OpenFlow Overview

GENI Engineering Conference 12

Nov 1 2011

Sections

- OpenFlow Introduction
- OpenFlow tools - hardware and software
- Deployments
  - GENI
  - NDDI
  - Campus
- Use Cases
- Discussion

InCNTRE

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

InCNTRE

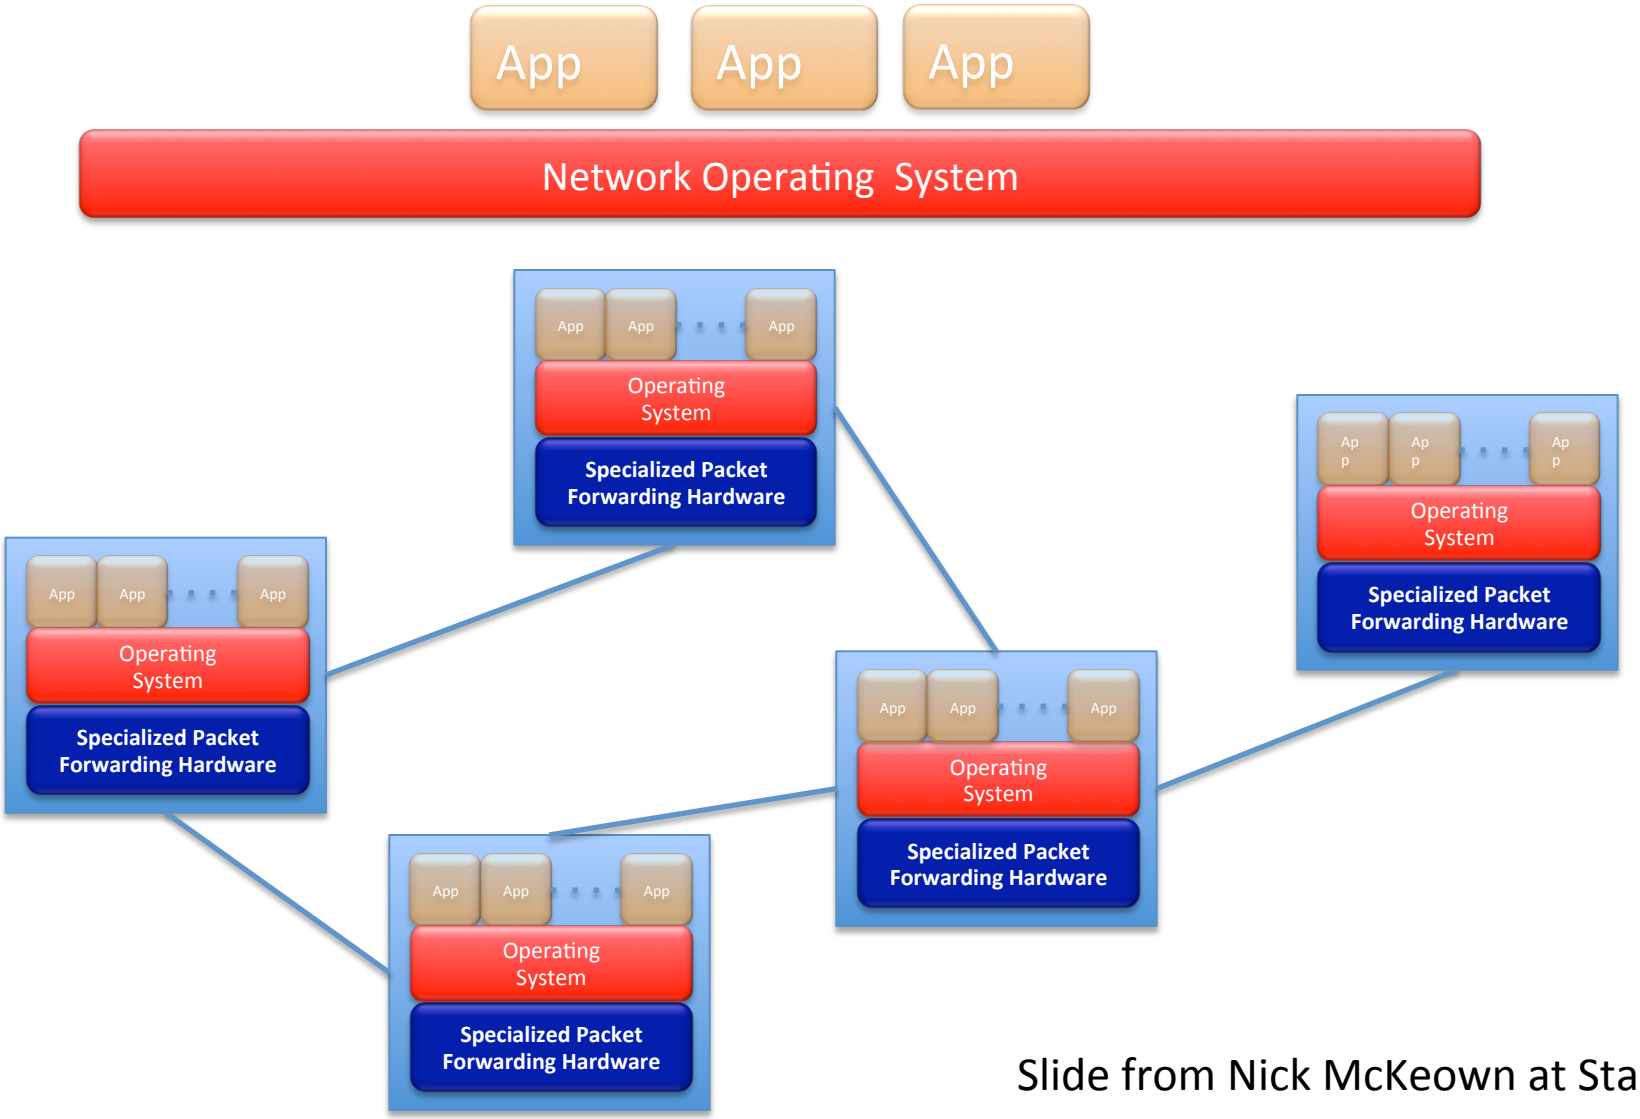

Slide from Nick McKeown at Stanford

The "Software-defined Network"

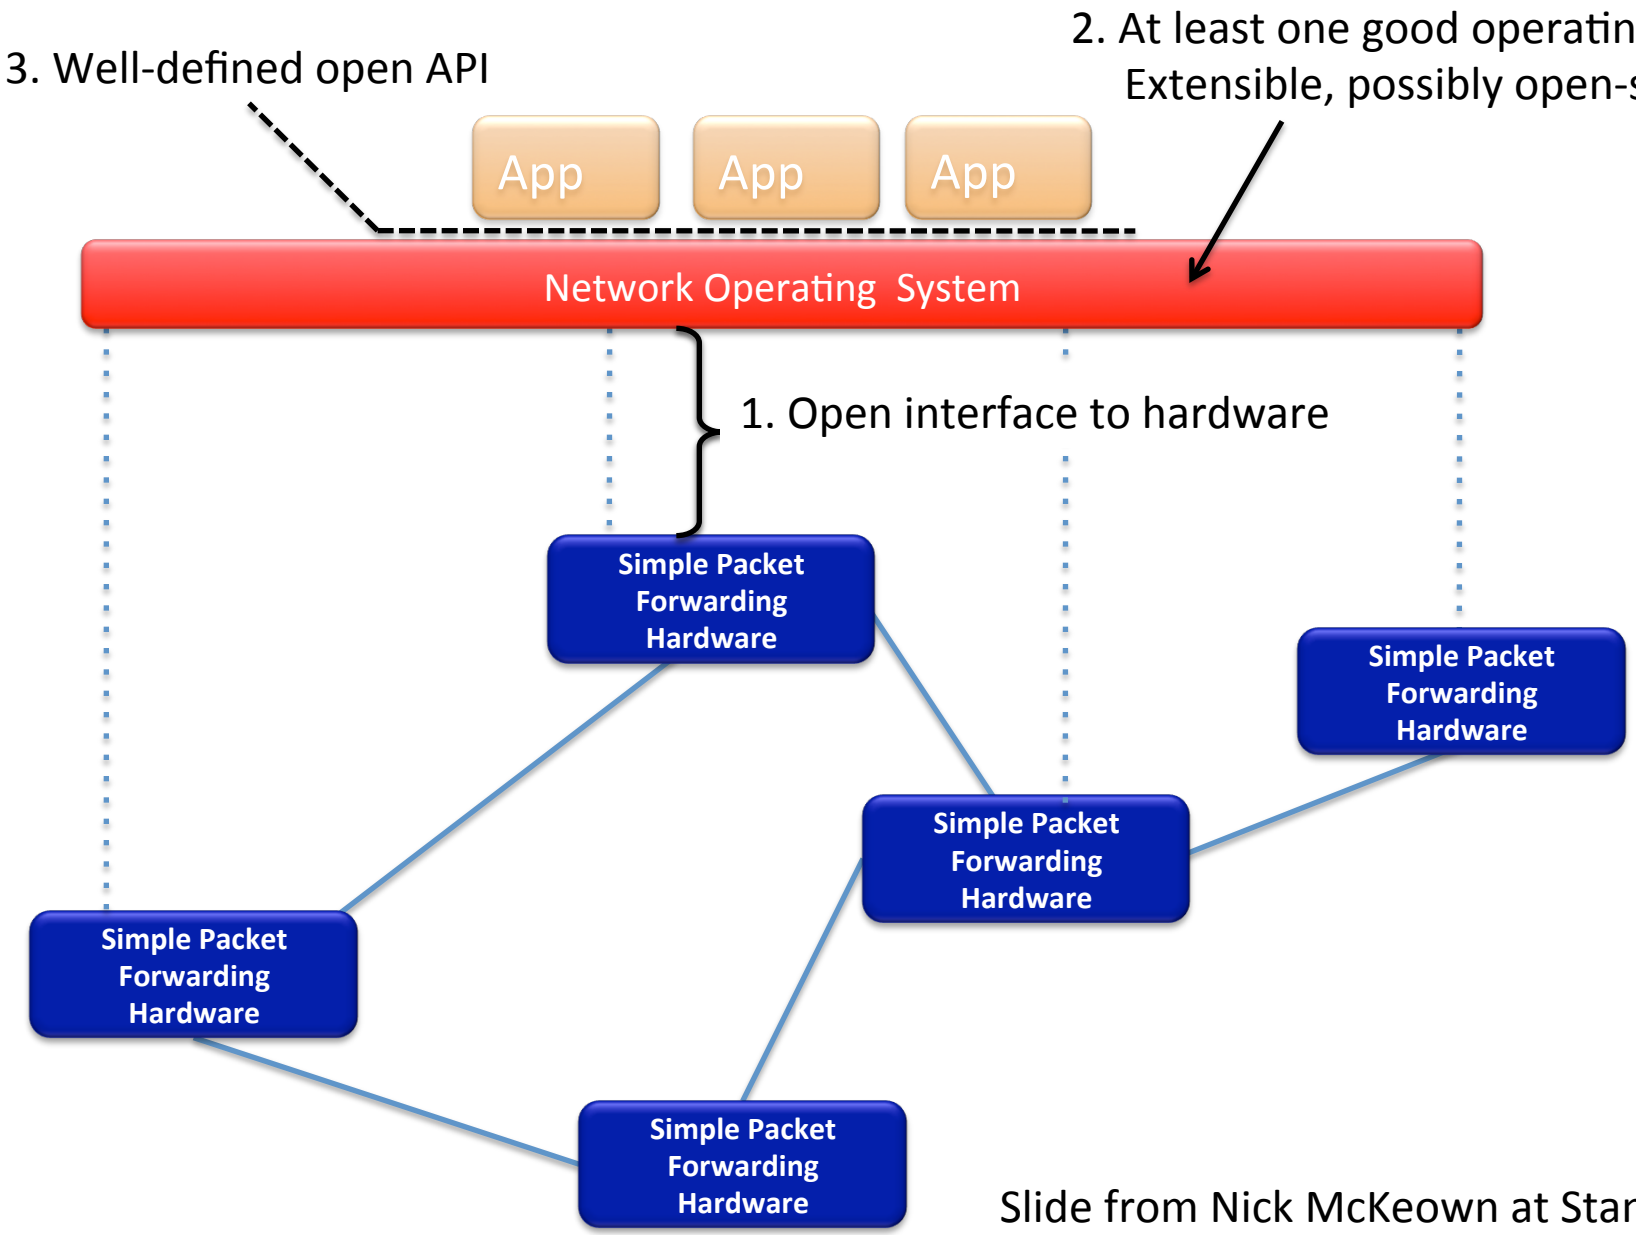

Slide from Nick McKeown at Stanford

Trend

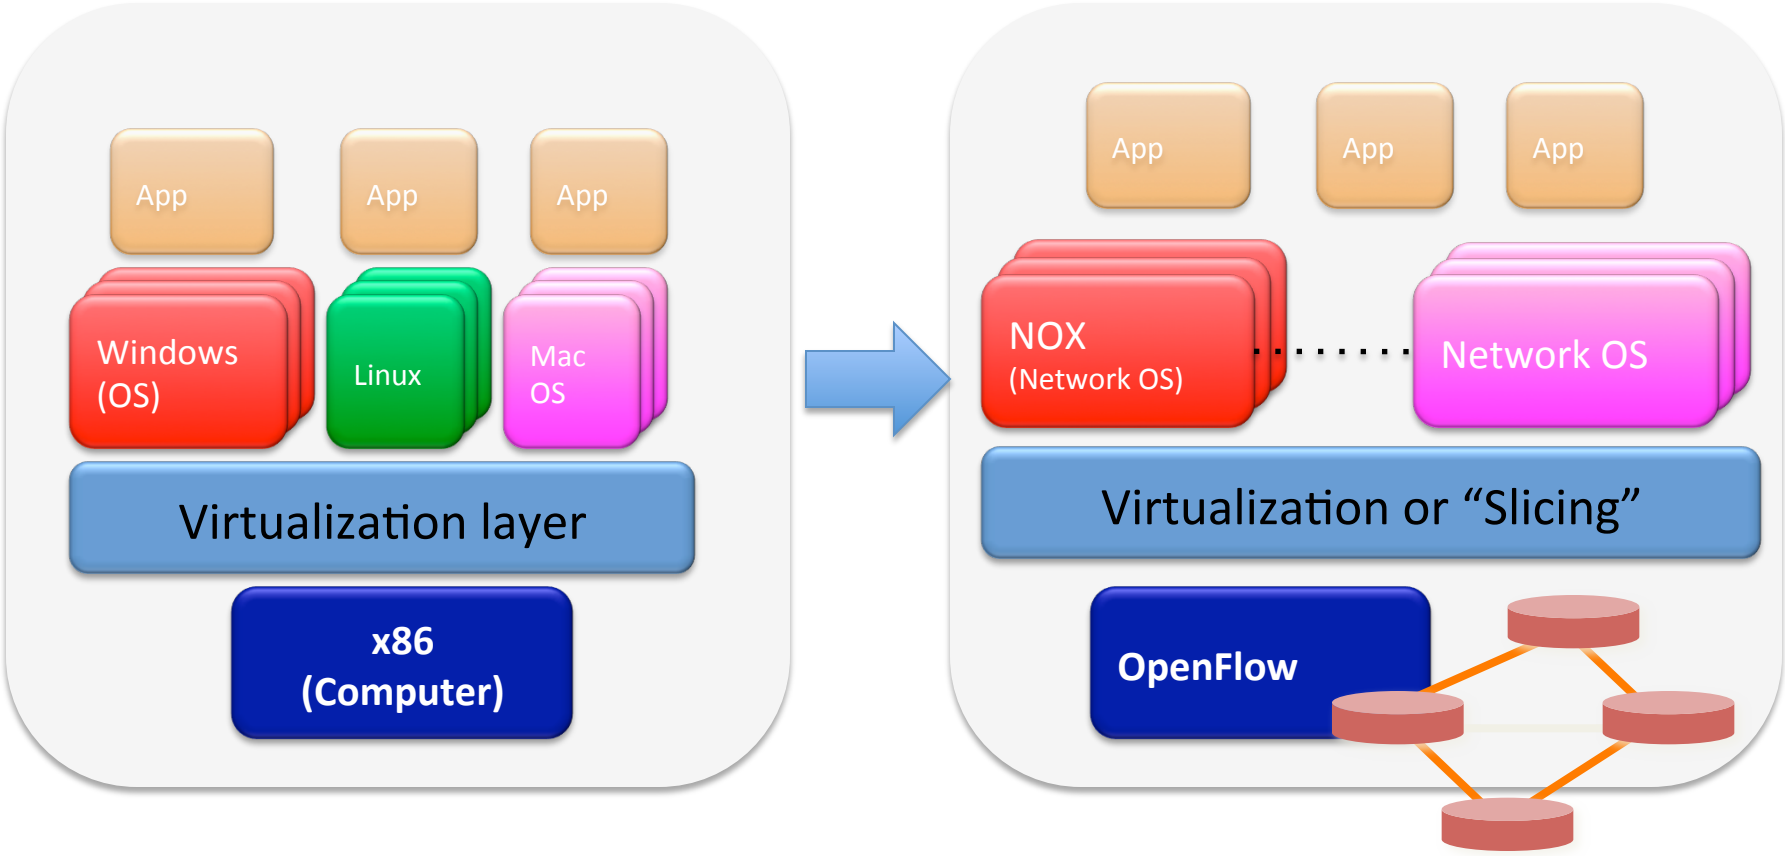

Computer Industry

Network Industry

Slide from Nick McKeown at Stanford

OpenFlow Basics

OpenFlow Basics (1)

Exploit the flow table in switches, routers, and chipsets

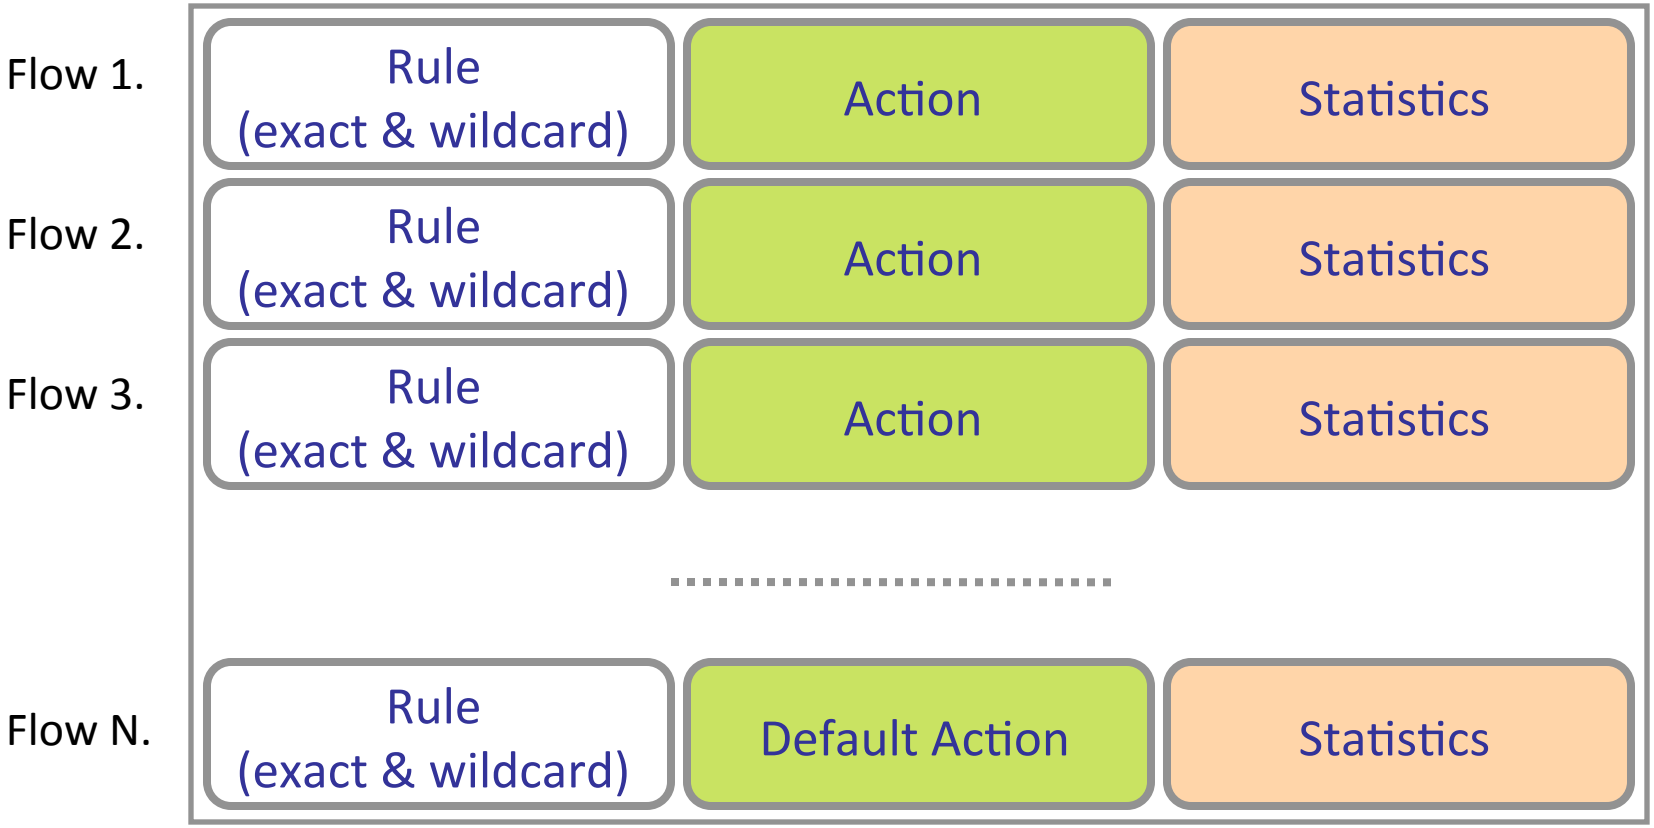

OpenFlow Basics (2)

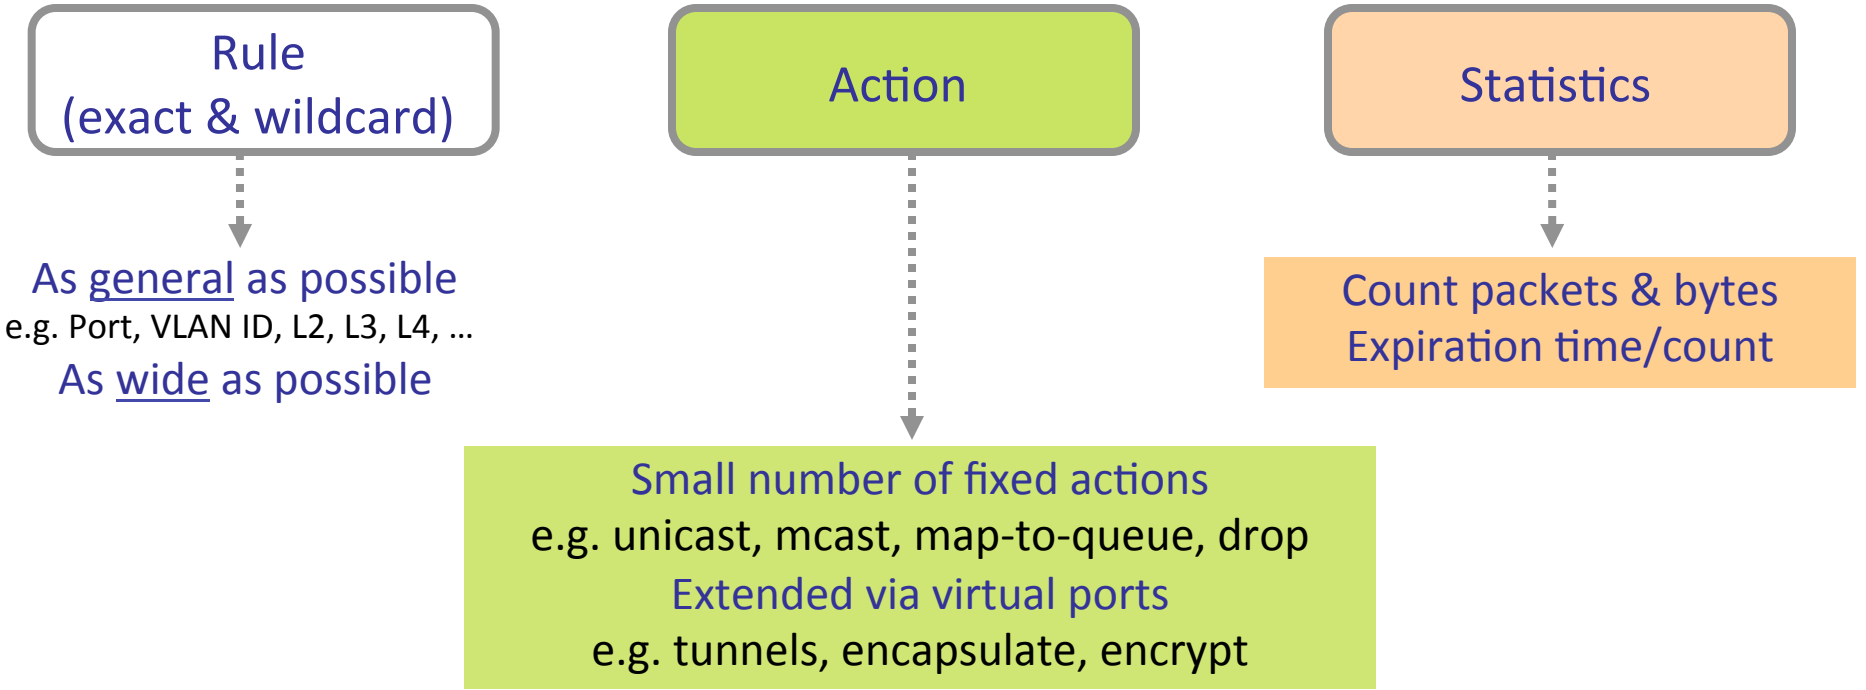

InCNTRE

Flow Table Entry

OpenFlow 1.0 Switch

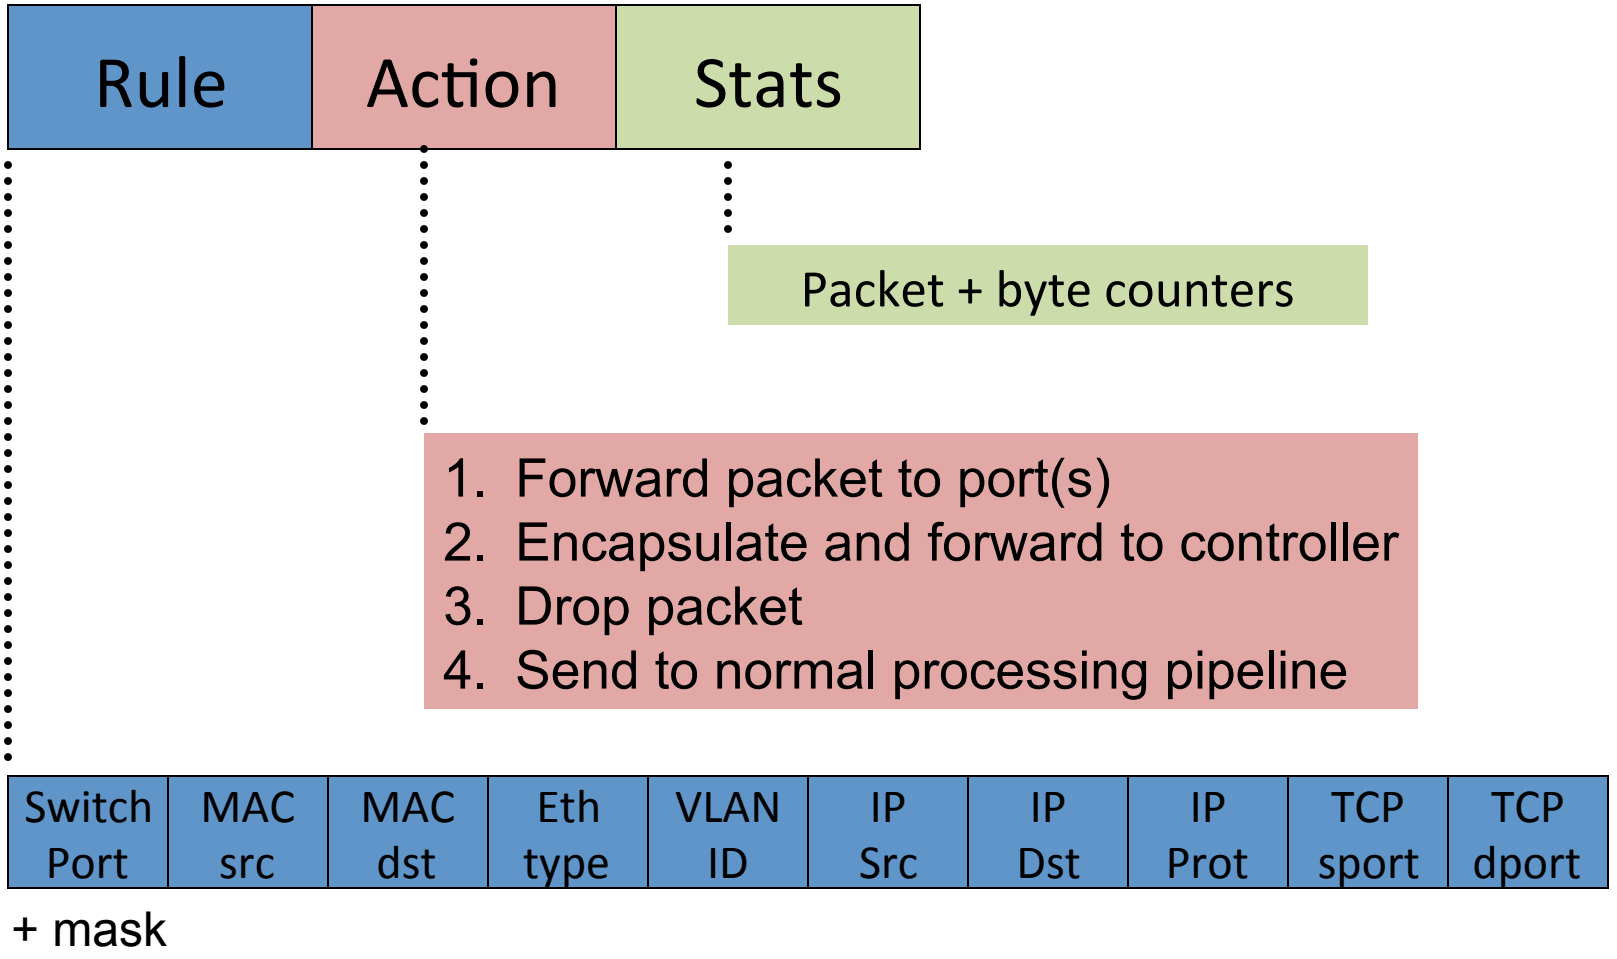

OpenFlow Basics (3)

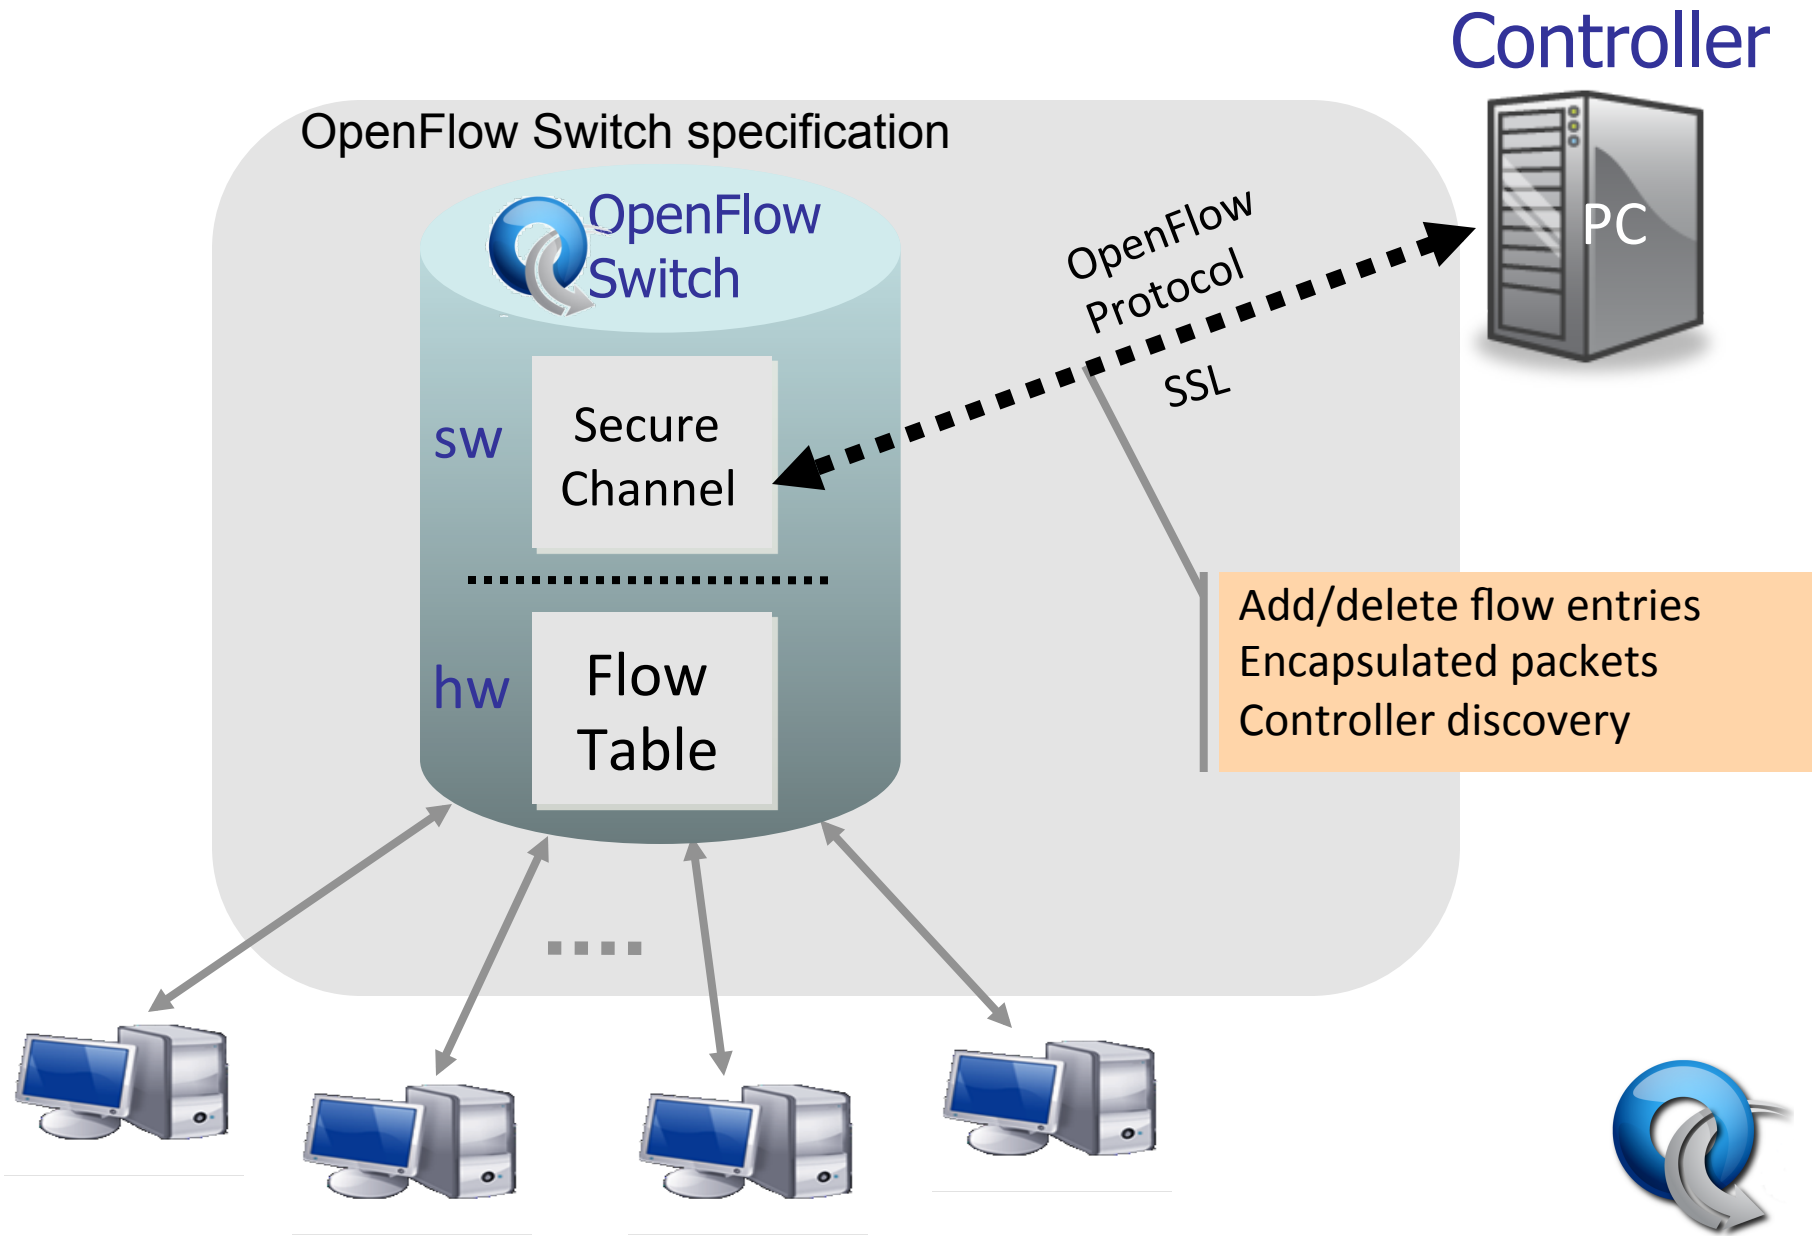

OpenFlow Usage

Dedicated OpenFlow Network

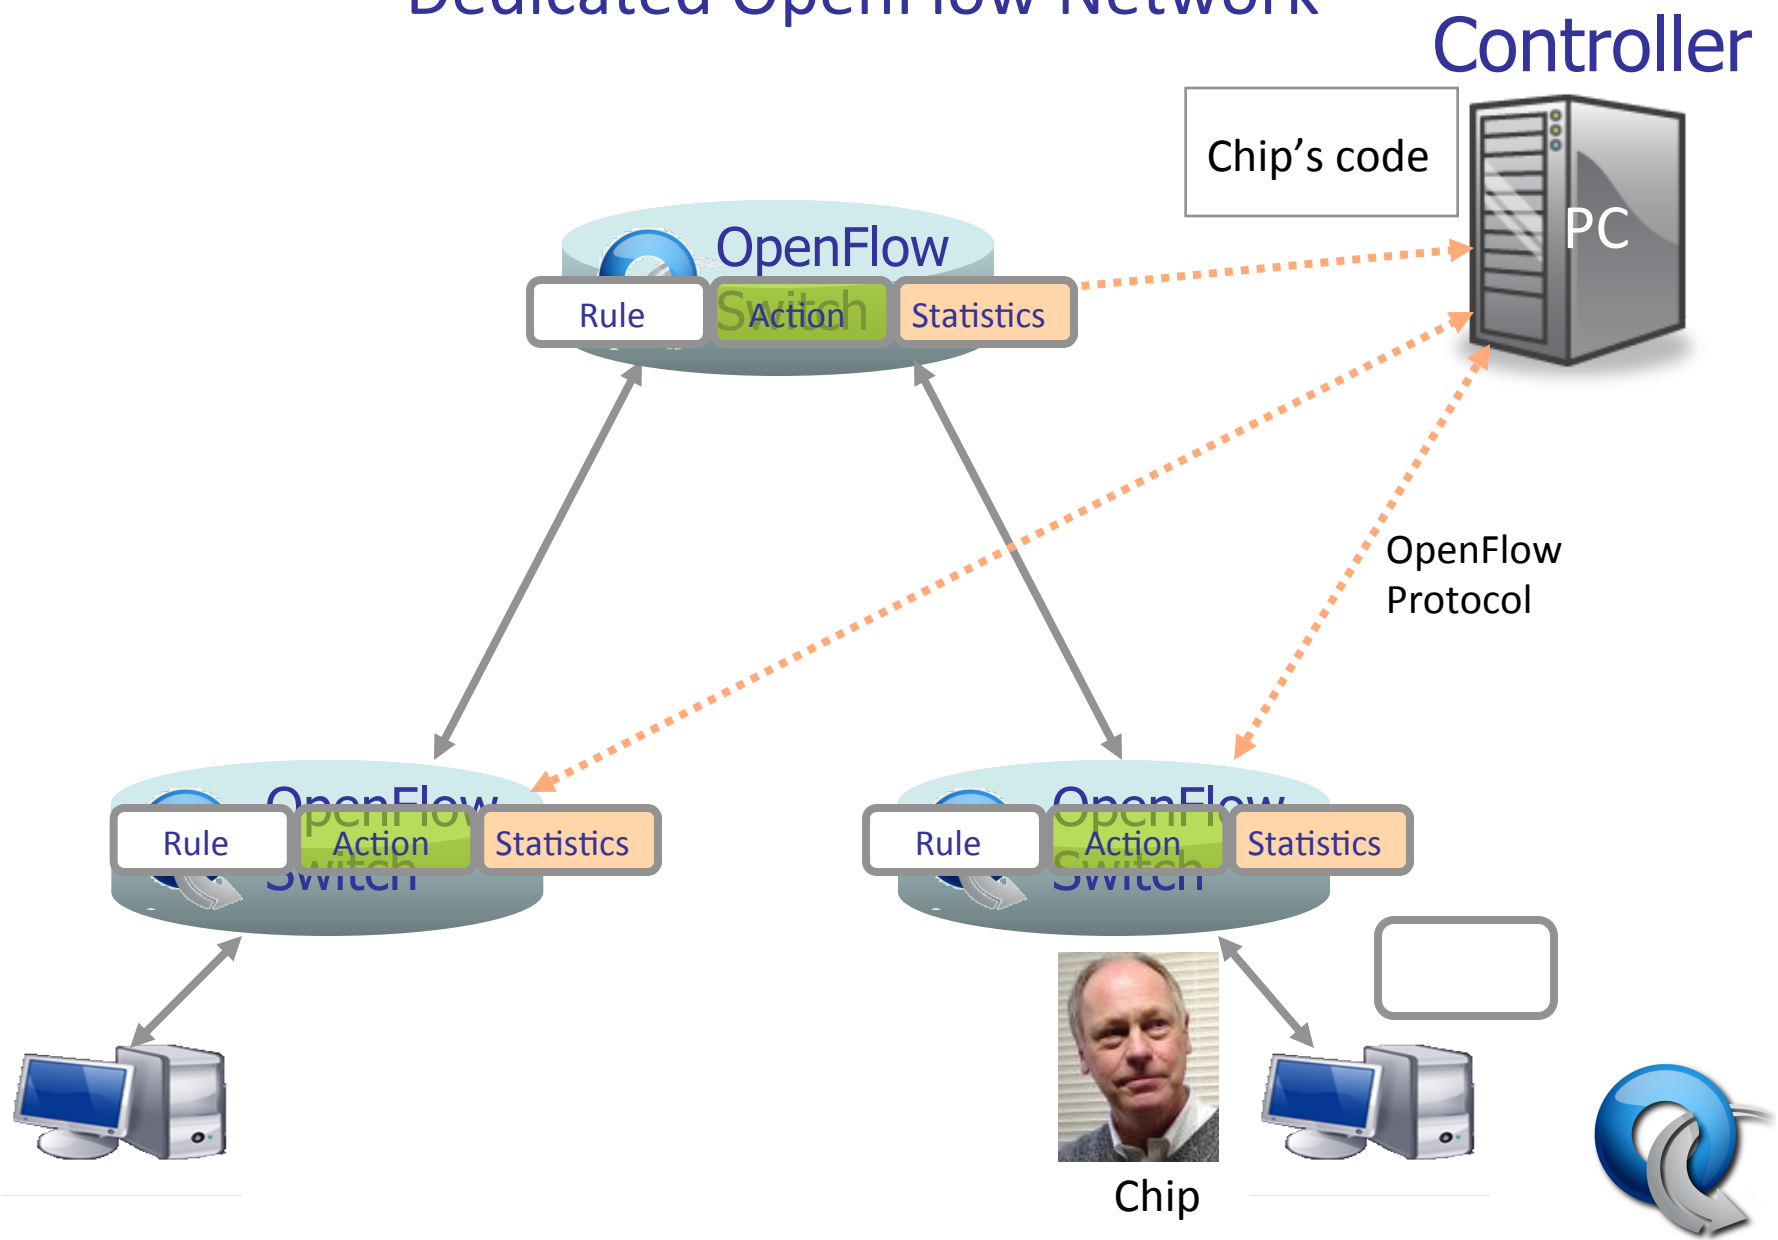

OpenFlow Deployments

GENI

- GENI OpenFlow deployment on 8 campuses
  - Will be increasing to 25 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 (200?)

InCNTRE

OpenFlow and GENI

8 Universities, GPO/BBN, & 2 National Backbones

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

InCNTRE

NLR – I2 OpenFlow Core

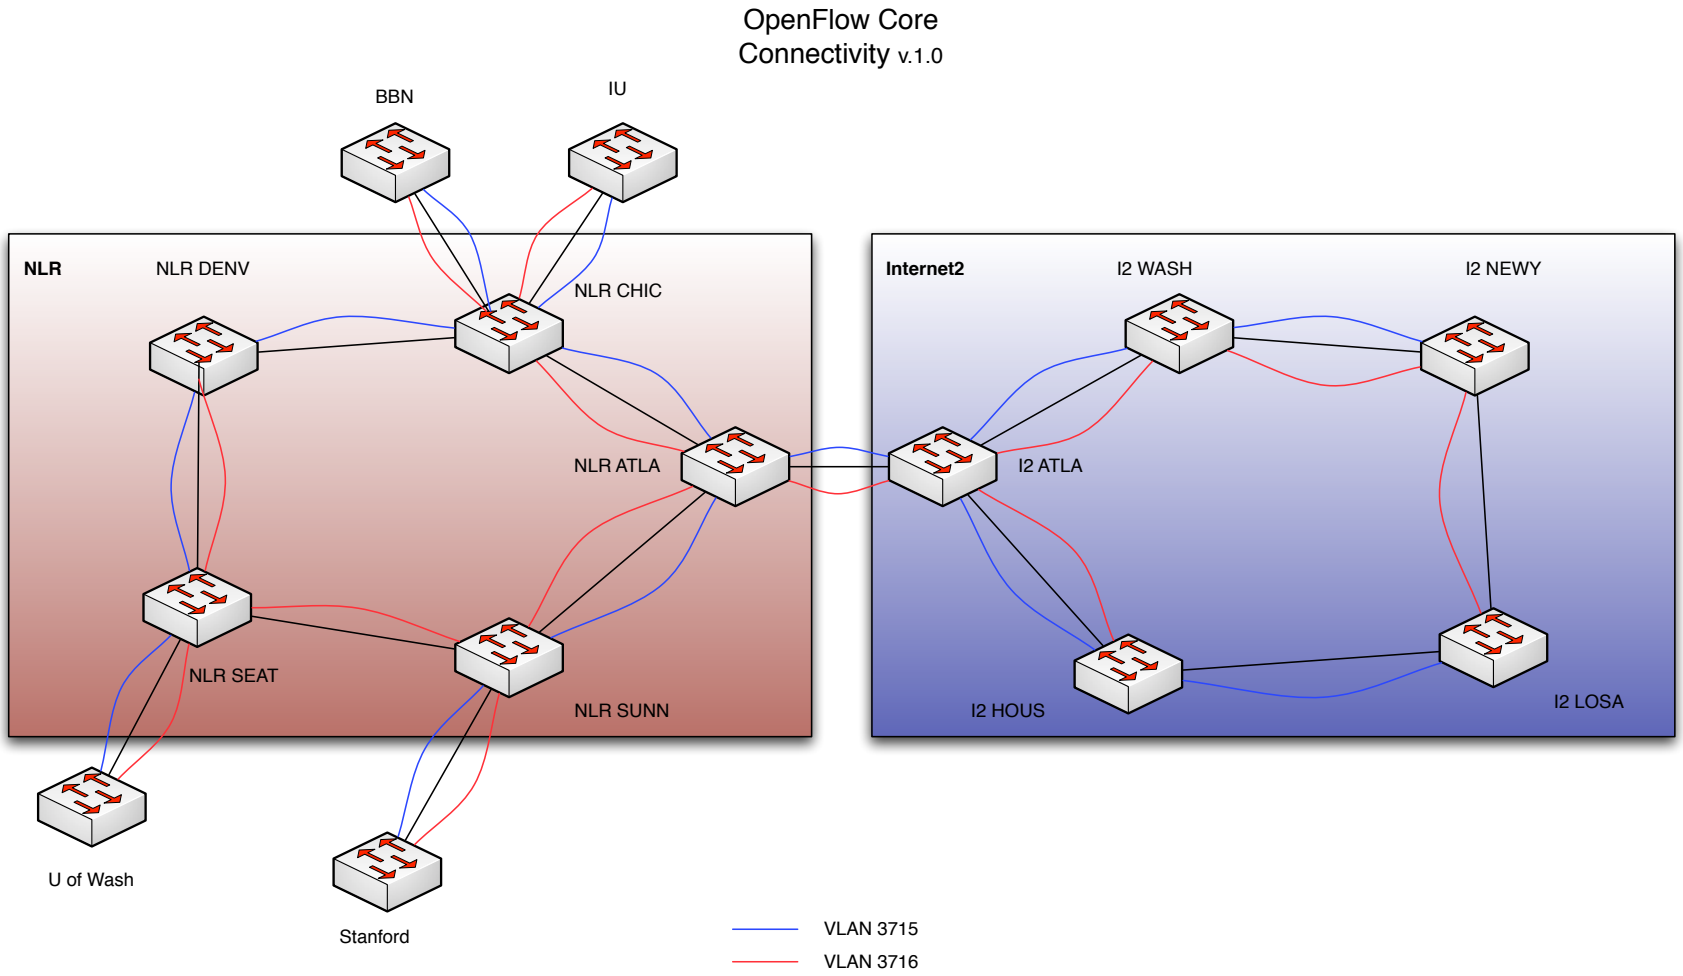

NLR – I2 Backbone

- 2 Main Vlans – 3515 and 3516
- Arranged so loops-free in the backbone over either vlan
- Campuses can cause loops though ☹

OS^{3}E

- OS^{3}E – Open Science, Scholarship and Services Exchange
  - Production Layer 2 Provisioning tool over Internet2 NDDI Infrastructure
  - OpenFlow based
  - Developed by IU
  - Demo'd at I2 Fall Members Meeting
    - https://os3e.net.internet2.edu
    - user: os3e pw: os3edemo

InCNTRE


OS3E
The Open Science, Scholarship & Services Exchange
Vancouver
Seattle
Portland
Missoula
Minneapolis
Chicago
Cleveland
Pittsburgh
Philadelphia
Boston
Buffalo
New York
Ashburn
Washington DC
Salt Lake City
Sunnyvale
Denver
Kansas City
Indianapolis
Louisville
Raleigh
Los Angeles
Albuquerque
Nashville
Memphis
Atlanta
Tucson
El Paso
Dallas
Jackson
Jacksonville
Houston
Baton Rouge
Miami
INTERNET2
INDIANA UNIVERSITY
STANFORD
UNIVERSITY

InCNTRE

IU Campus Deployment

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

Sample Campus Deployment (Indiana University)

- HP switches in Testlab and Production
  - 2 6600s in Bloomington testlab
  - 2 6600s in Indy SDN lab
  - 1 5406 Testlab/Wireless
  - 2 5406 used by Engineering
  - 3500 in Gigapop, IUPUI Wireless
- 5 Pronto switches (3290s) in testlab
- 2 Pronto 10G switches in production soon
- Netgear, NetFPGA 10G and 1G

InCNTRE

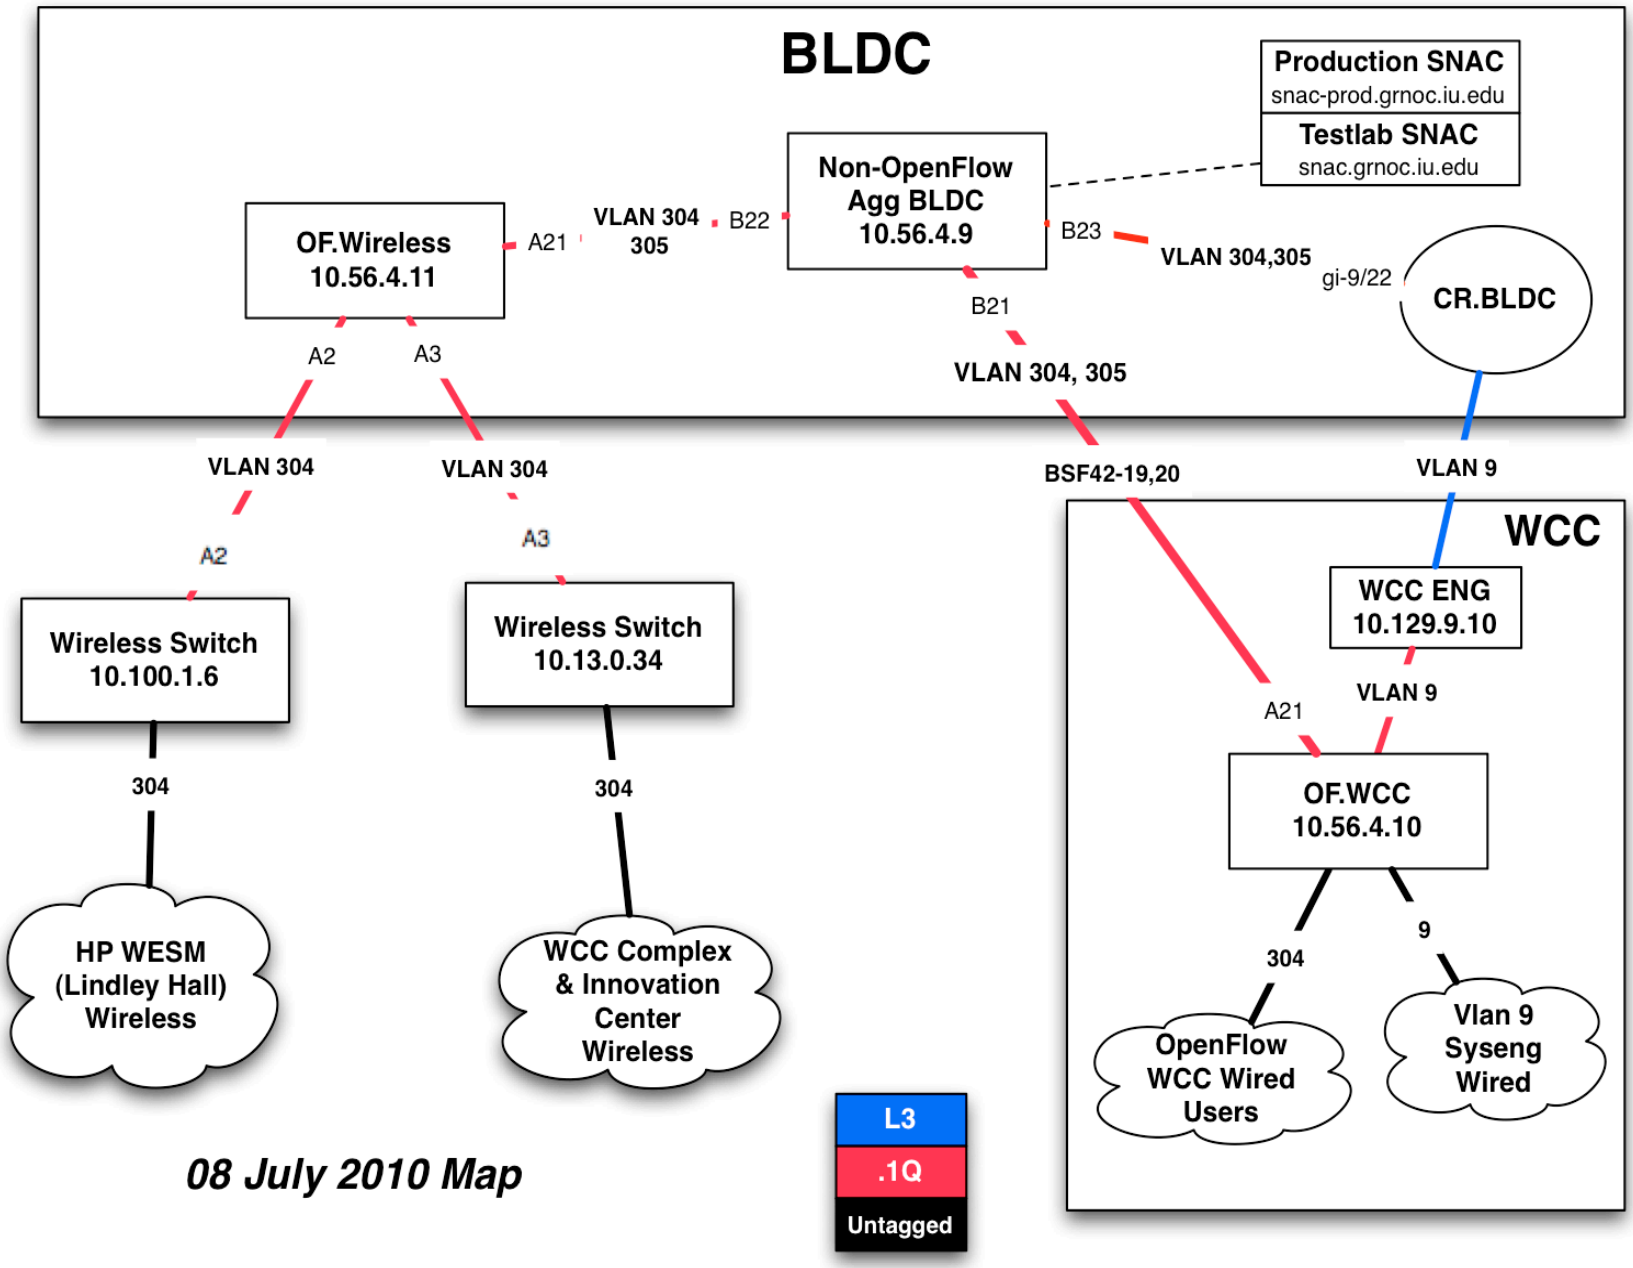

08 July 2010 Map

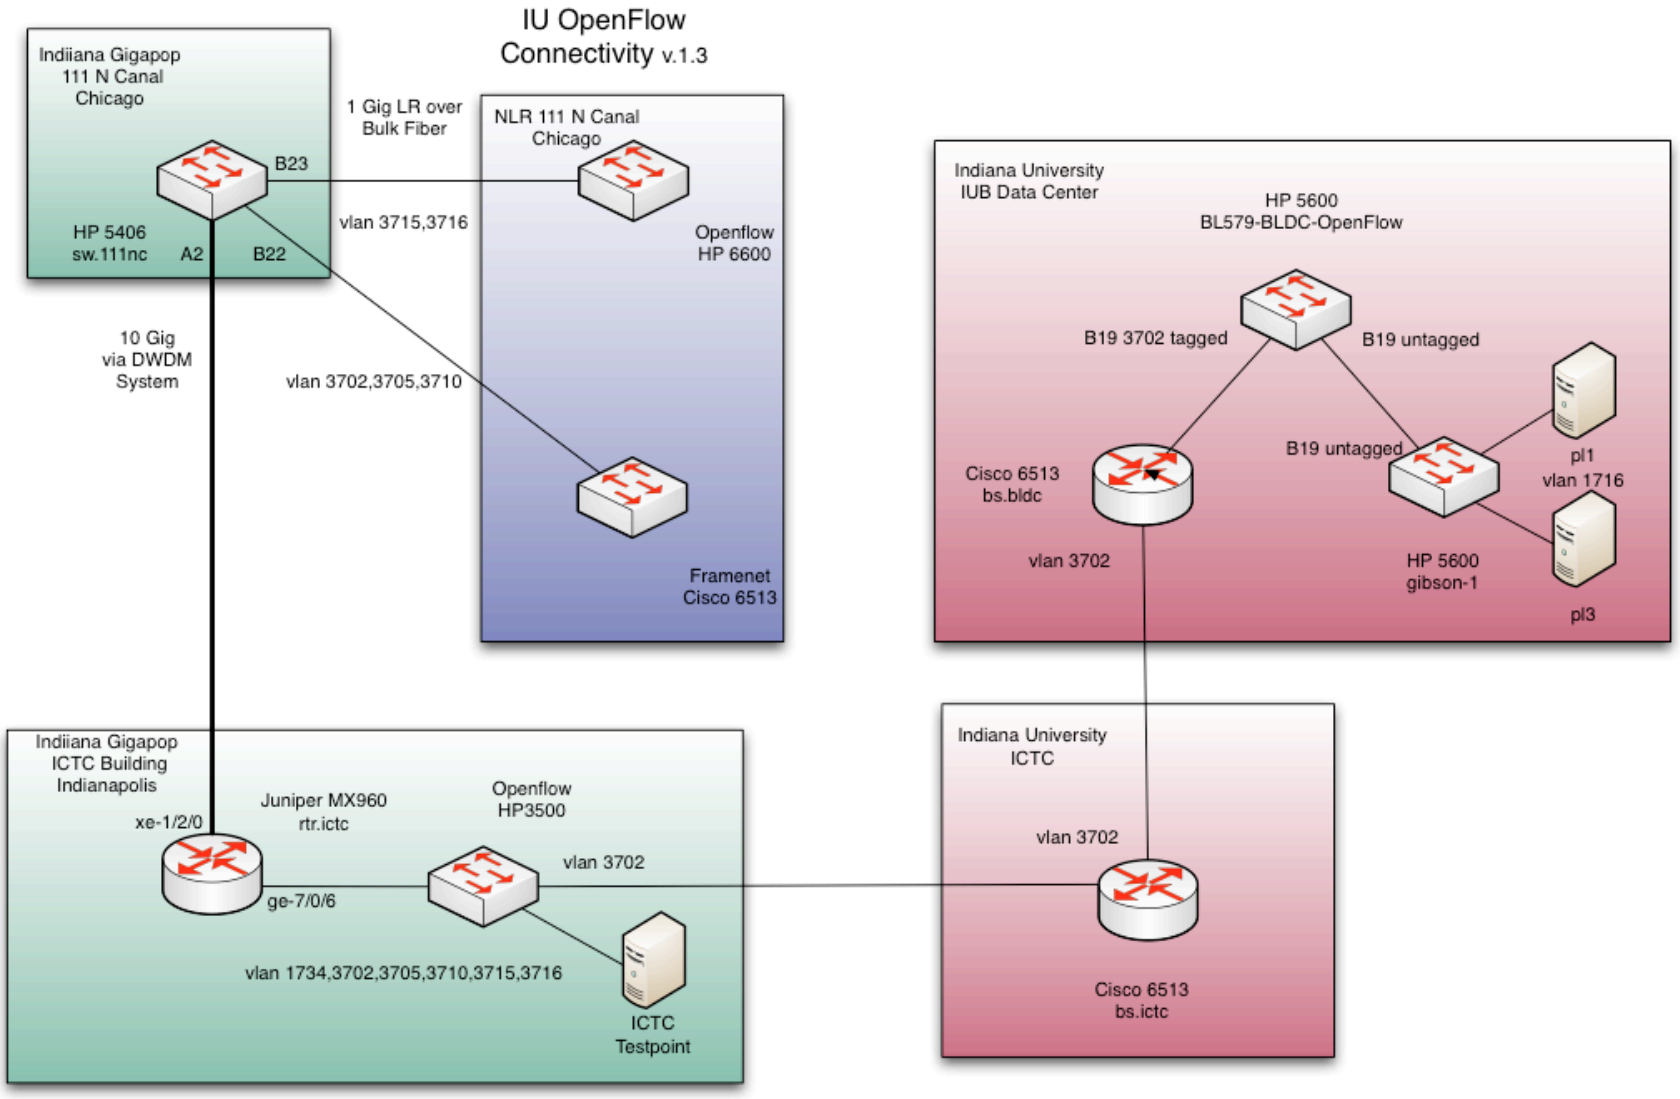
IU OpenFlow
Connectivity v.1.3
Indiiana Gigapop
111 N Canal
Chicago
B23
HP 5406
sw.111nc
A2
B22
1 Gig LR over
Bulk Fiber
NLR 111 N Canal
Chicago
Openflow
HP 6600
vlan 3715,3716
10 Gig
via DWDM
System
vlan 3702,3705,3710
Framenet
Cisco 6513
Indiana University
IUB Data Center
HP 5600
BL579-BLDC-OpenFlow
B19 3702 tagged
B19 untagged
B19 untagged
pl1
vlan 1716
Cisco 6513
bs.bldc
vlan 3702
HP 5600
gibson-1
pl3
Indiiana Gigapop
ICTC Building
Indianapolis
xe-1/2/0
Juniper MX960
rtr.ictc
Openflow
HP3500
vlan 3702
ge-7/0/6
vlan 1734,3702,3705,3710,3715,3716
ICTC
Testpoint
Indiana University
ICTC
vlan 3702
Cisco 6513
bs.ictc

3 New EU Projects: OFELIA, SPARC, CHANGE

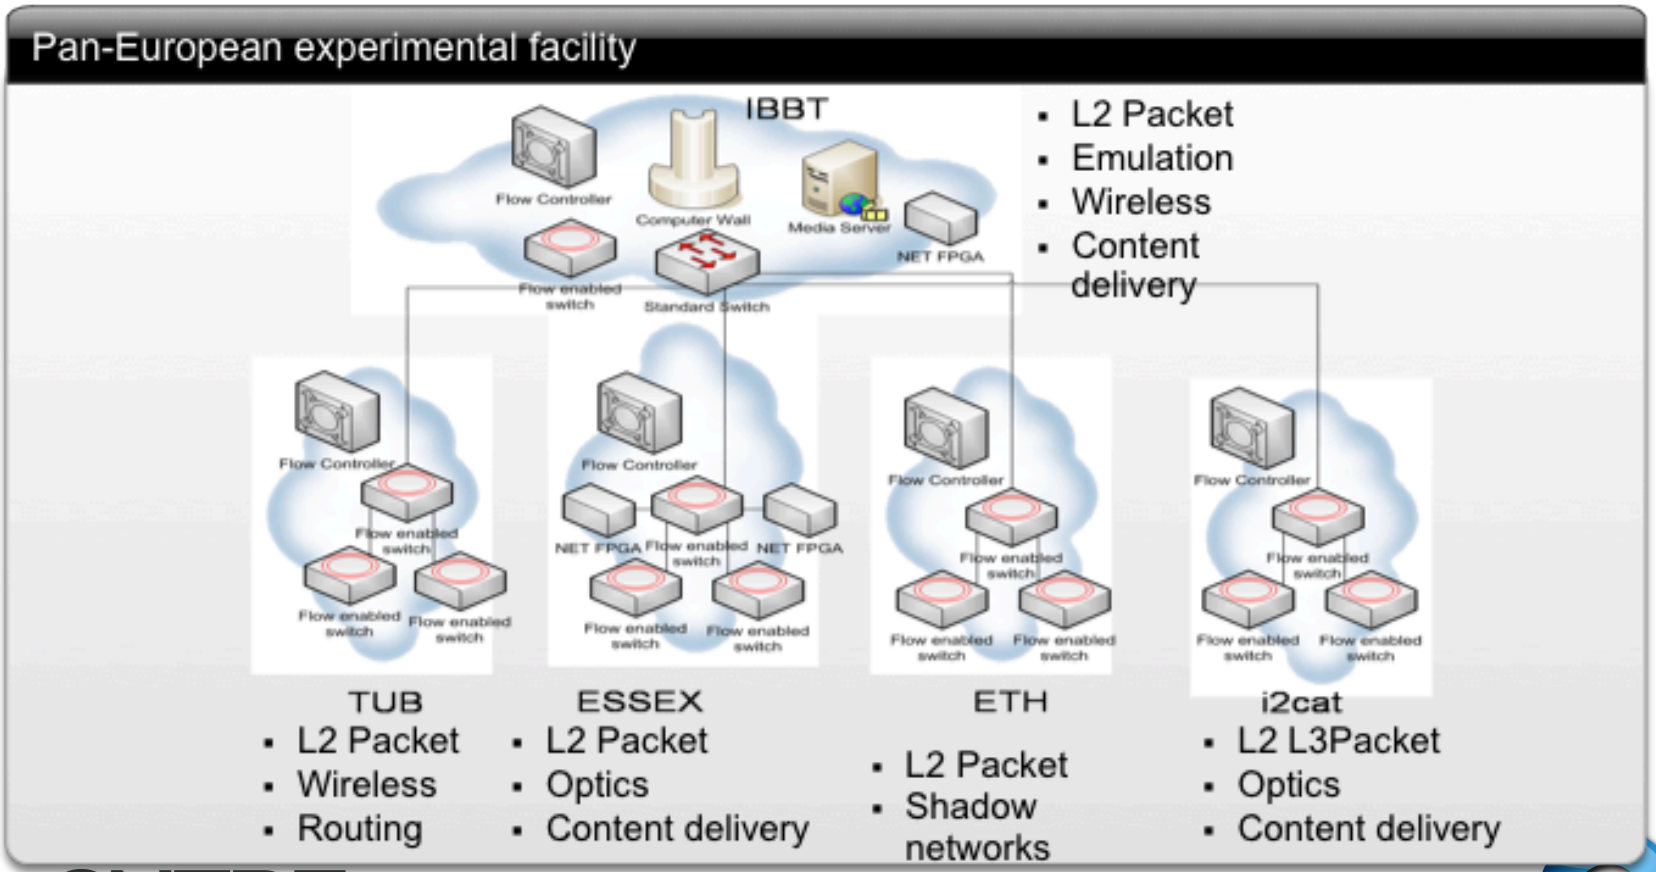

EU Project Participants

- 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)

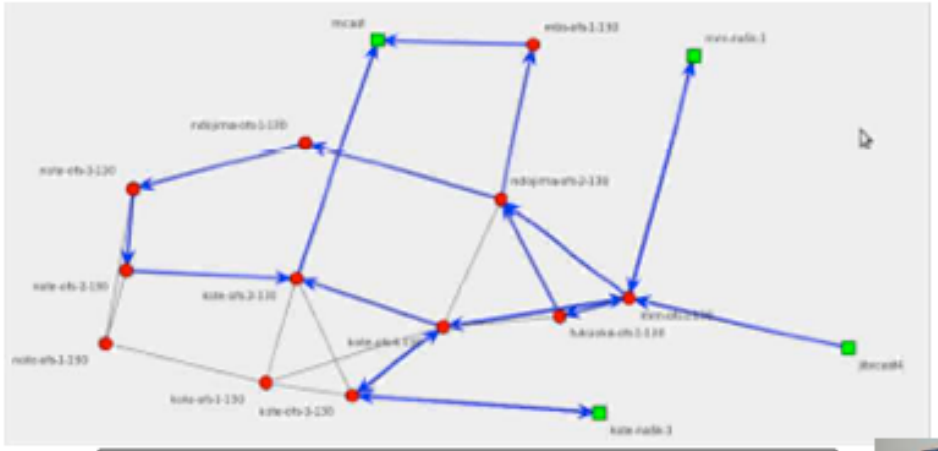

Flow and topology visualization GUI

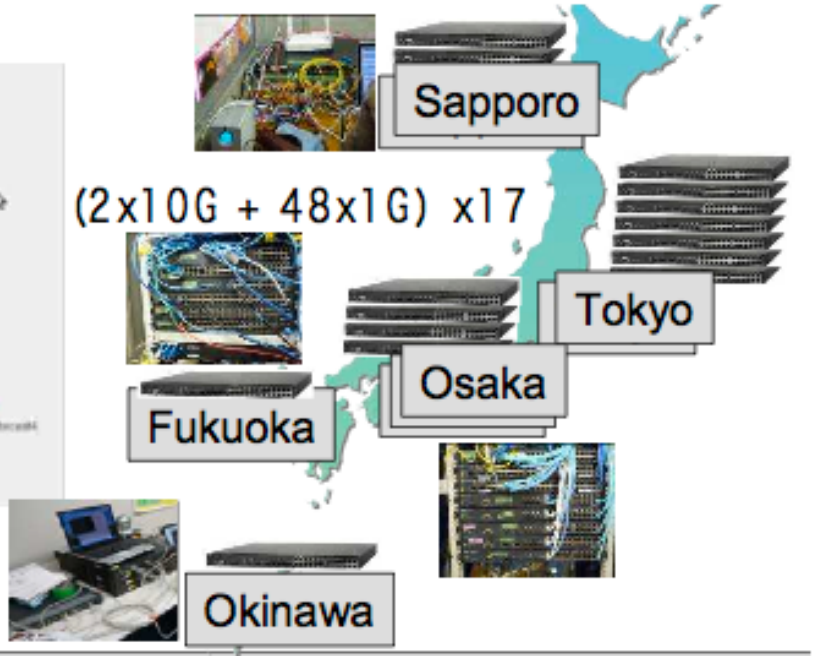

- Network virtualization and slicing
- HD video distribution in different slices
  - Baseball game
  - Snow festival

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%	58.00%
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 Concepts, Hardware and Software

Controllers

- The Network "OS"
- Open Source
  - NOX
  - Beacon/Floodlight
- Commercial
- Watch for overload on the word "Controller"

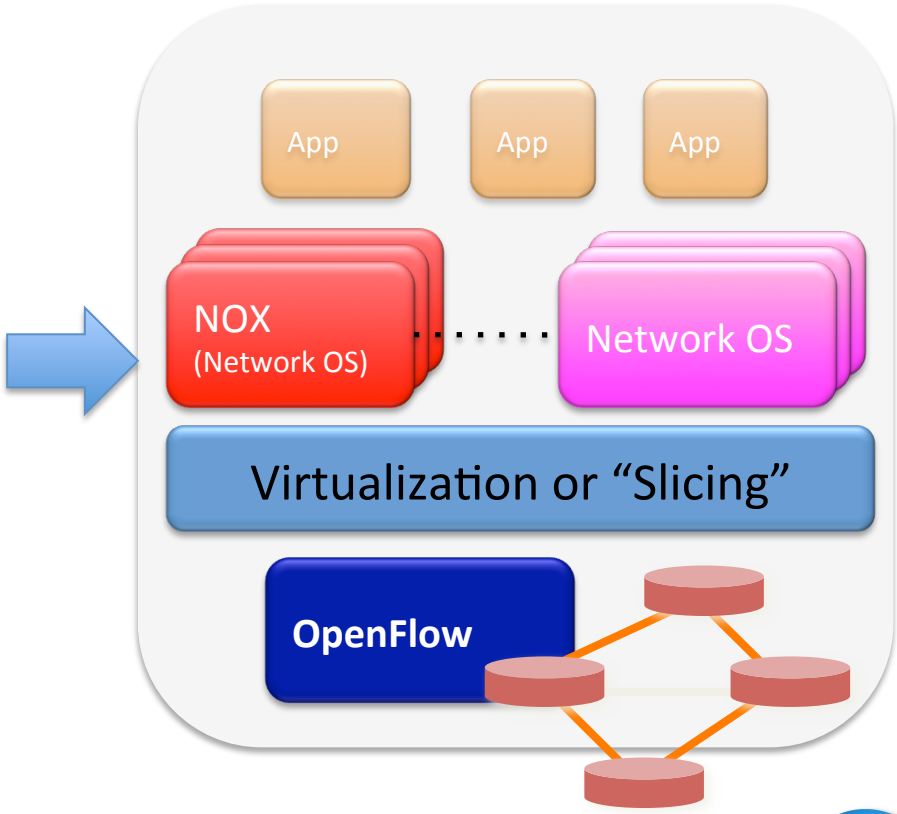

InCNTRE

Applications

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

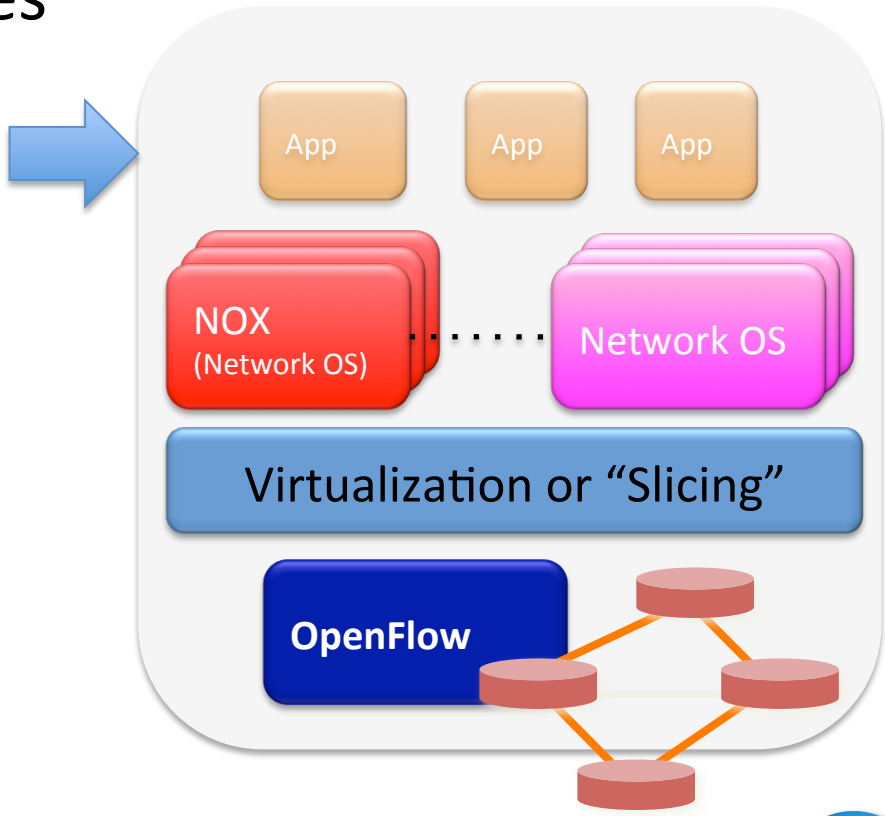

GENI Backbone Components

- Switches
- Flowvisor - Hypervisor
- Opt-In Manager – (Replacement by FOAM)
- Expedient
- User/Experimenter Applications

InCNTRE

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

InCNTRE

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],]

InCNTRE

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

InCNTRE

Experiment Workflow

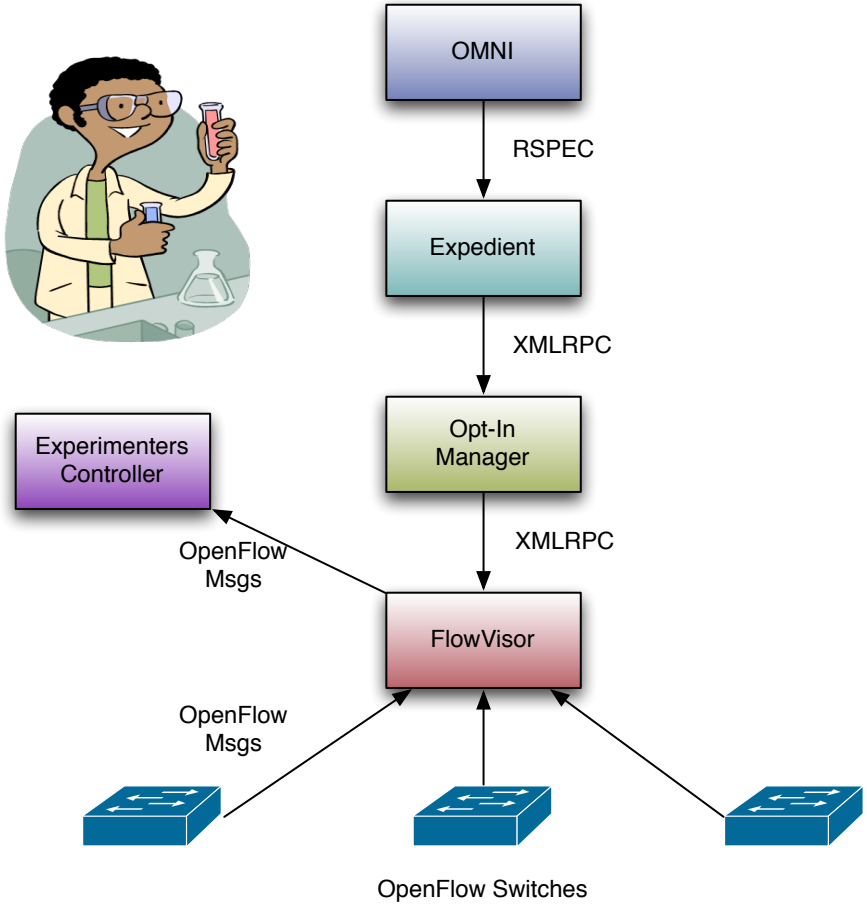

InCNTRE

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

InCNTRE

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

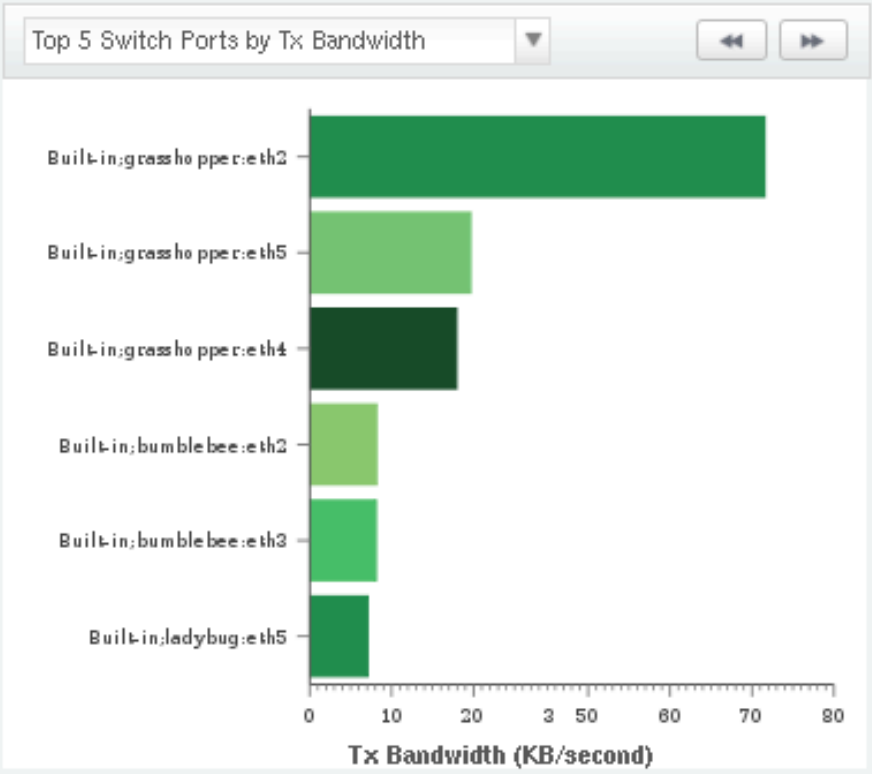

Internal Network | 6 Active Admins

Version: 0.3.0+build43

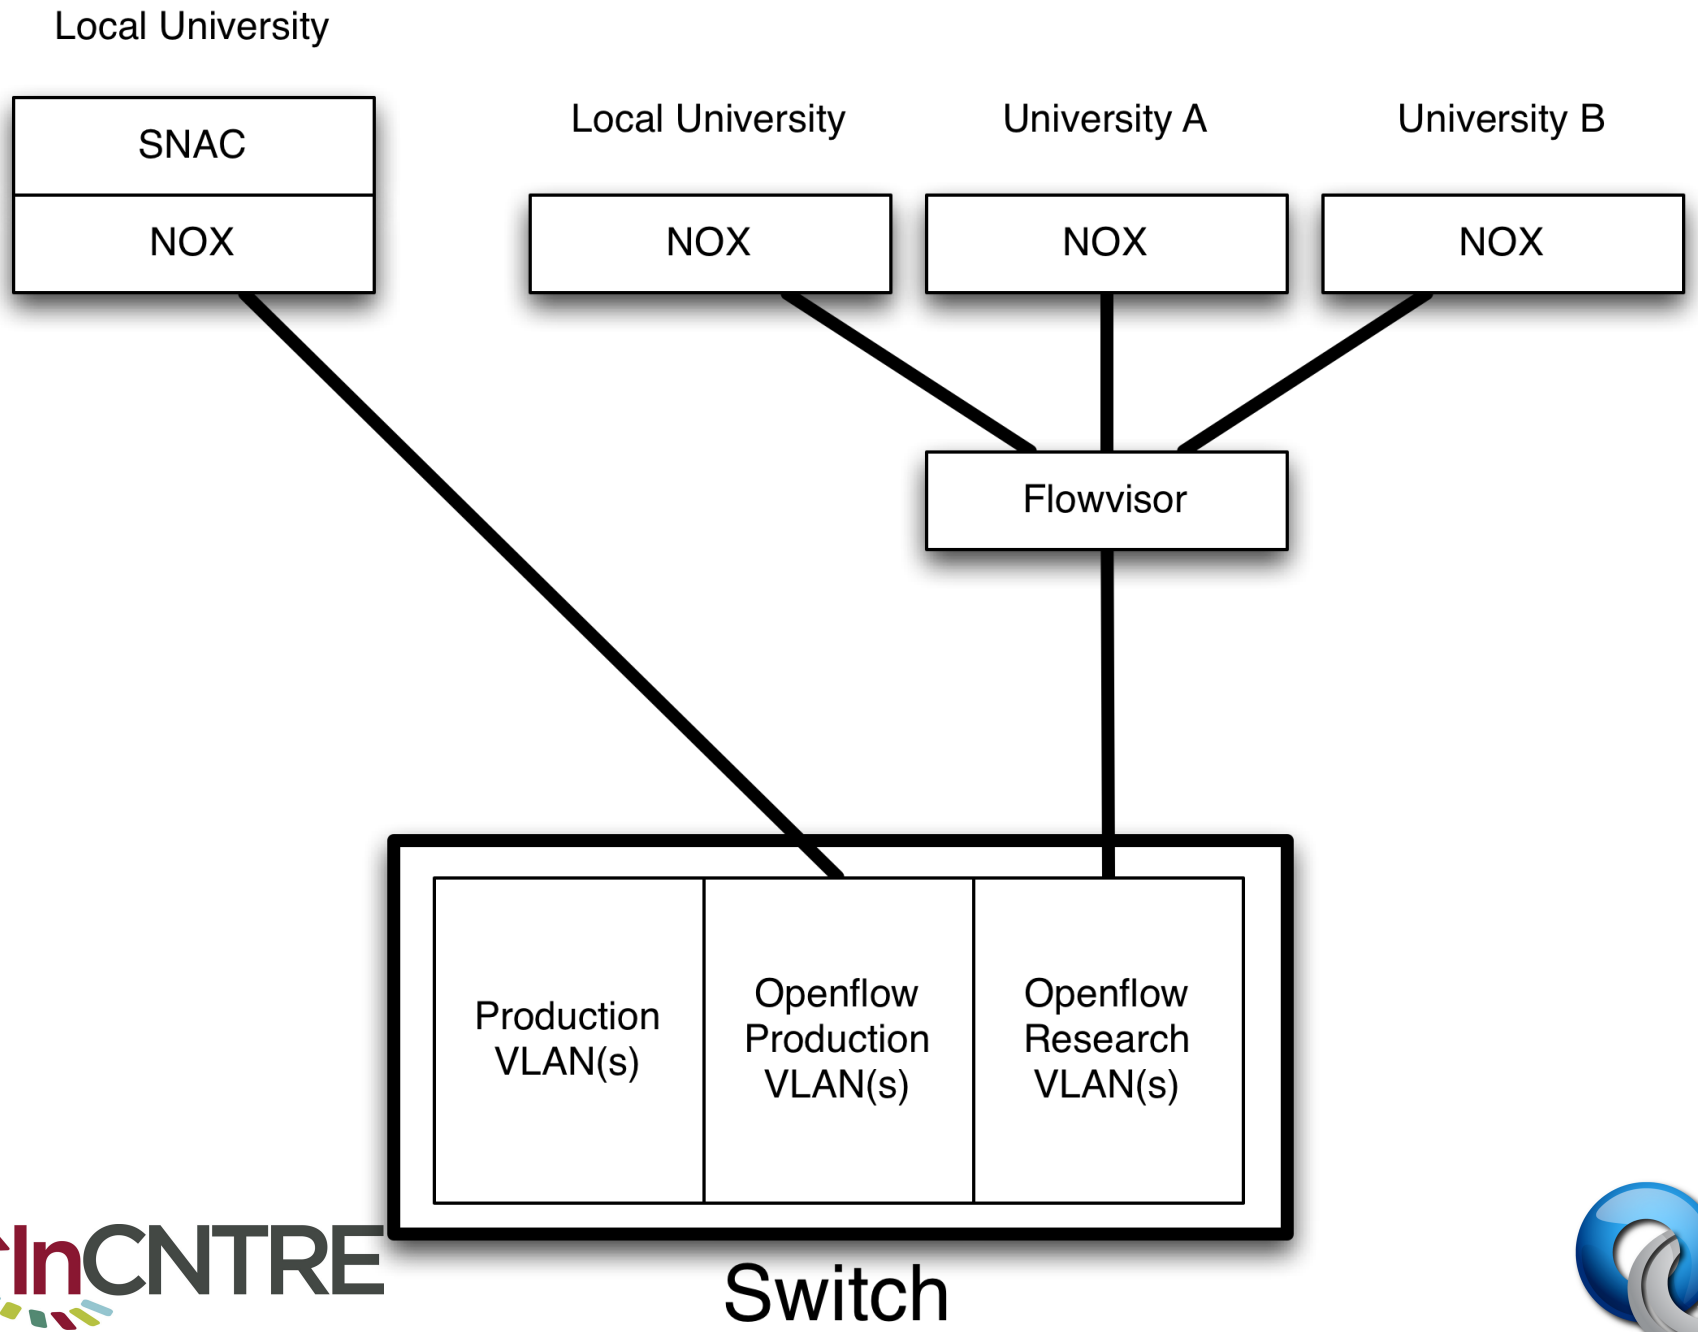
Local University
SNAC
NOX
Local University
NOX
University A
NOX
University B
NOX
Flowvisor
Production VLAN(s)
Openflow Production VLAN(s)
Openflow Research VLAN(s)
Switch

OpenFlow Hardware

Juniper

NEC

HP

Quanta/Pronto

Cisco Nexus 3k

Extreme

IBM/Blade Networks

Netgear

Netronome
EZChip
Broadcom/Marvell ref. designs
Linux
OpenWRT
WiMax
More Coming Soon..

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

InCNTRE

Switch Issues

- Hw vs Sw rules
- Optional items in OF Spec
- Control Channel resource exhaustion
- CPU exhaustion and isolation
  - Preventing OF traffic affecting production vlans
- Security
- 48bit vs 64 bit DPIDs

InCNTRE

OpenVSwitch

http://openvswitch.org

VM-aware virtual switch, run distributed over hardware;

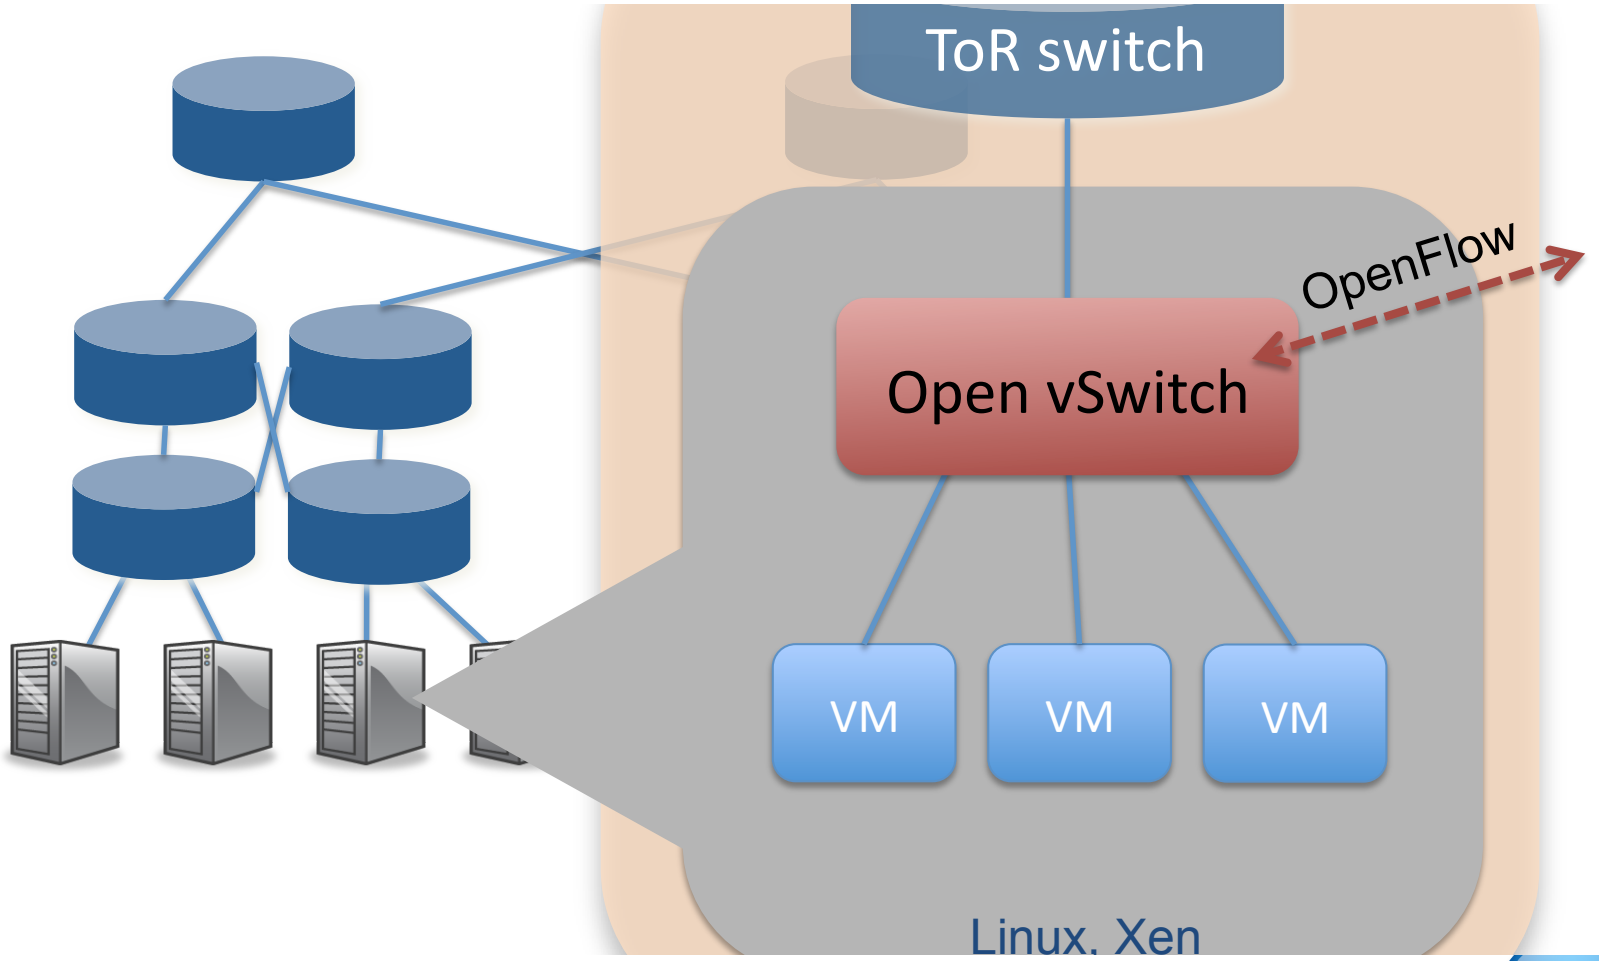

OpenFlow Spec process

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

InCNTRE

Measurement Manager

Measurement Manager

- Software built by IU for monitoring OpenFlow networks
- Ties into Flowvisor to get list of devices and topology (using LLDP)
- Acts as OF Controller to gather statistics
- Outputs Nagios, GMOC, SNAPP formats

InCNTRE

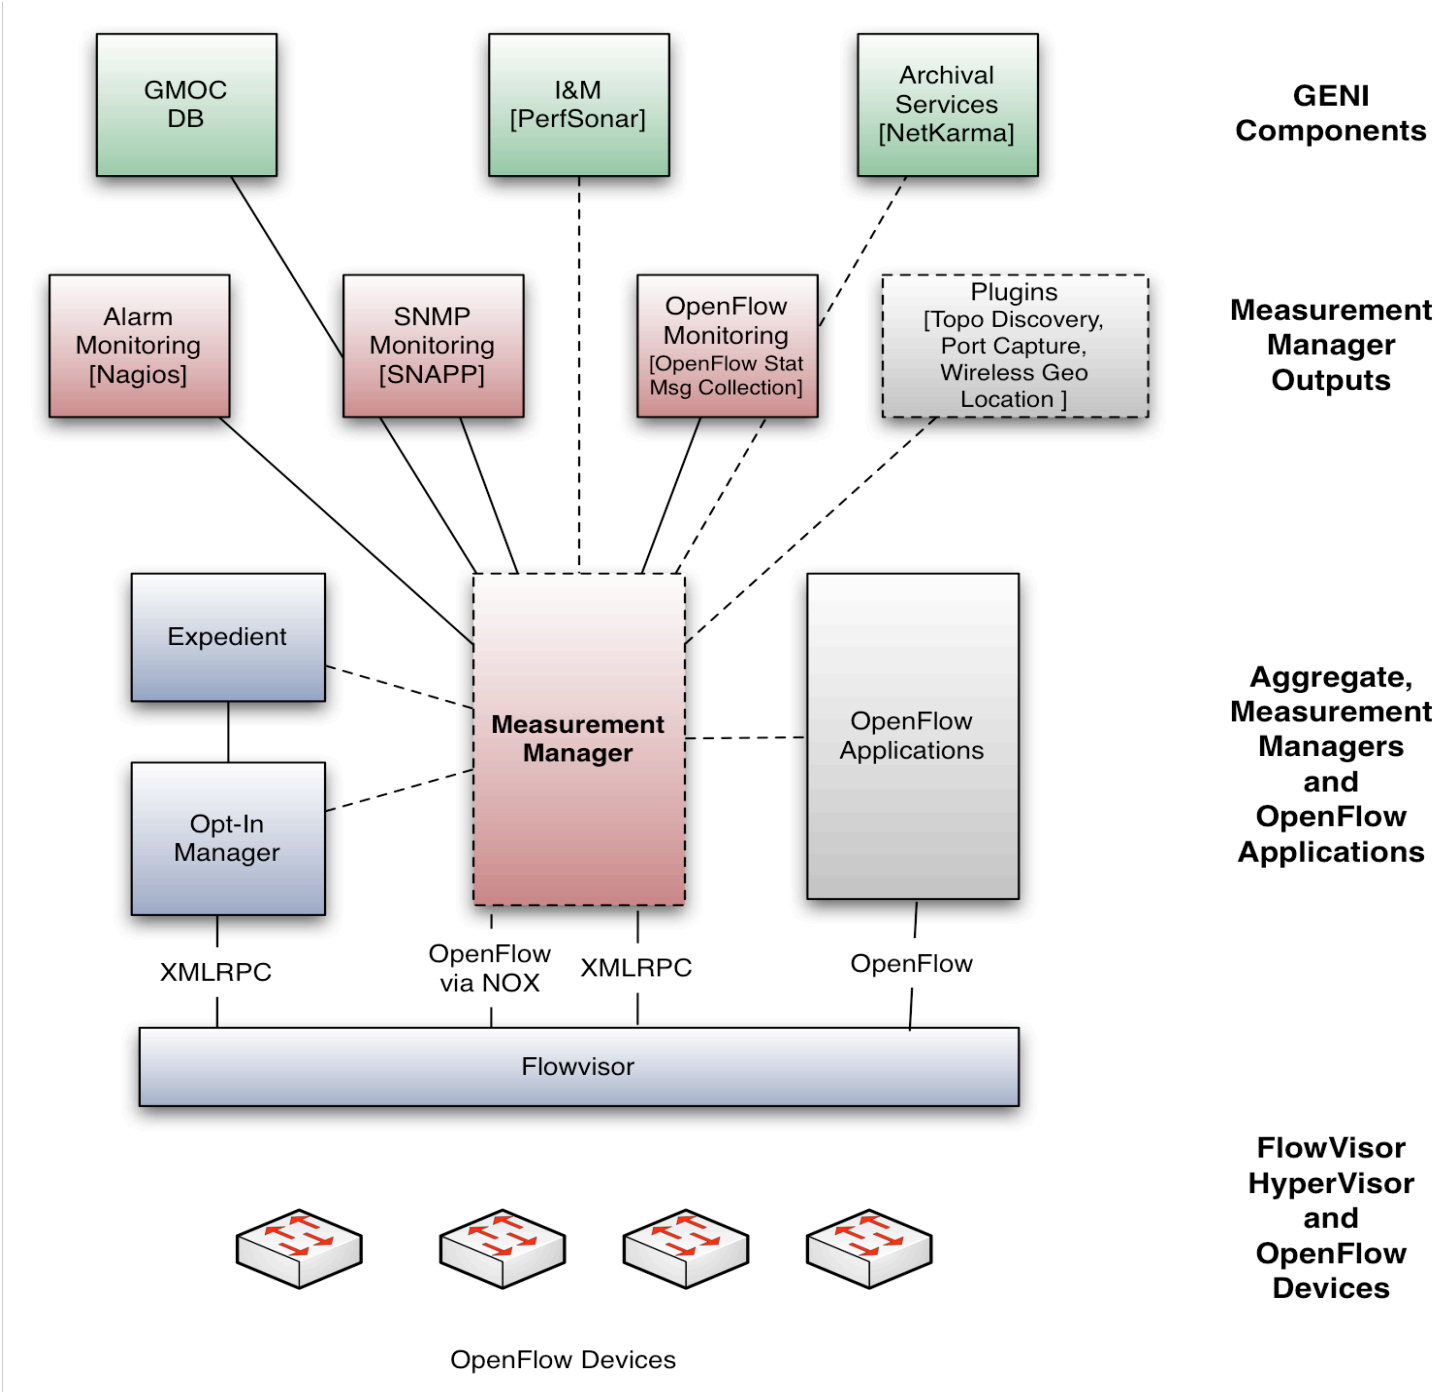
GMOC DB
I&M [PerfSonar]
Archival Services [NetKarma]
GENI Components
Alarm Monitoring [Nagios]
SNMP Monitoring [SNAPP]
OpenFlow Monitoring [OpenFlow Stat Msg Collection]
Plugins [Topo Discovery, Port Capture, Wireless Geo Location]
Measurement Manager Outputs
Expedient
Opt-In Manager
Measurement Manager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

Measurement Manager

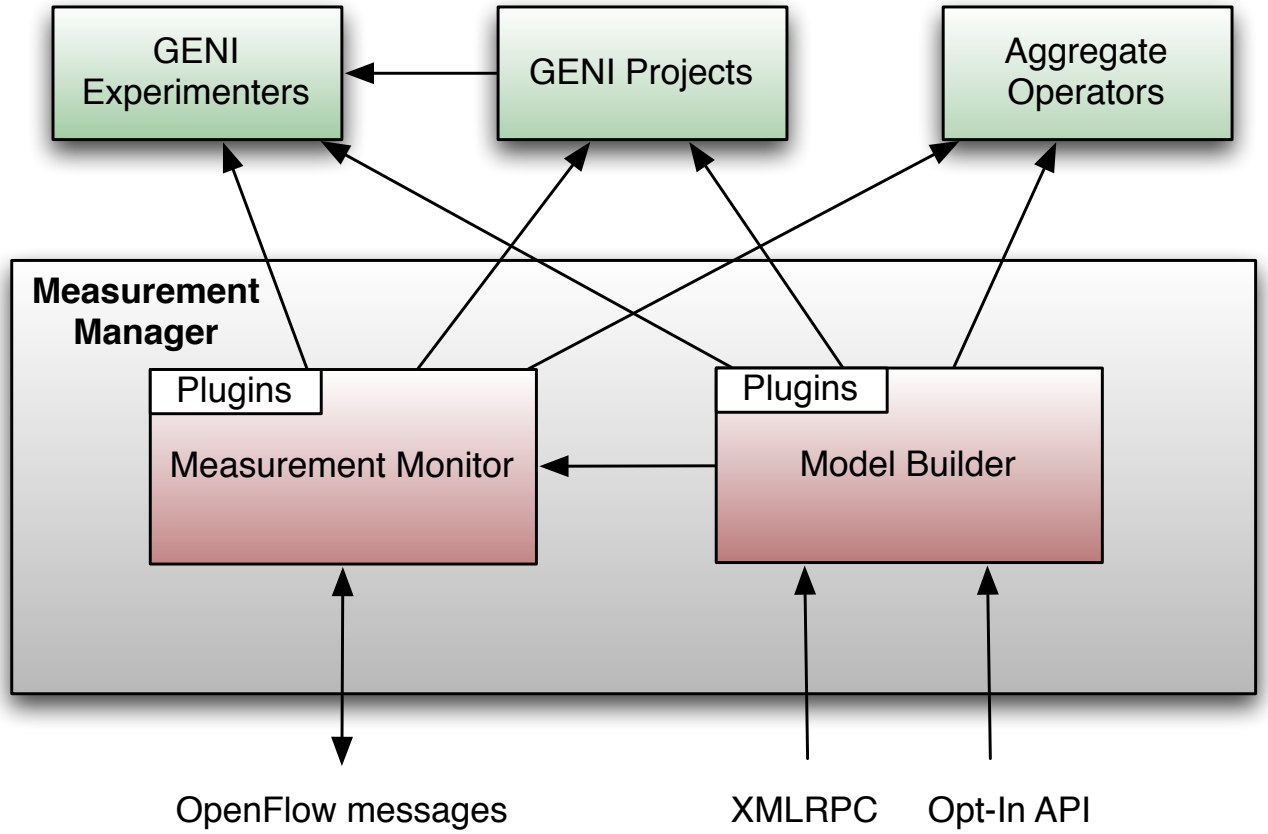

What will we do with OpenFlow ?

- 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 ?

InCNTRE

Existing Uses of Openflow

- Flow based network provisioning
  - OS^{3}E
- Load Balancing
  - FlowScale
- L3 Routing
  - RouteFlow
- Policy/ACL distribution
  - SNAC

InCNTRE

Uses of OpenFlow

- Security Applications
  - NAC – Georgia Tech
  - 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

InCNTRE

Intercontinental VM Migration

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

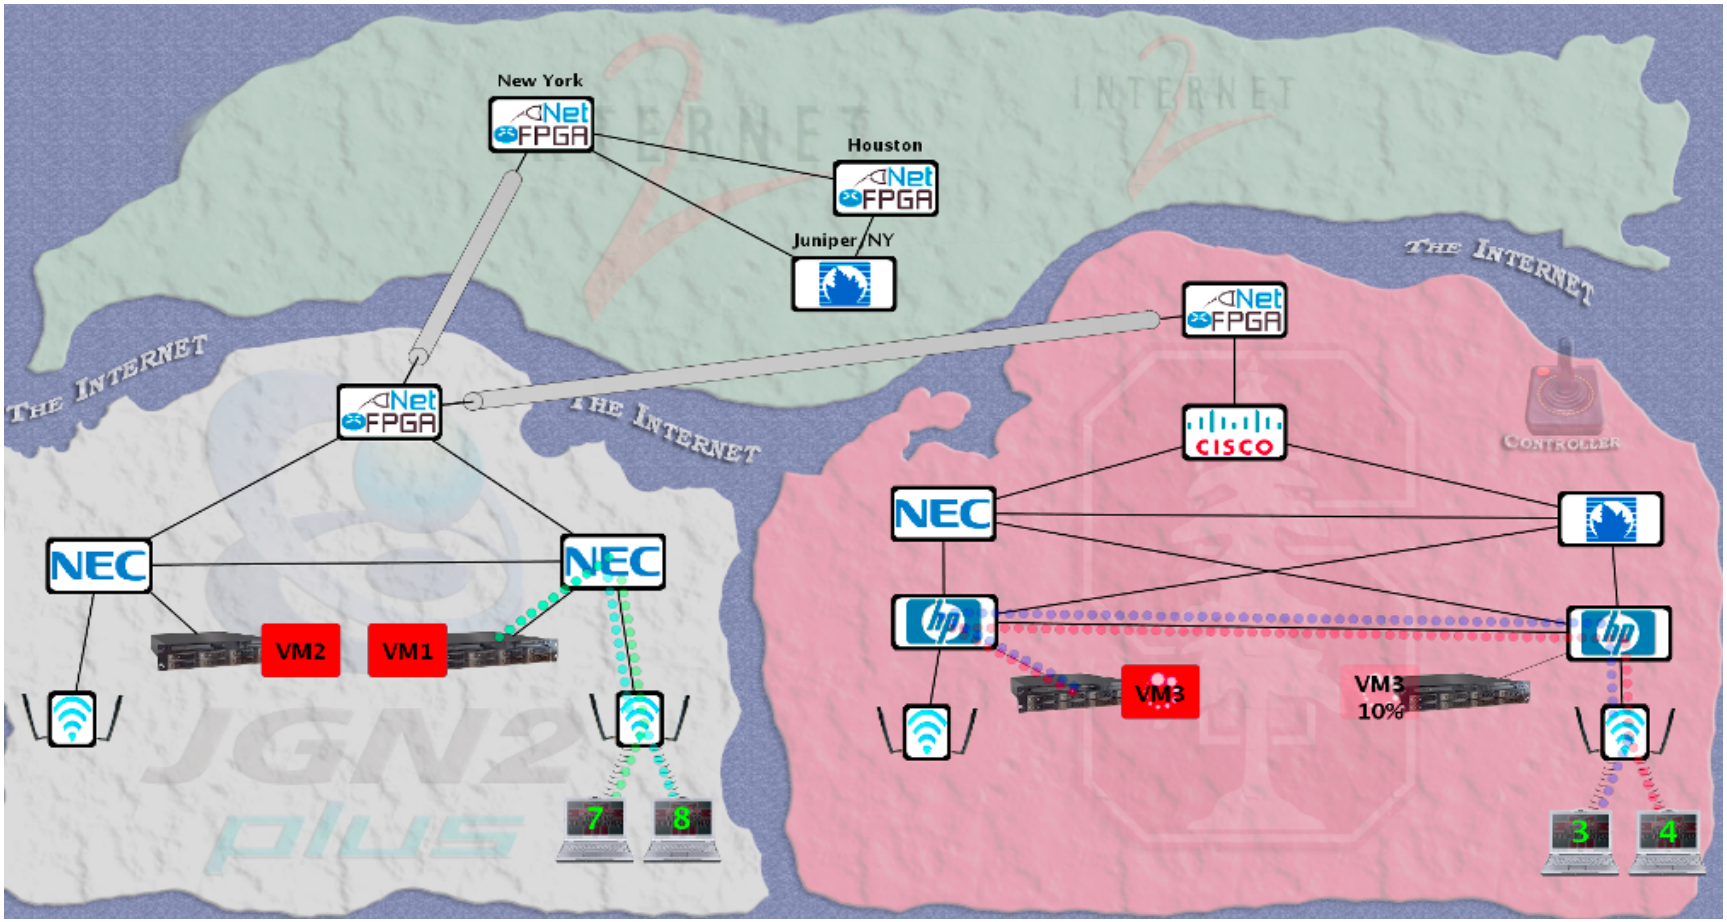

Possible Uses of Openflow (Quick Wins)

- 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.

InCNTRE

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 ?

Resources for OpenFlow

http://www.openflow.org/wk/index.php/OpenFlow_Tutorial

InCNTRE

More Tools

- dpctl
  - Can view/add flow rules on switch
  - dpctl dump-flows tcp:127.0.0.1:6633

[flowvisor@ofc-vm2 chsmall]$ dpctl dump-flows tcp:156.56.5.43:6633

stats_reply (xid=0xc13a63bc): flags=[more] type=1(flow)

cookie=0, duration_sec=0s, duration_nsec=113000000s, table_id=0, priority=65535, n_packets=0, n_bytes=0, idle_timeout=15,hard_timeout=0,tcp,in_port=30,dl_vlan=0xffff,dl_vlan_pcp=0x00,dl_src=00:1d:70:83:5b:40,dl_dst=00:0e:d6:2b:08:00,nw_src=156.56.193.180,nw_dst=67.164.35.76,nw_tos=0x00,tp_src=49211,tp_dst=53981,actions=

Wireshark

- Wireshark with OpenFlow Plugin
  - Diagnose OF problems by looking at OF messages
  - More for Cntroller debugging but can also solve issues with Flowvisor
  - Can see the “insides” of OpenFlow messages to the controller

InCNTRE

Additional Sources of Info

- Resources for OpenFlow
  - Openflow.org
  - GENI Wiki
  - Openflow section of WebGUI
  - OpenFlow Tutorial
    - http://www.openflow.org/wk/index.php/OpenFlow_Tutorial

InCNTRE

How to get involved

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

InCNTRE