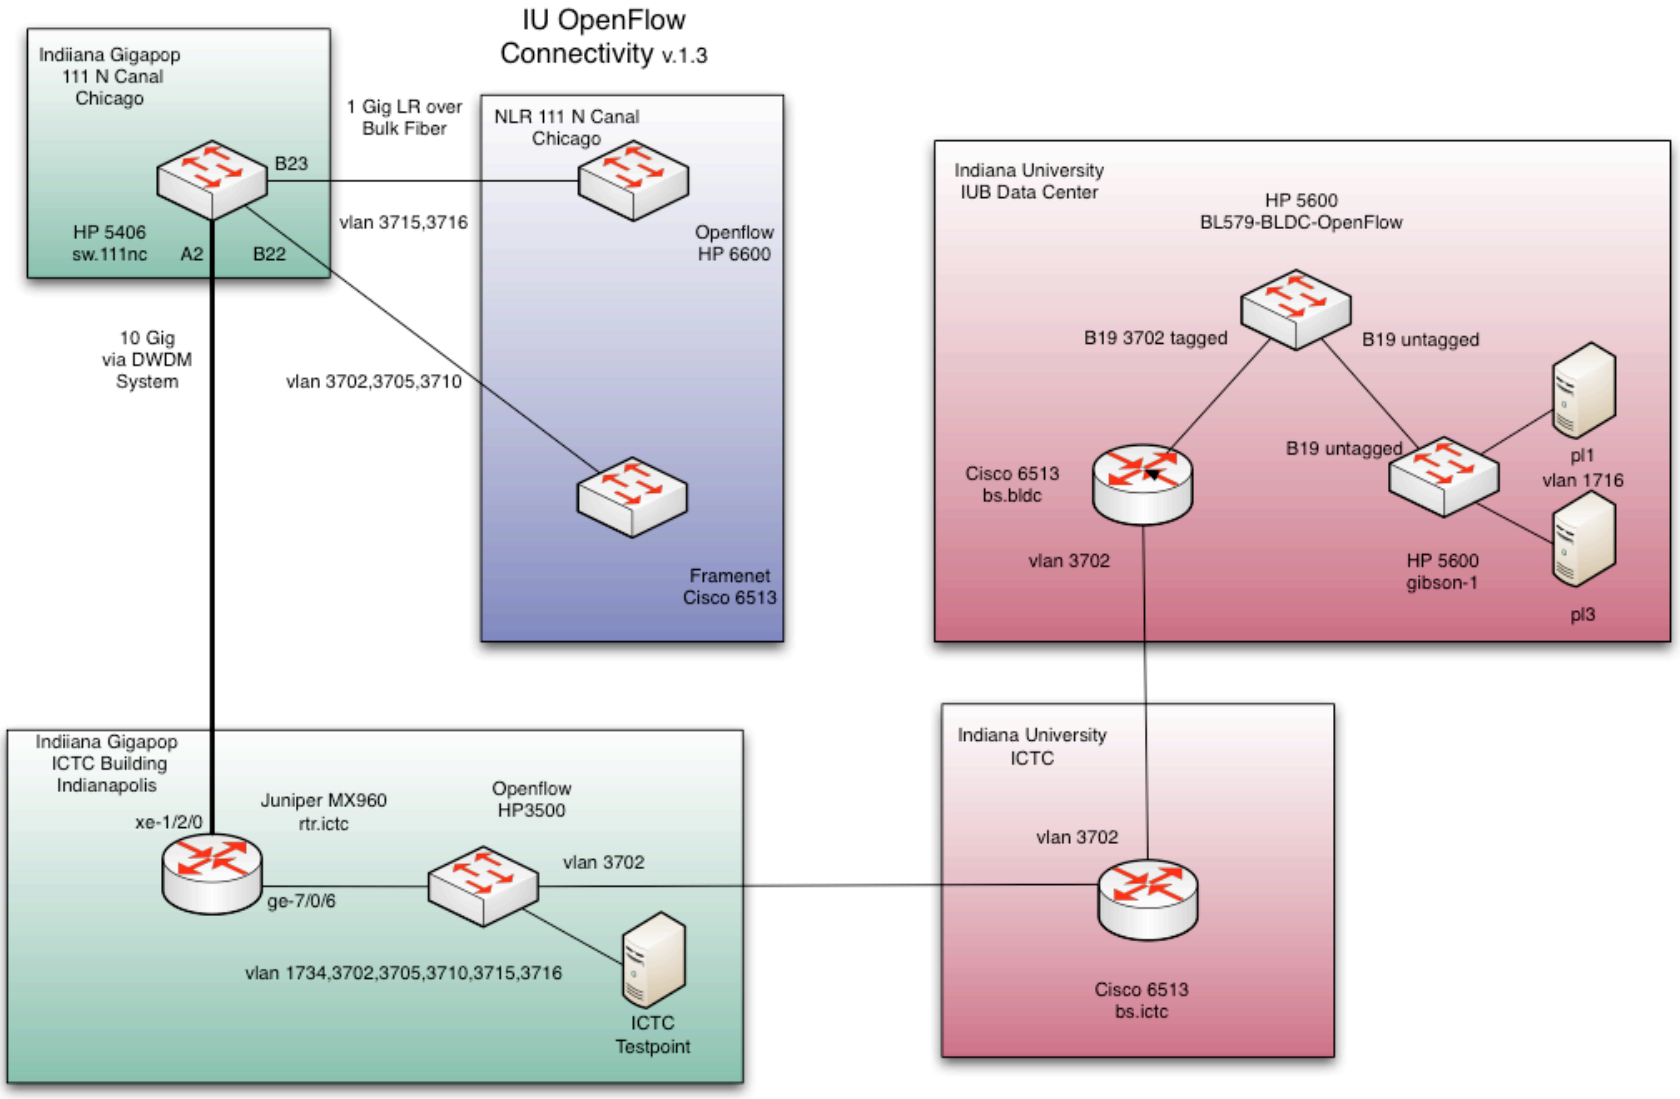

GlobalNOC
Global Research Network Operations Center

3 New EU Projects: OFELIA, SPARC, CHANGE

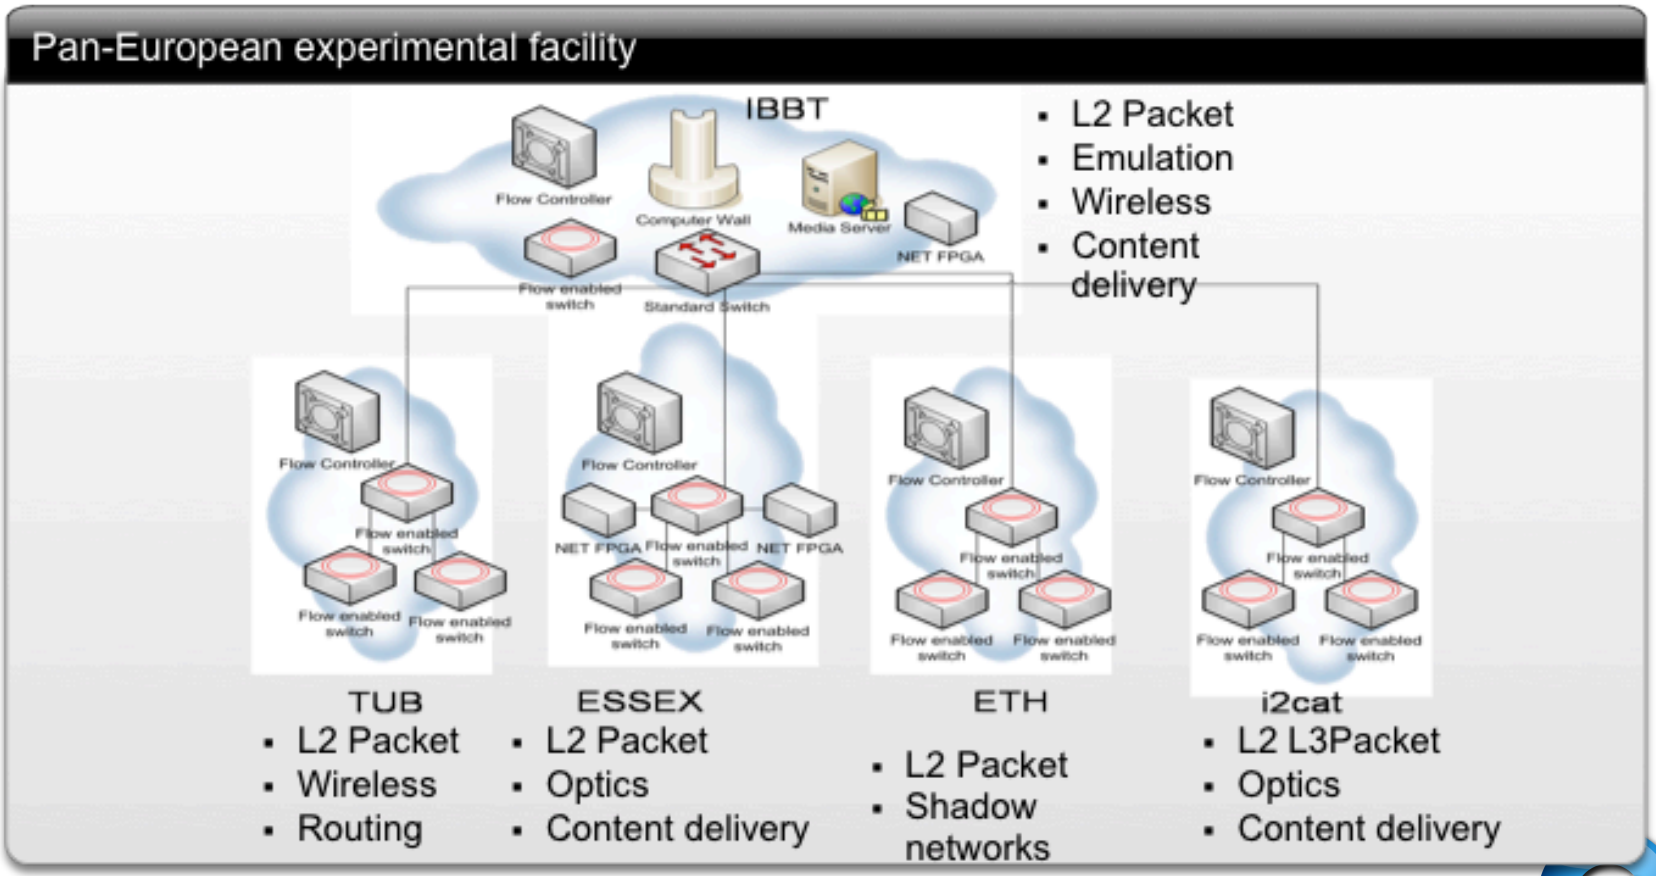

GlobalNOC
Global Research Network Operations Center

EU Project Participants

- Germany
  - Deutsch Telekom Laboratories
  - Technishche Universitat Berlin
  - European Center for ICT
  - ADVA AG Optical Networking
  - NEC Europe Ltd.
  - Eurescom
- United Kingdom
  - University of Essex
  - Lancaster University
  - University College London
- Spain
  - i2CAT Foundation
  - University of the Basque Country, Bilbao
- Romania
  - Universitatea Politehnica Bucuresti
- Sweden
  - ACREO AB (Sweden)
  - Ericsson AB Sweden (Sweden)
- Hungary
  - Ericsson Magyarorszag Kommunikacios Rendszerek KFT
- Switzerland
  - Dreamlab Technologies
  - Eidgenossische Technische Hochschule Zurich
- Italy
  - Nextworks
  - Universita` di Pisa
- Belgium
  - Interdisciplinary Institute for Broadband Technology
  - Universite catholique de Louvain

OpenFlow Deployment in Japan
NEC and JGN2Plus (NICT)

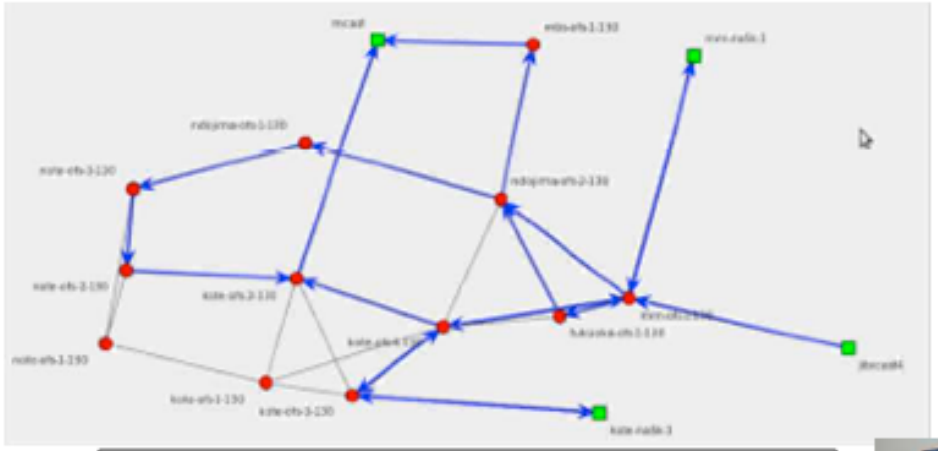

Flow and topology visualization GUI

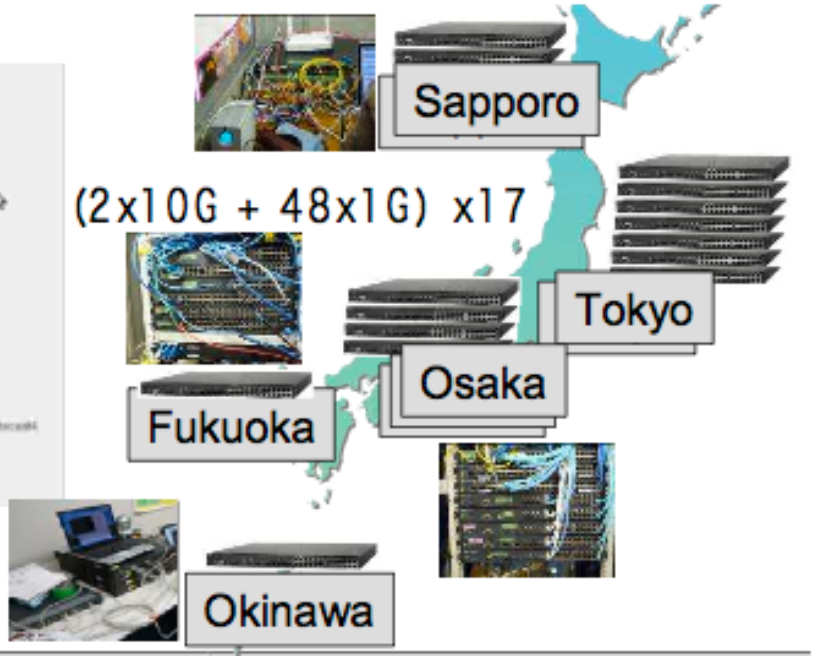

- Network virtualization and slicing
- HD video distribution in different slices
  - Baseball game
  - Snow festival

Global Interest

11,129 visits came from 1,252 cities

Detail Level: City | Country/Territory | Sub Continent Region | Continent Dimension: None

Site Usage | Goal Set 1

Views:

Visits	Pages/Visit	Avg. Time on Site	% New Visits	Bounce Rate
11,129	2.85	00:05:33	41.05%	49.09%
% of Site Total: 100.00%	Site Avg: 2.85 (0.00%)	Site Avg: 00:05:33 (0.00%)	Site Avg: 40.89% (0.37%)	Site Avg: 49.09% (0.00%)

	Detail Level: City	Visits ↓	Pages/Visit	Avg. Time on Site	% New Visits	Bounce Rate
1.	Shibuya	530	2.85	00:02:47	36.79%	45.09%
2.	Hanoi	519	3.24	00:06:16	25.01%	40.85%
3.	San Jose	381	3.01	00:03:36	39.63%	46.19%
4.	Stanford	368	3.61	00:03:39	12.23%	41.85%
5.	Tokyo	290	2.78	00:03:53	45.17%	56.90%
6.	Atlanta	230	2.86	00:05:26	18.79%	43.04%
7.	San Ch'Ung	226	3.53	00:05:45	23.89%	31.86%
8.	San Francisco	185	2.63	00:02:58	39.46%	52.43%
9.	Mountain View	176	3.24	01:36:33	22.73%	48.43%
10.	Bangalore	167	2.37	00:05:11	39.52%	47.90%

Current Trials and Deployments

68 Trials/Deployments - 13 Countries

Current Trials and Deployments

USA-Academia

Stanford University, CA
University of Washington, WA
Rutgers University, NJ
Princeton University, NJ
Clemson University, SC
Georgia Tech, GA
University of Wisconsin at Madison, WI
Indiana University
ICSI Berkeley, CA
University of Massachusetts at Lowell
Clarkston University
Columbia University (course offered)
University of Kentucky
UC San Diego
UC Davis
iCAIR/Northwestern
Rice University
Purdue University
Northern Arizona University

USA-Industry

Internet2
Cisco
Juniper
HP
Ciena
Deutsche Telekom R&D Lab
Marvell
Broadcom
Google
Unnamed Data Center Company
Toroki
Nicira
Big switch networks
Orange Labs

USA-Government

BBN
Unnamed Federal Agency

Current Trials and Deployments

Brazil
University of Campinas
Federal University of Rio de Janeiro
Federal University of Amazonas
Foundation Center of R&D in Telecomm.

Canada
University of Toronto

Germany
T-Labs Berlin
Leibniz Universität Hannover

France
ENS Lyon/INRIA

India
VNIT
Mahindra Satyam

Italy
Politecnico di Torino

United Kingdom
University College London
Lancaster University
University of Essex

Taiwan
National Center for High-Performance Computing
Chunghwa Telecom Co

Japan
NEC
JGN Plus
NICT
University of Tokyo
Tokyo Institute of Technology
Kyushu Institute of Technology
NTT Network Innovation Laboratories
KDDI R&D Laboratories
Unnamed University

South Korea
KOREN
Seoul National University
Gwangju Institute of Science & Tech
Pohang University of Science & Tech
Korea Institute of Science & Tech
ETRI
Chungnam National University
Kyung Hee University

Spain
University of Granada

Switzerland
CERN

OpenFlow and GENI

8 Universities, GPO/BBN, & 2 National Backbones

OpenFlow Concepts, Hardware and Software

OpenFlow Hardware

Juniper MX-series

NEC IP8800

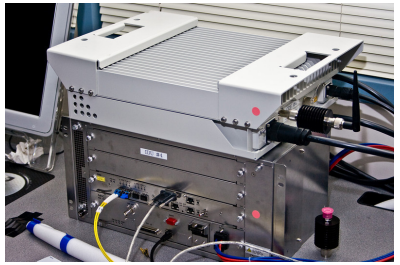

WiMax (NEC)

HP Procurve 5400

Cisco Catalyst 6k

PC Engines

Quanta LB4G

Netgear

More Equipment Soon

Controllers

- The Network "OS"
- Open Source
  - NOX
    - Nicira
    - C++/Python
  - Beacon
    - BigSwitch
  - Maestro
    - Rice
- Commercial
  - NEC

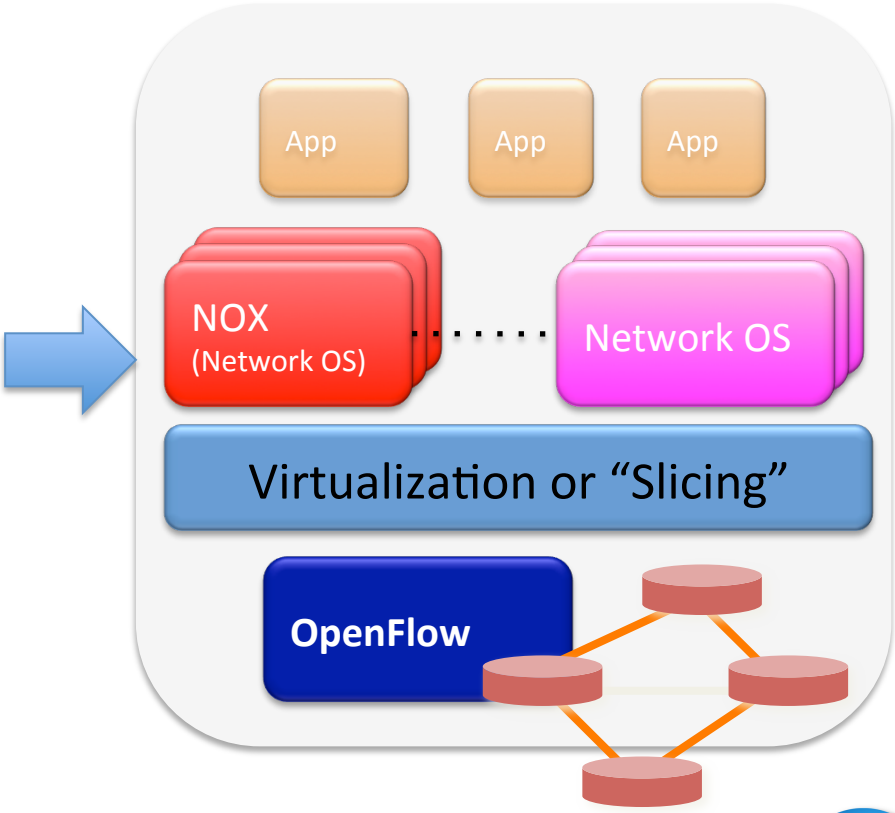

Applications

- Use controller software to build applications
- Possible operational uses
  - Layer 2 provisioning
  - Layer 3 routing
  - Load Balancing
  - Distributed Firewall
  - Monitoring / IDS
- Research use on production networks

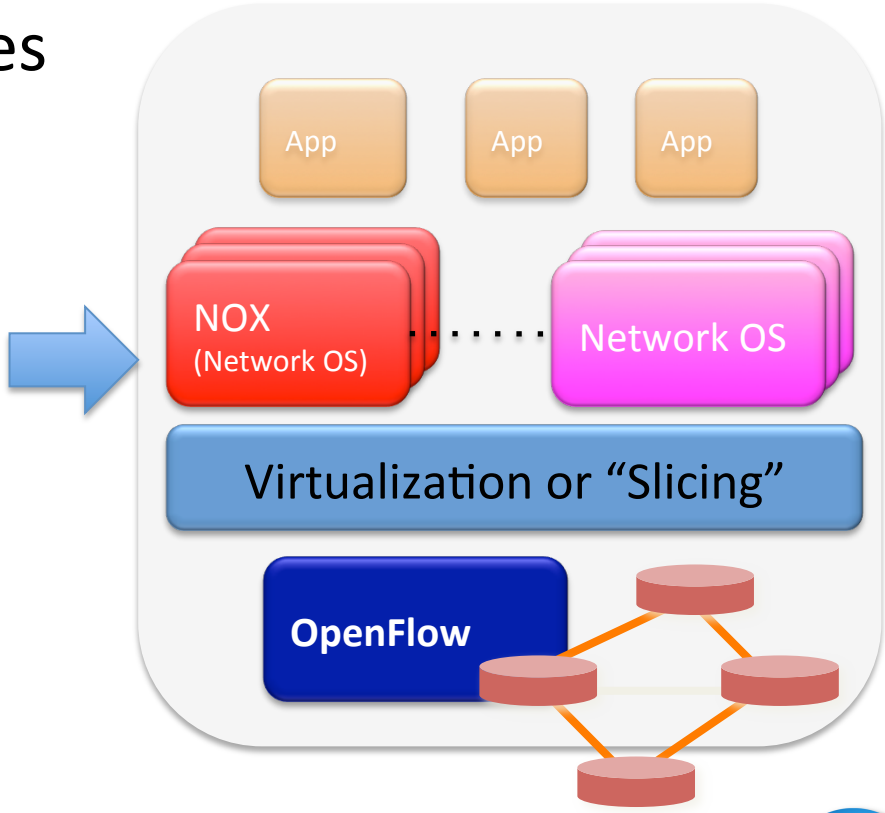

GlobalNOC
Global Research Network Operations Center

Flowvisor

- Sends traffic from the same switch(es) to multiple controllers
- Acts like a Hypervisor for network equipment
- Rule set similar to OpenFlow rules that send traffic to multiple controllers
- Most research shared infrastructure will use Flowvisor to have multiple controllers control the same switches

Fvctl

- Fvctl used to control flowvisor (over XMLRPC)
- Can create slice, direct traffic to “slices”, see
- Flowspace is the set of mapping rules
- Devices Identified by DPID

chsmall@flowvisor:~$ fvctl listDevices
Device 0: 0e:83:00:23:47:c8:bc:00
Device 1: 0e:83:00:26:f1:40:a8:00

chsmall@flowvisor:~$ fvctl listFlowSpace
rule 0: FlowEntry[dpid=[all_dpids],ruleMatch=[OFMatch[]],actionsList=
[Slice:meas_manager=4],id=[236],priority=[10],]

SNAC

- Simple Network Policy Controller
- Web-Based Policy manager
- IU production SNAC at snac-prod.grnoc.iu.edu
- Can provide distributed firewall services
- Some statistics collected

Monitors | Policy | Settings | Help | user: admin | Logout

Overview

- Network Overview
- Switches
- Hosts
- Users
- Locations
- Groups
- Network Events Log

Server Information

Uptime:	1 day 16 hours 52 minutes 53 seconds
CPU Load:	1%
Flows/sec:	7

Entity Counts (Active/Total /Unregistered)

Switches:	3 / 3 / 0
Locations:	18 / 18 / 0
Hosts:	17 / 22 / 4
Users:	1 / 3 / 0

Policy Statistics

Total Drops:	47698
Total Rules:	15

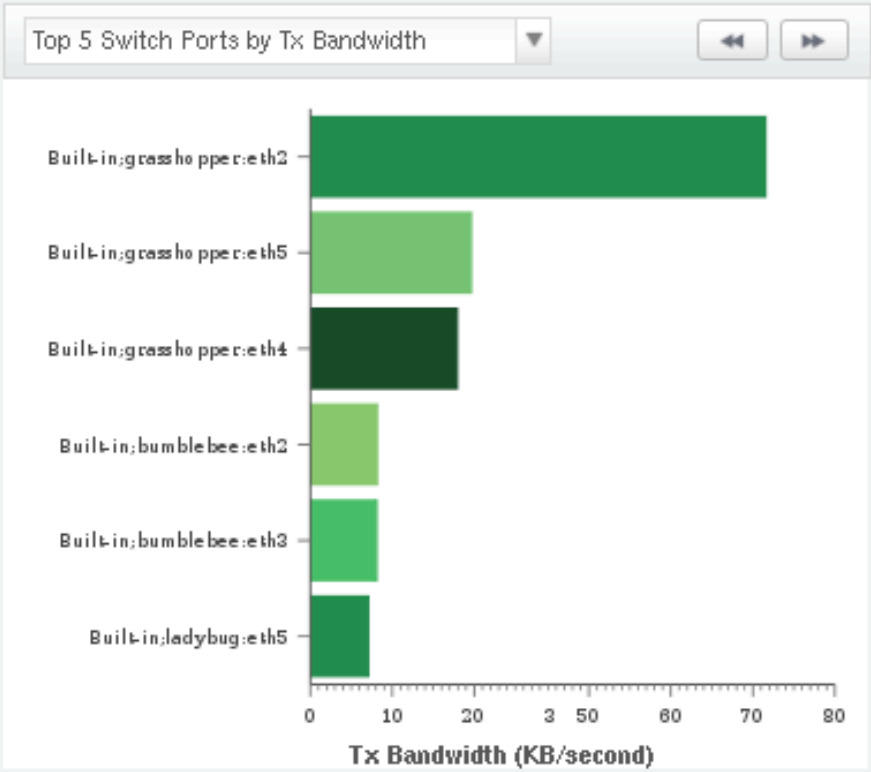

Internal Network | 6 Active Admins

Version: 0.3.0+build43

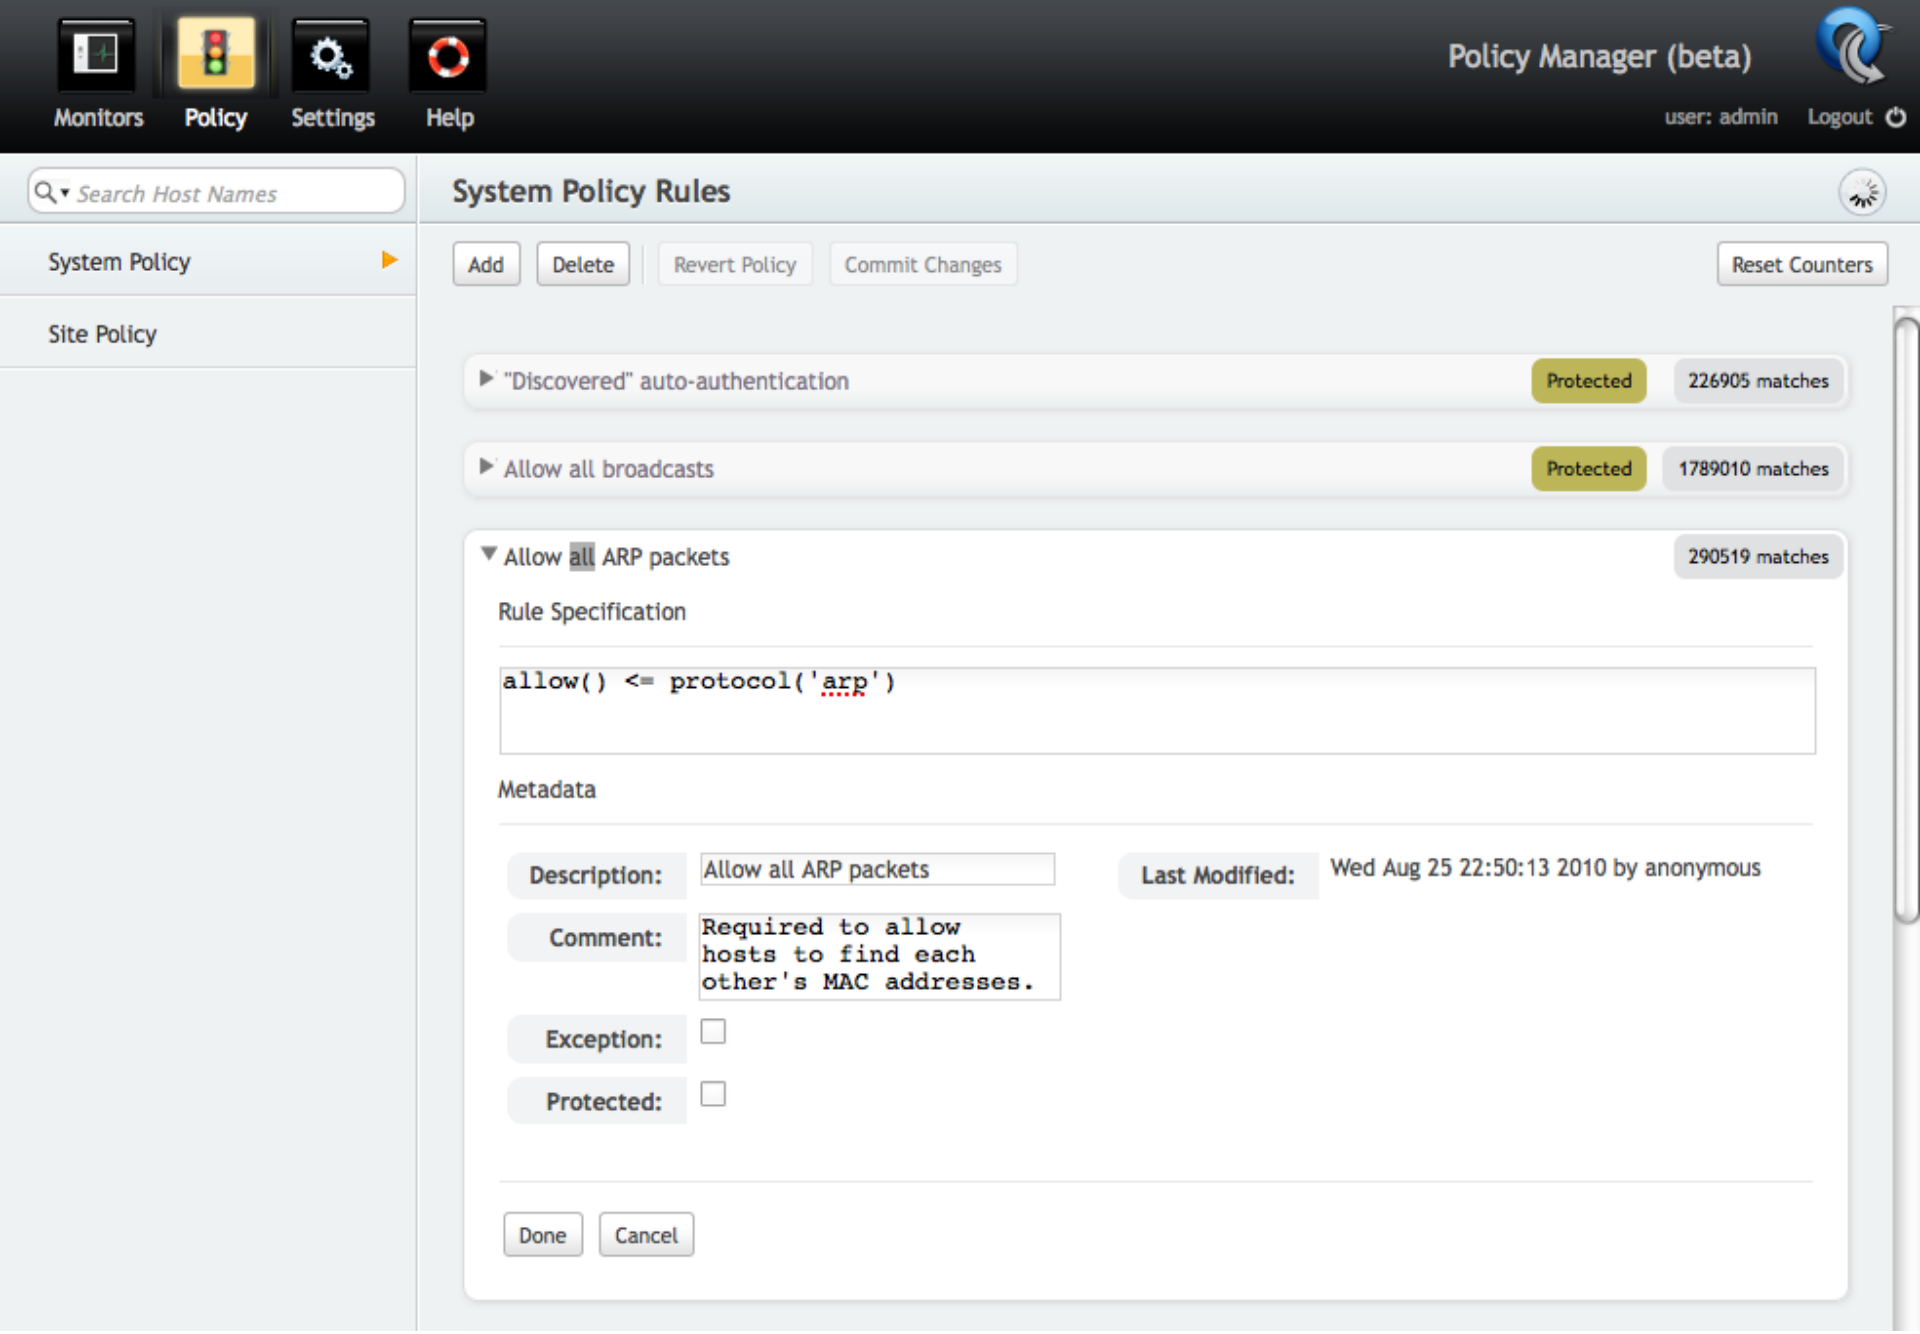
Monitors
Policy
Settings
Help
Policy Manager (beta)
user: admin
Logout
Search Host Names
System Policy
Site Policy
System Policy Rules
Add
Delete
Revert Policy
Commit Changes
Reset Counters
"Discovered" auto-authentication
Protected
226905 matches
Allow all broadcasts
Protected
1789010 matches
Allow all ARP packets
290519 matches
Rule Specification
allow() <= protocol('arp')
Metadata
Description:
Allow all ARP packets
Last Modified:
Wed Aug 25 22:50:13 2010 by anonymous
Comment:
Required to allow hosts to find each other's MAC addresses.
Exception:
Protected:
Done
Cancel

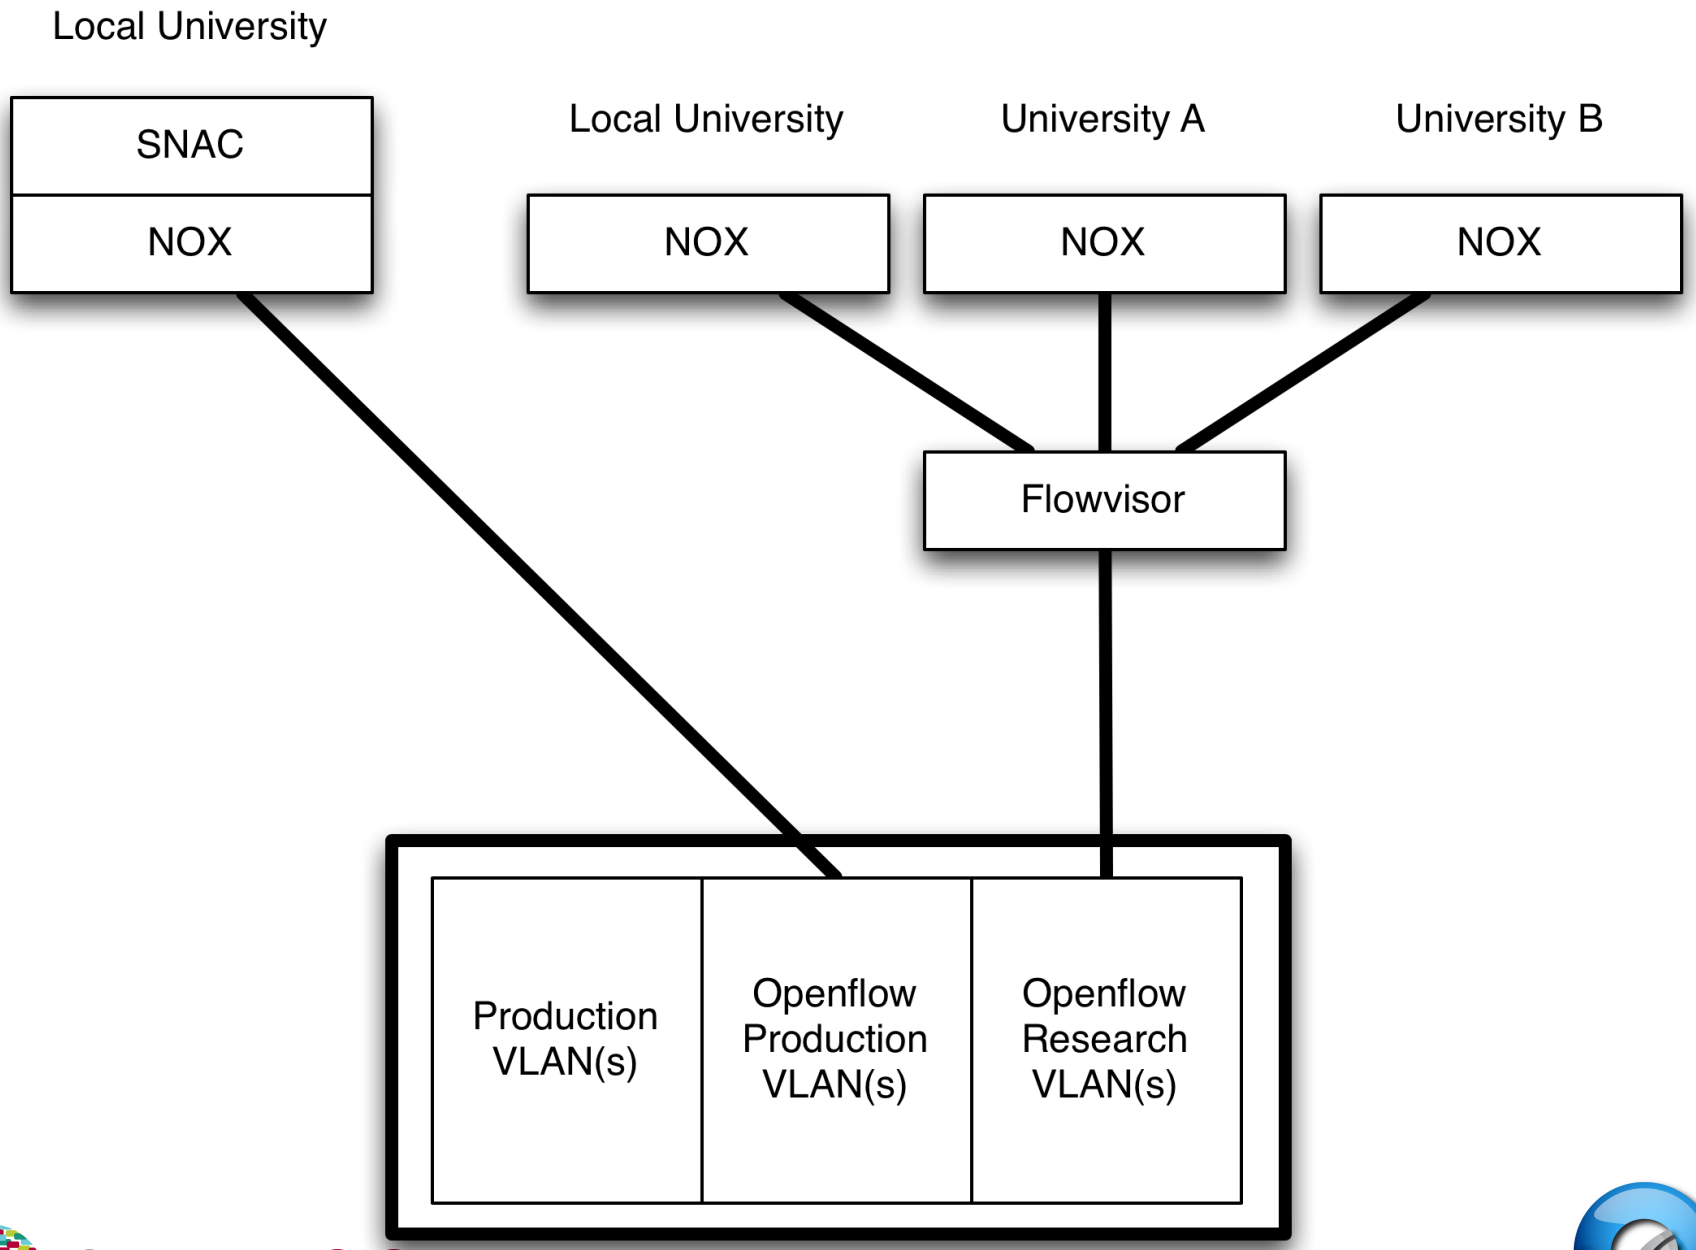

GlobalNOC
Global Research Network Operations Center

Expedient / Opt-In manager

- Software to tie campus OpenFlow deployments to GENI Infrastructure.
- Allows Aggregate Providers (Campus) to make a “sliver” of a switch available to researchers
- Integrates with Flowvisor XMLRPC interface and GENI AAA infrastructure
  - http://www.openflowswitch.org/foswiki/bin/view/OpenFlow/Deployment/HOWTO/ProductionSetup/InstallingExpedientOIM

NetFPGA and Indigo

- NetFPGA
  - FPGA card to test protocols in hardware
  - 4 x 1G and 4 x 10G models
  - OpenFlow 1.0 implementation
  - Google used it for testing OpenFlow-MPLS code
    - http://www.nanog.org/meetings/nanog50/presentations/Monday/NANOG50.Talk17.swhyte_Opensource_LSR_Presentation.pdf
- Indigo
  - Userspace Firmware Reference Release
  - Support for Broadcom chips used in Pronto/Quanta

Switch Issues

- Hw vs Sw rules
- Optional items in OF Spec
  - No one is really implementing rewrite right now
- Control Channel resource exhaustion
- CPU exhaustion and isolation
  - Preventing OF traffic affecting production vlans
- Security
- 48bit vs 64 bit DPIDs
- General strangeness
  - HPs built off live train, NEC uniqueness

Issues

- Inter-operation of different hardware and software
  - Optional items in OF Spec
- Resource exhaustion on switches (CPU, Control channel)
  - Preventing OF traffic affecting production vlans
- Security
- IPv6 Support

OpenVSwitch

http://openvswitch.org

VM-aware virtual switch, run distributed over hardware;

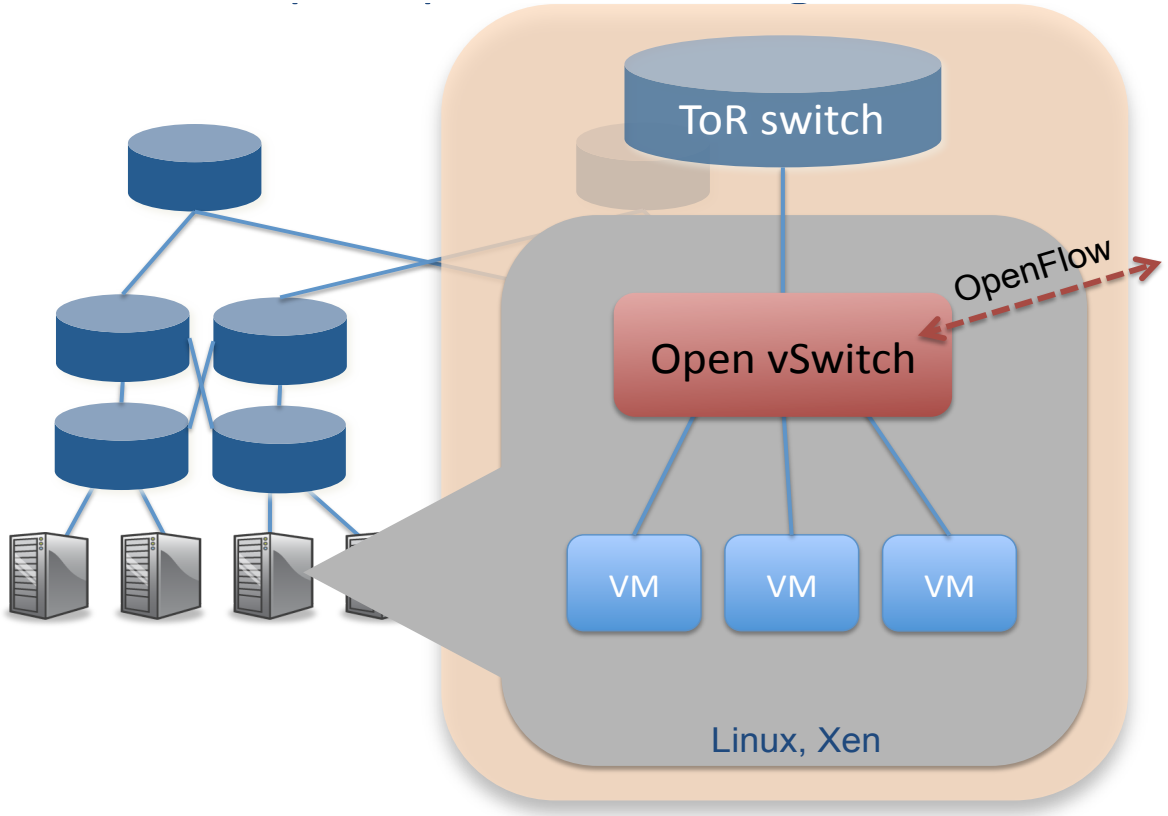

OpenFlow Spec process

http://openflow.org

- V1.0: December 2009
- V1.1: November 2010
  - Open but ad-hoc process among 10-15 companies
- Future

  Planning a more “standard” process from 2011

Measurement Manager

Measurement Manager

- Software built by Indiana University for monitoring OpenFlow networks
- Ties into Flowvisor to get list of devices and topology (using LLDP)
- Acts as OF Controller to gather statistics
- Outputs formats for other tools
  - Nagios (Alarms)
  - GMOC (Topology)
  - SNAPP (Measurement Collector)

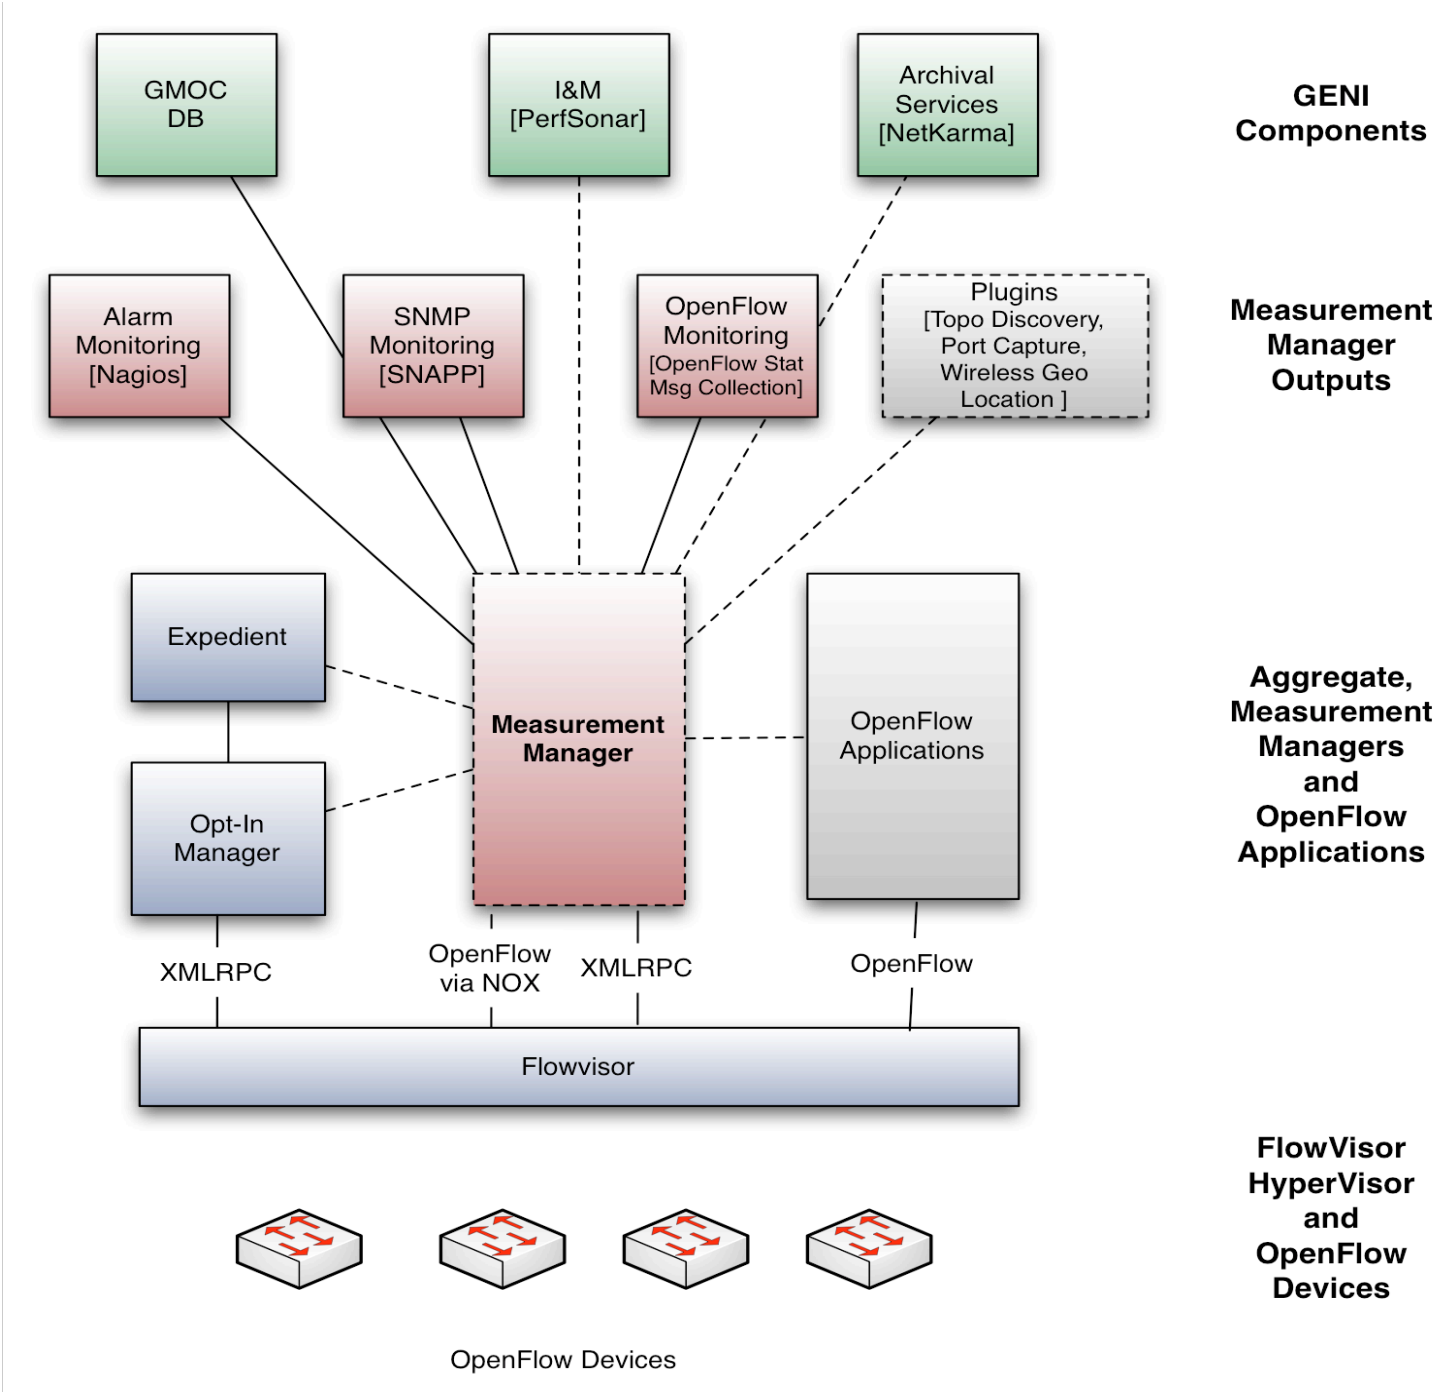
GMOC DB
I&M [PerfSonar]
Archival Services [NetKarma]
GENI Components
Alarm Monitoring [Nagios]
SNMP Monitoring [SNAPP]
OpenFlow Monitoring [OpenFlow Stat Msg Collection]
Plugins [Topo Discovery, Port Capture, Wireless Geo Location]
Measurement Manager Outputs
Expedient
Opt-In Manager
Measurement Manager
OpenFlow Applications
Aggregate, Measurement Managers and OpenFlow Applications
XMLRPC
OpenFlow via NOX
XMLRPC
OpenFlow
Flowvisor
FlowVisor HyperVisor and OpenFlow Devices
OpenFlow Devices

Measurement Manager

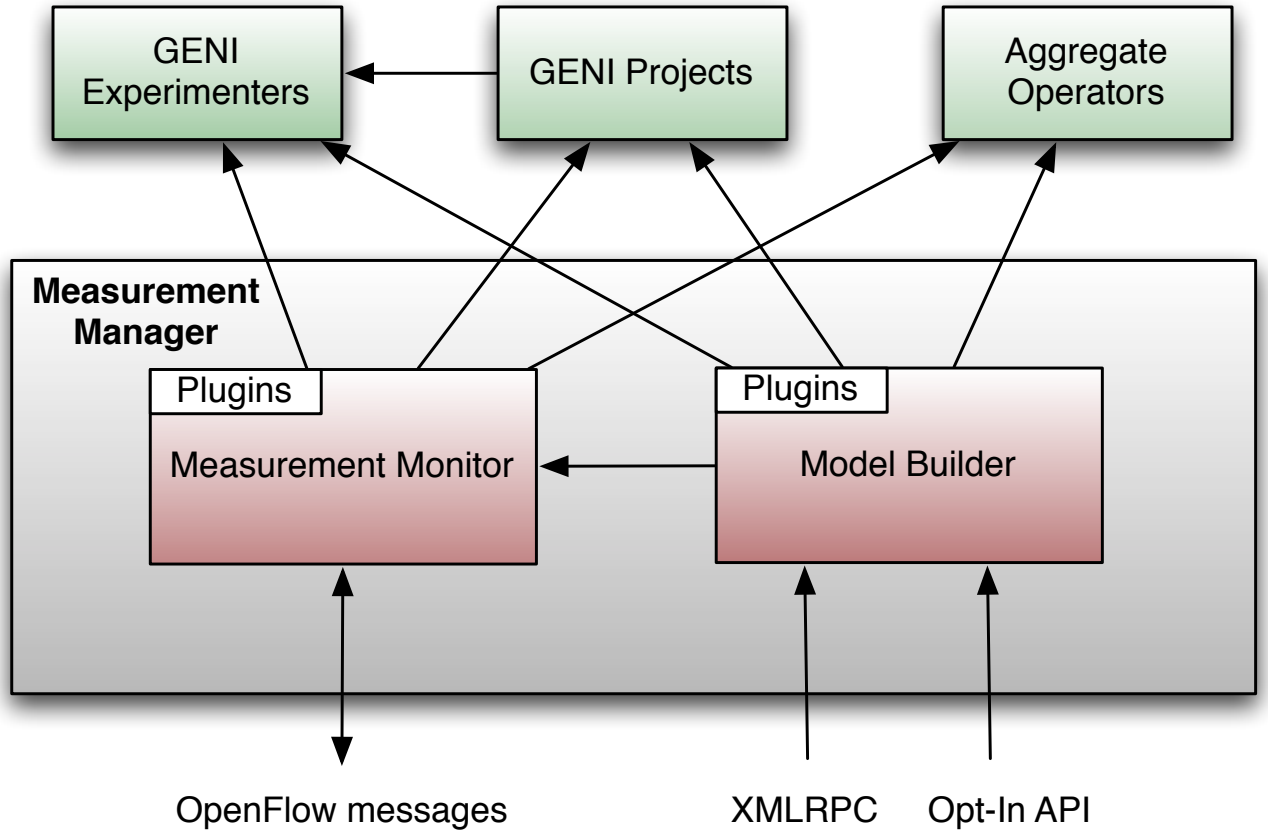

What will can do with OpenFlow ?

- 1k-3k TCAM Entries in Typical Edge Switch
- Difficult to take advantage of:
  - Manual Config, SNMP Writes, RADIUS
  - Limited Actions (allow/deny)
  - Vendor Specific
- But what if you could program these through a standard API ?

Possible Uses of Openflow (Quick Wins)

- Security Applications
  - NAC
  - IDS/IPS
  - Remote Packet Capture & Injection
- VM Mobility
  - Redirect specific application traffic to remote site
  - Flow-based forwarding – no need to extend entire broadcast domain – no STP issues

Other Applications

- Load Balancing
- n-cast
  - multiple streams over lossy networks
- Policy (Firewall)
- Flow based network provisioning

Intercontinental VM Migration

Moved a VM from Stanford to Japan without changing its IP.
VM hosted a video game server with active network connections.

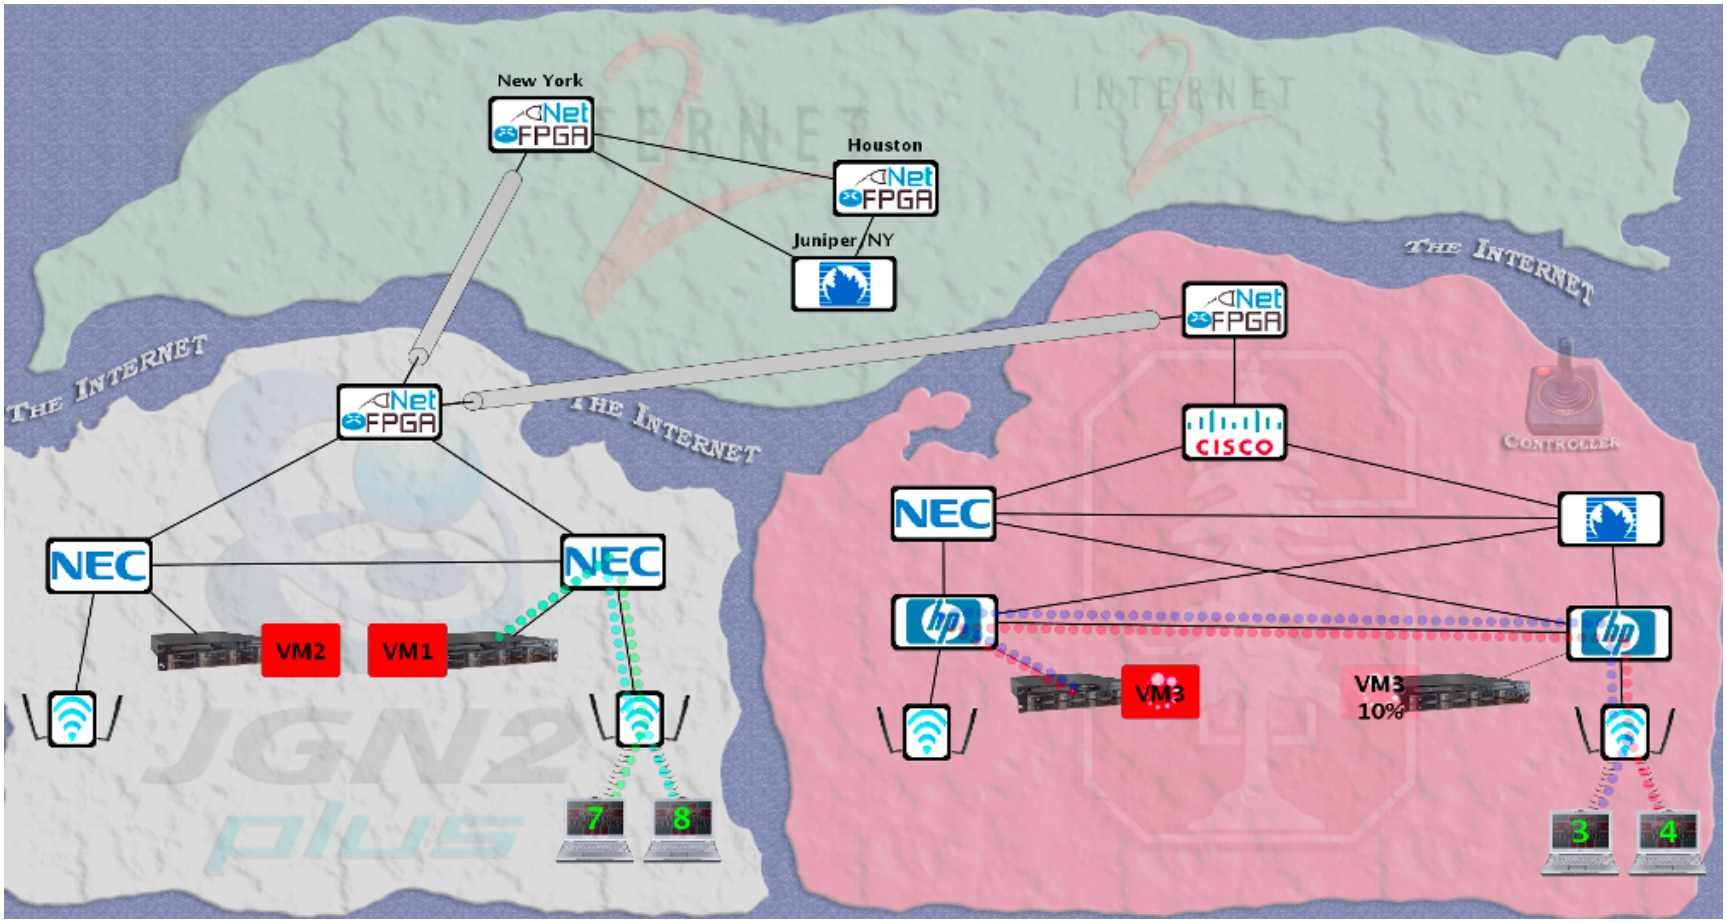

GlobalNOC
Global Research Network Operations Center

Possible Uses of Openflow (Quick Wins)

- Dynamic Circuit Provisioning
  - Don't need to extend layer-2 end-to-end
  - Simply direct specific flows down a engineered path with guaranteed priority
  - Don't have to rely on scripted SSH sessions, SNMP or other sub-optimal ways to programmatically configure switches/routers.

Possible Uses of Openflow (Grand Challenges)

- Distributed Control-Plane Architecture Requires a Lot of State to be Synchronized Across Many Devices
- Many Protocols Needed for Synchronization Internally to Networks (OSPF, RSVP, STP, etc)
- Can these “internal” protocols eventually be removed entirely with only BGP for inter-domain route advertisements ?

OpenFlow Paradigm shifts

- "Wireless like" management of wired switches
- Manipulate virtual switches over many physical devices
  - VM Migration demo
- OSI model breakdown
- Control at the flow level

Workshop Demos

- Mininet Introduction – Tutorial VM
  - http://www.openflowswitch.org/wk/index.php/HOTITutorial2010
- Multiple switch control using single CLI
- VM Migration Demo
  - Moving a VM between subnets
  - Simplified version of other VM migration demos
- Measurement Manager showing Backbone Deployments
  - Topology and Statistic collection in a controller based environment

Mininet Demo

OpenFlow Tutorial: 3hosts-1switch topology

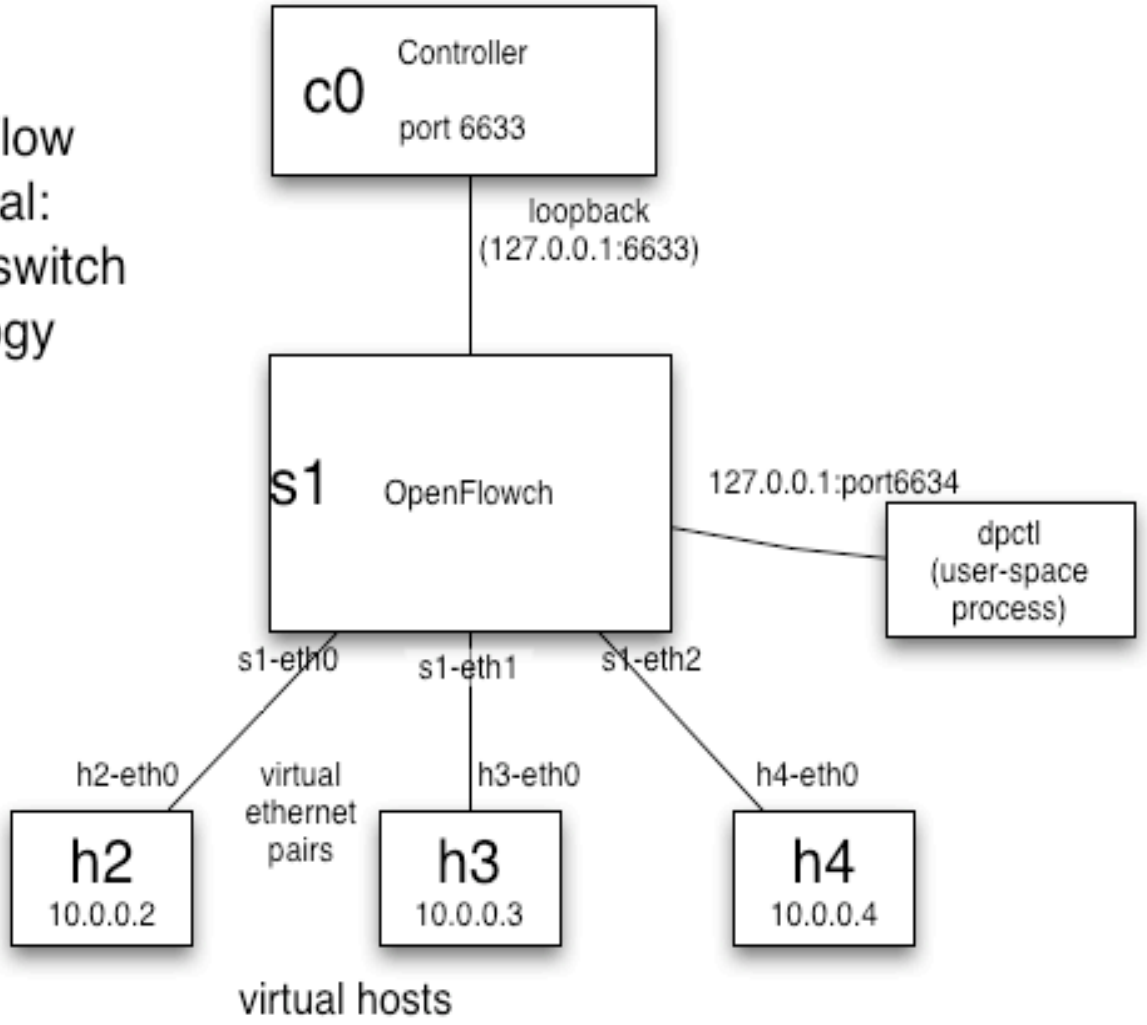

GlobalNOC
Global Research Network Operations Center

Single CLI Demo

- Run a CLI commands over multiple physical switches
- Manipulate flow rules to block certain traffic

VM Migration Demo

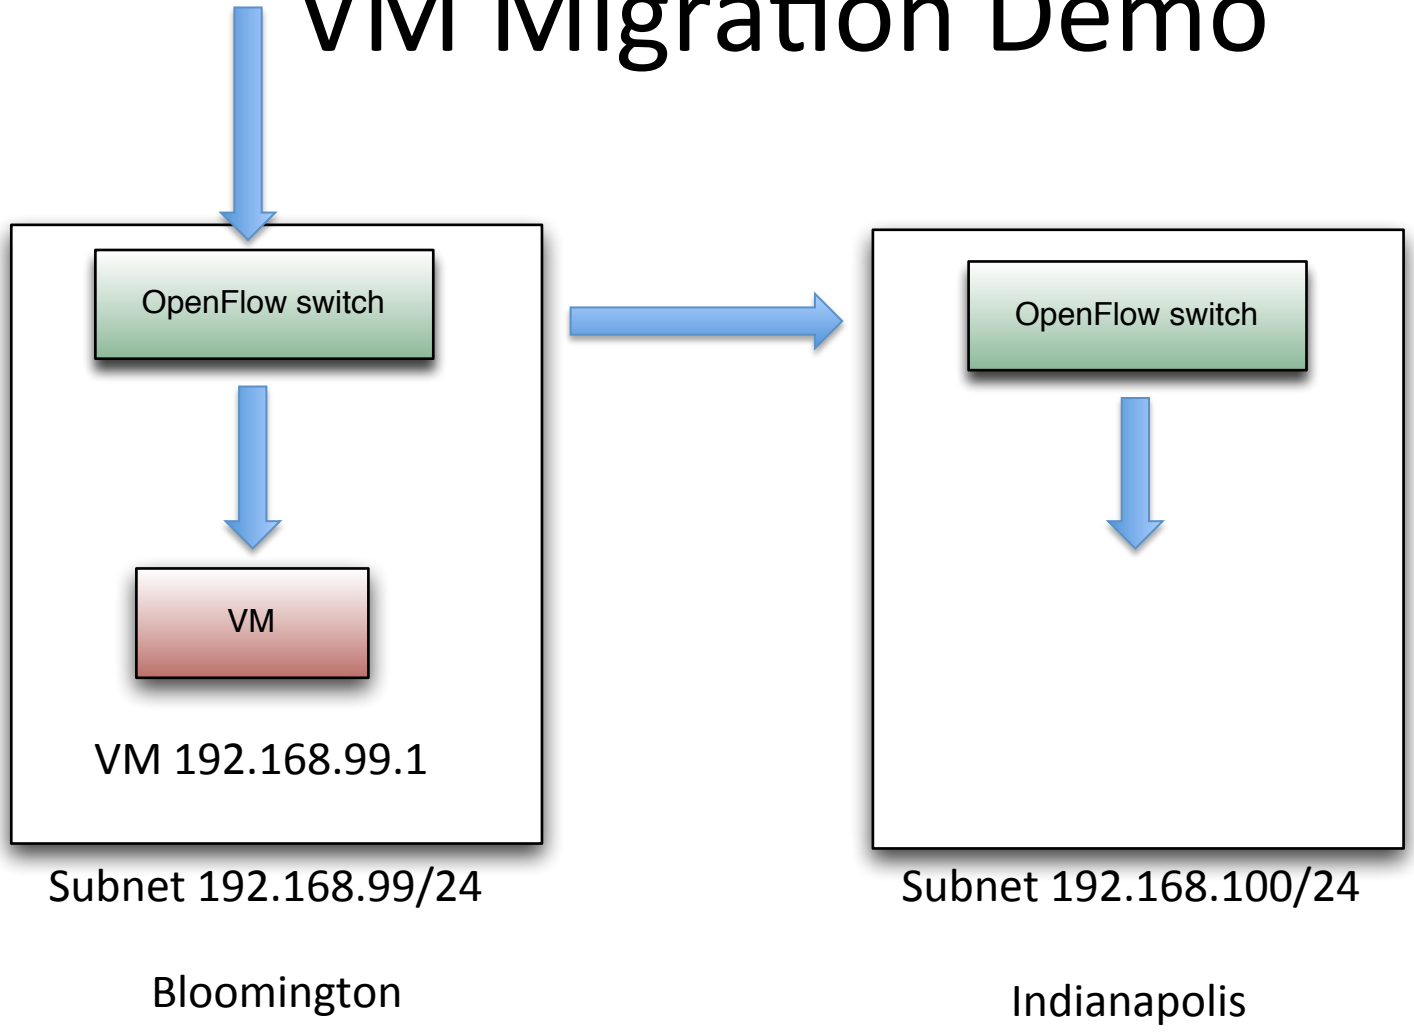

Measurement Manager Demo

- Topology – using Google Earth
  - http://gmoc-db.grnoc.iu.edu
  - Select OpenFlow Aggregate
- Nagios data collection
  - http://gmoc-db.grnoc.iu.edu/nagios
- SNAPP Statistics
  - http://gmoc-db.grnoc.iu.edu/nlr-of/

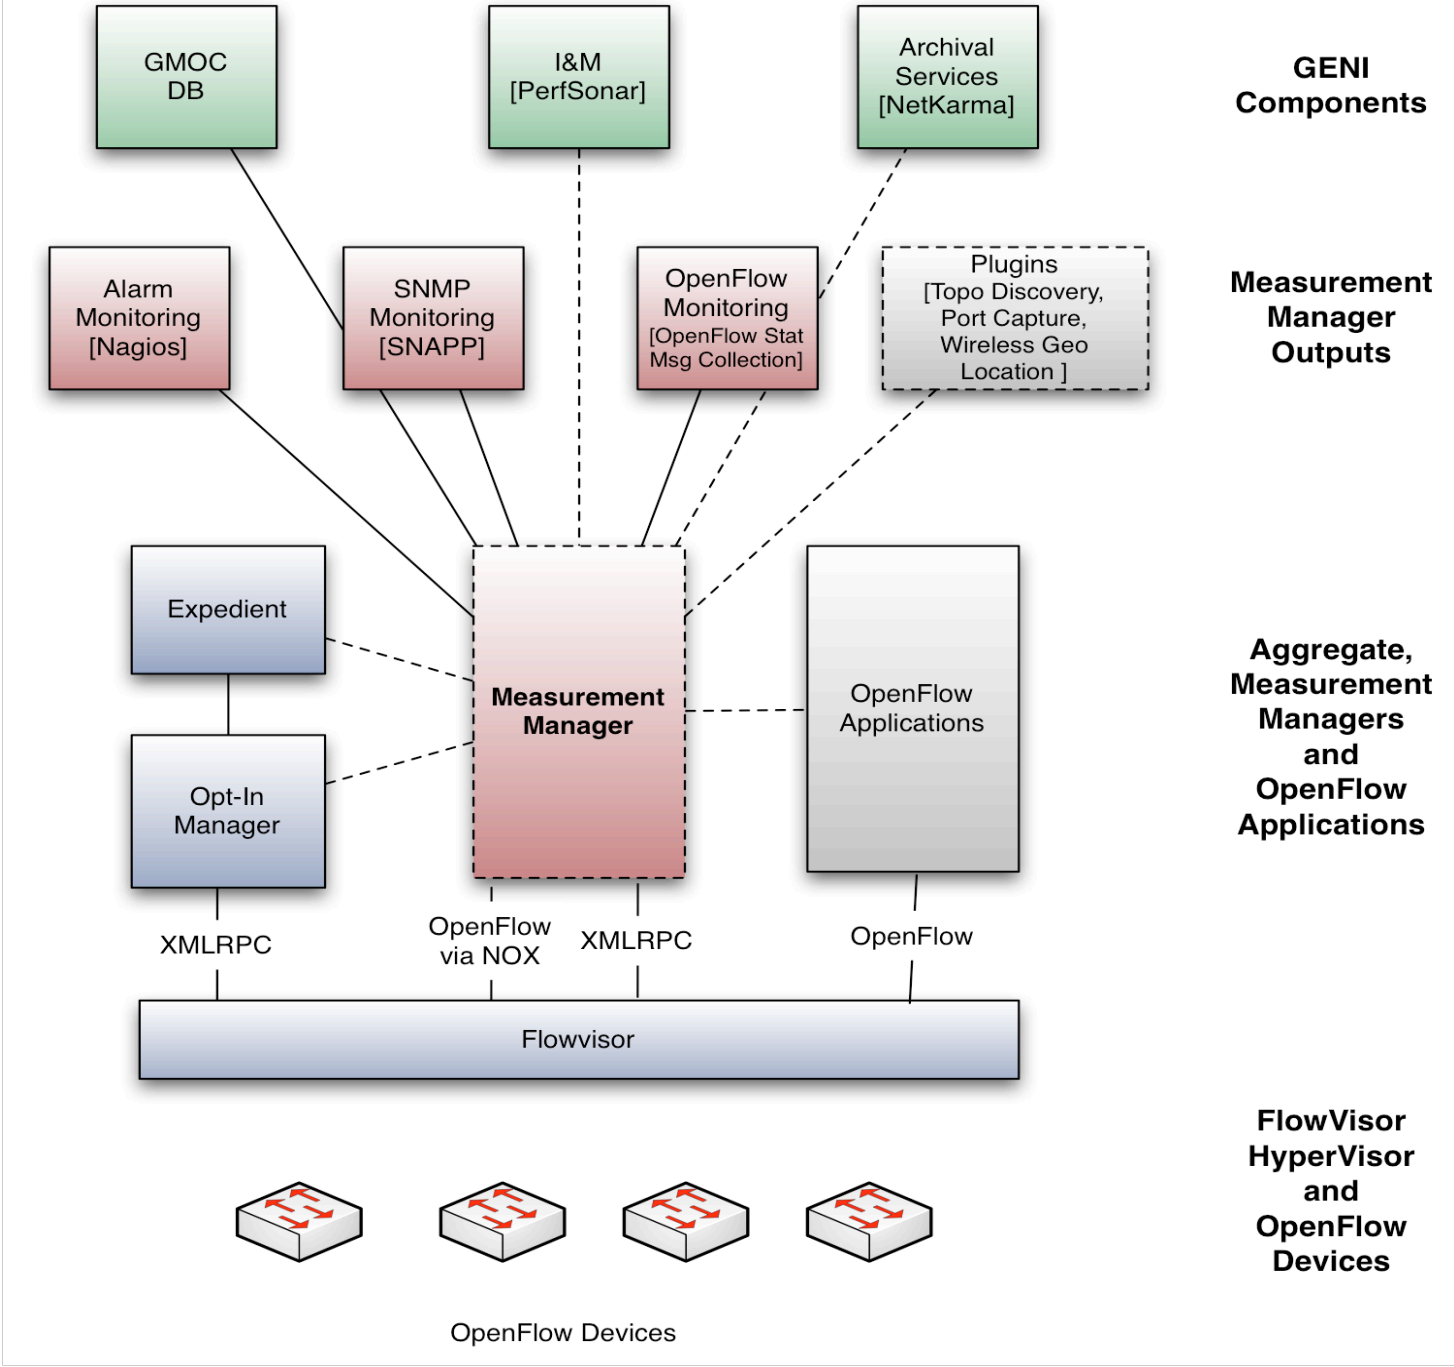
GMOC DB
I&M [PerfSonar]
Archival Services [NetKarma]
GENI Components
Alarm Monitoring [Nagios]
SNMP Monitoring [SNAPP]
OpenFlow Monitoring [OpenFlow Stat Msg Collection]
Plugins [Topo Discovery, Port Capture, Wireless Geo Location]
Measurement Manager Outputs
Expedient
Measurement Manager
OpenFlow Applications
Opt-In Manager
Aggregate, Measurement Managers and OpenFlow Applications
XMLRPC
OpenFlow via NOX
XMLRPC
OpenFlow
Flowvisor
FlowVisor HyperVisor and OpenFlow Devices
OpenFlow Devices

How to get involved

- Experiment with Controllers
  - NOX: http://noxrepo.org
  - Beacon: http://www.openflowhub.org/
- Switches
  - Soft switches / Mininet
  - Hardware switches you already may have
  - Deploy Applications

More Information sources

- OpenFlow
  - http://openflowswitch.org

- My contact info

Chris Small -- Indiana University

E-mail: chsmall@indiana.edu

Discussion and Questions?