

# OpenFlow Campus Trials

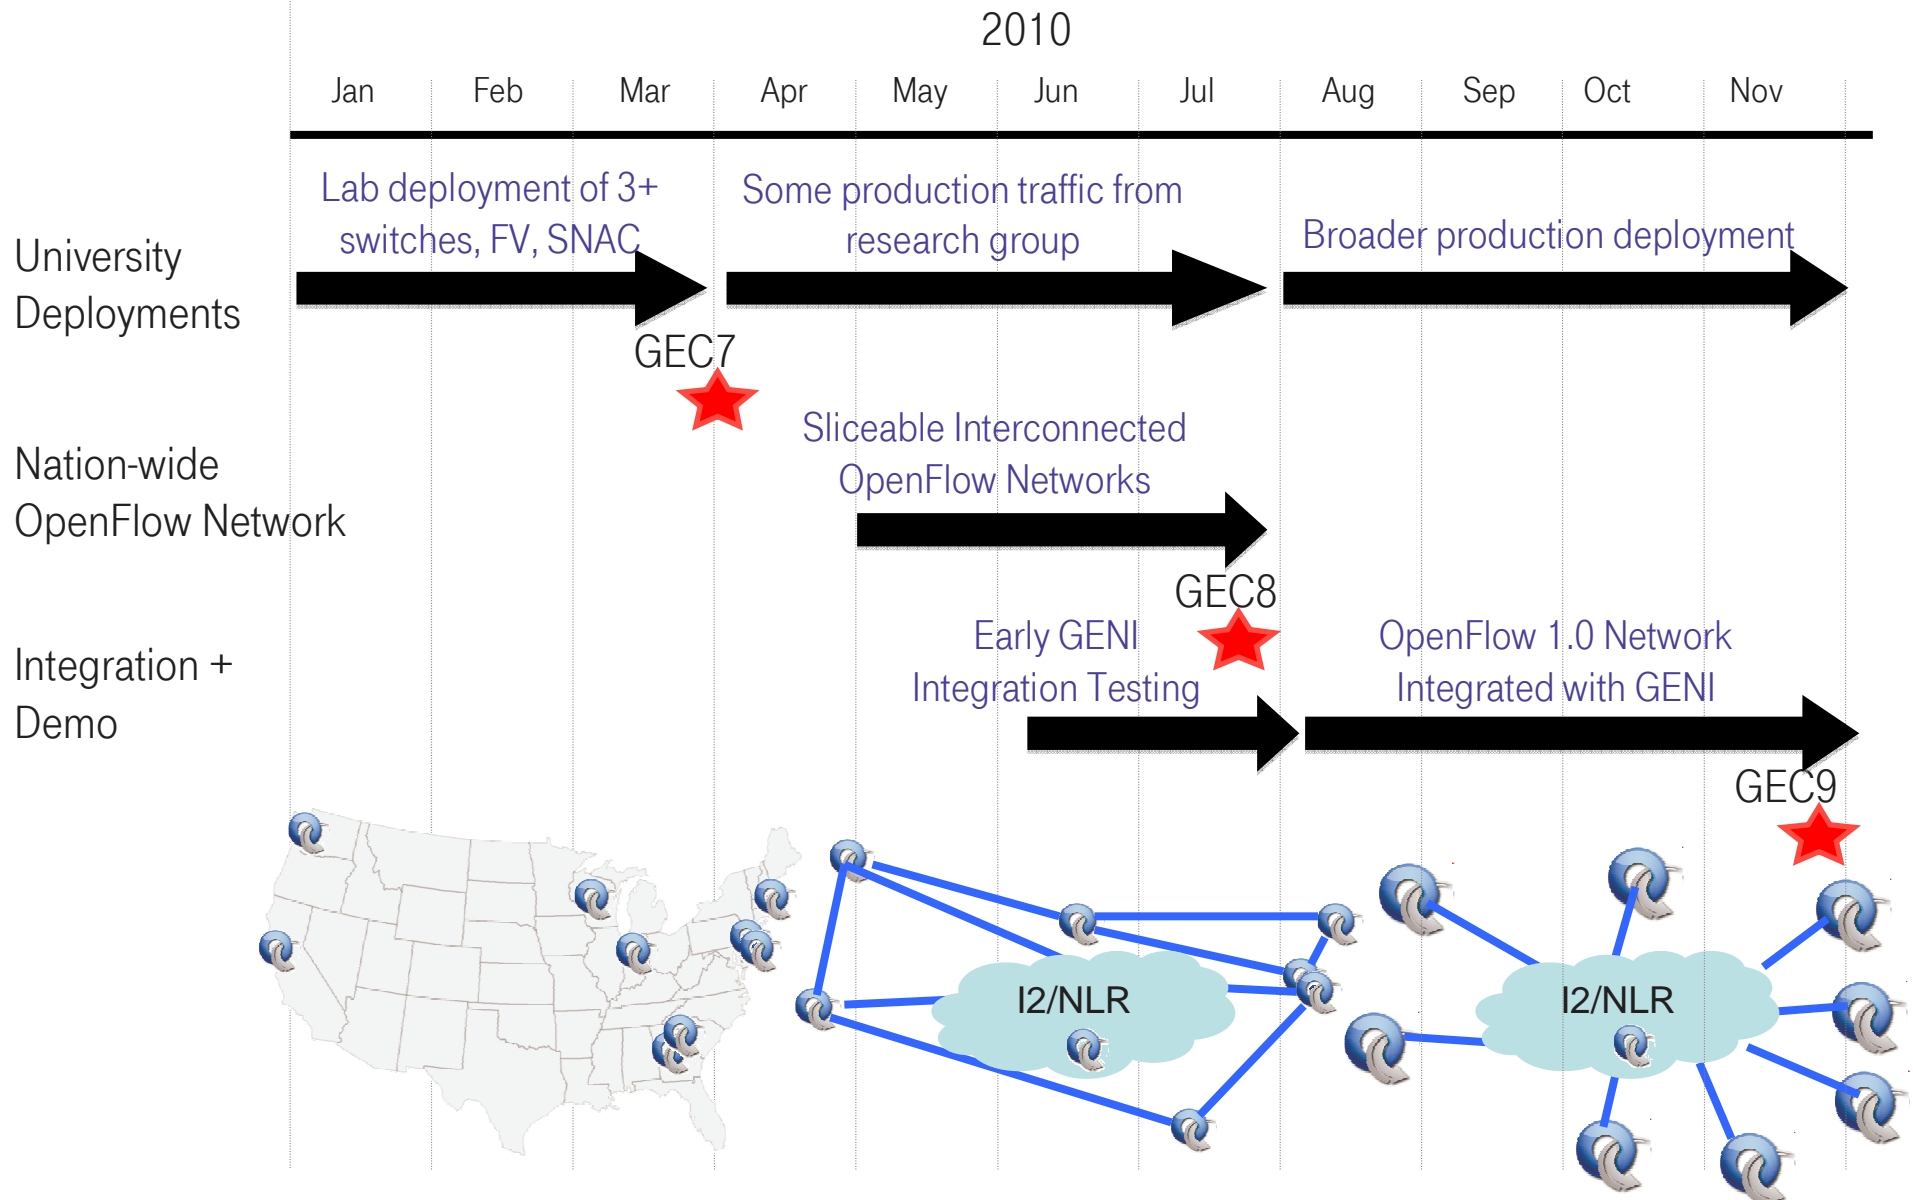
GEC7

Stanford University

# Continued progress

- **OpenFlow 1.0**
  - Spec released in Dec 2009
  - Reference implementations and early vendor implementations available
- **Increasing vendor interest**
  - HP support
  - NEC moving aggressively
  - Toroki
  - Quanta + Stanford software
  - Extreme networks (?)
  - More vendors in the pipeline
- **Increasing provider interest and engagement**
  - Google, Amazon, Yahoo, Microsoft, ...
  - DT, Verizon, Level3,
- **EU**
  - Funded three large projects
- **China**
  - CERNET, CSTNET, and others interested

# OpenFlow GENI roadmap



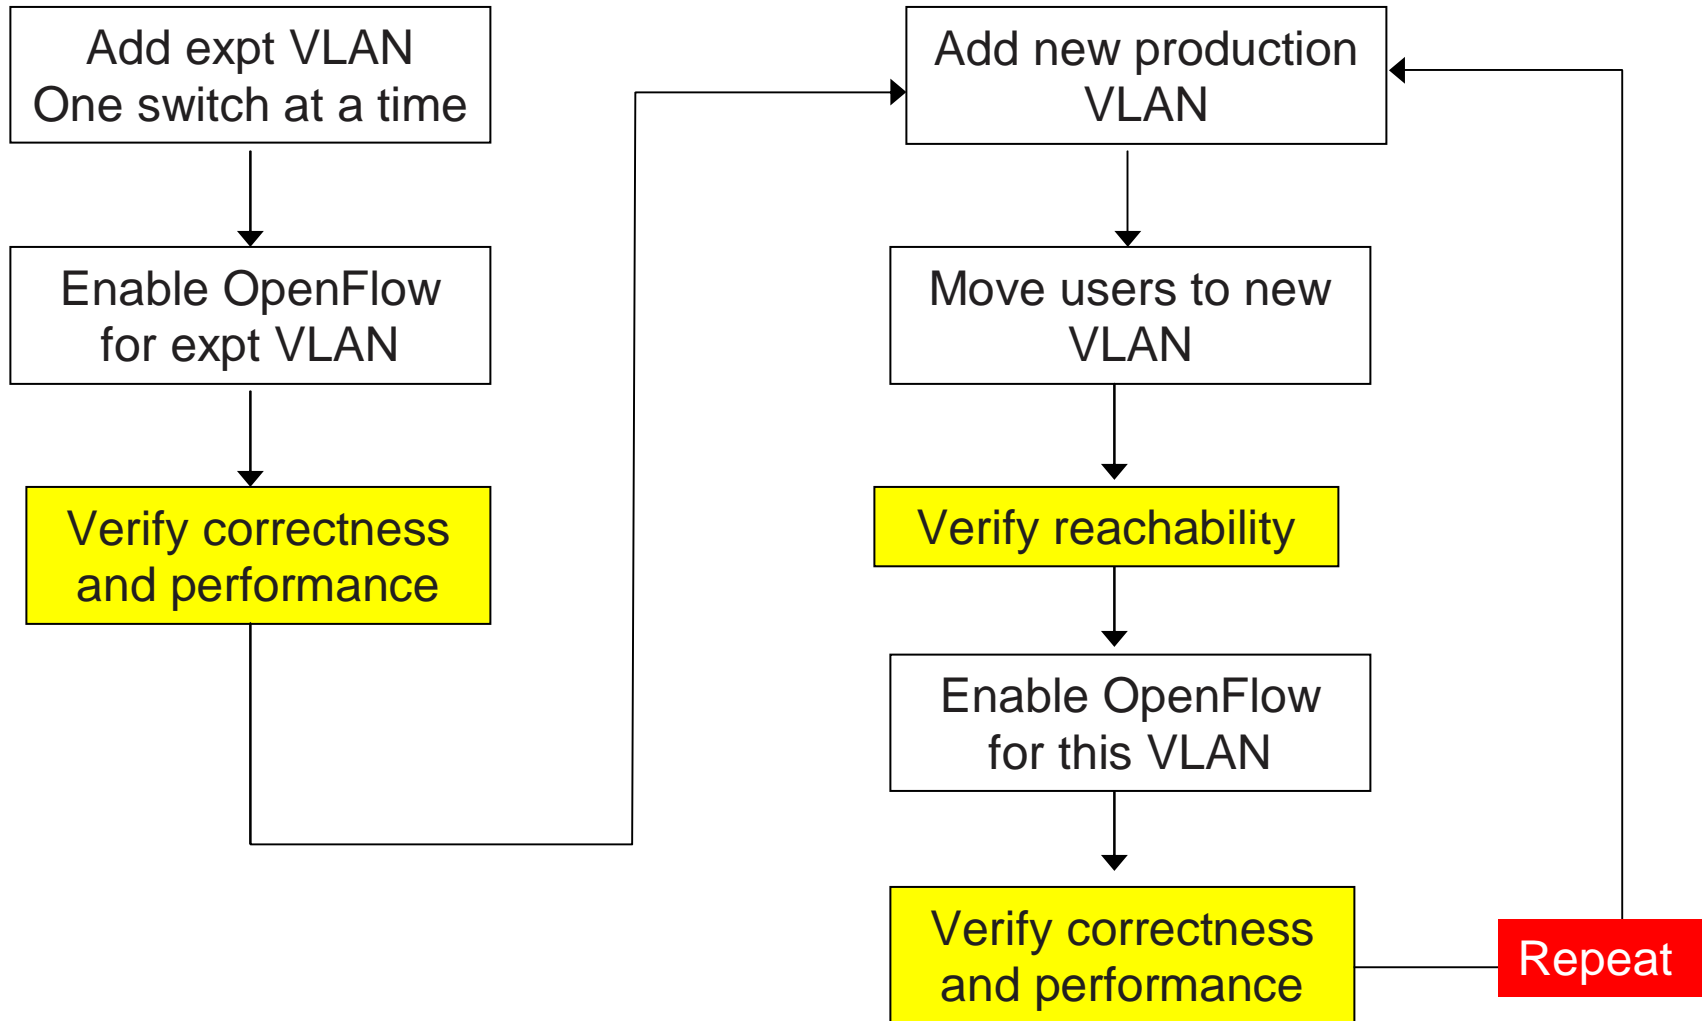
# GEC8: Nation-wide OpenFlow network

- 6+ OpenFlow switches, operated by campuses
- OpenFlow VLAN A:
  - Handles all research group traffic
  - Controlled by FlowVisor + SNAC
- OpenFlow VLAN B is sliced by FV into 3 or more slices:
  - For research and experimentation
- Early integration testing with GENI control plane
- Demo: Show expt spanning 2 or more campuses at GEC8 meeting, along with FV GUI for local aggregate.

# Key challenges

- Scale OpenFlow deployment
  - Add more switches and WiFi APs
  - Add slicing for production & experimentation
- Achieve network stability with experimentation
  - Keep users and experimenters happy
- Connect campus OpenFlow network to I2/NLR OpenFlow backbone
- Start integration with GENI control plane
- GEC8 not that far off and during summer

# Solution: Staged deployment



# Resources

- Support system
  - People, online resources, and more
- Stanford deployment experience
  - OpenFlow becoming production ready, but expect issues and plan well
- Goals within our reach if we plan well
  - Specific deployment plan for each campus
  - Customize support plan accordingly

# Support System



# Support team

## Stanford

Masa



Srini



Paul




Johan



## GPO/BBN

Josh



Heidi



# Support system

- **Bi-weekly calls:**
  - Help debug deployment issues
  - Help prepare a customized deployment / demo plan
- **Website:**  
[www.openflowswitch.org/foswiki/bin/view/OpenFlow/Deployment/](http://www.openflowswitch.org/foswiki/bin/view/OpenFlow/Deployment/)
- **Mailing lists:**  
openflow-discuss, openflow-spec, openflow-dev,  
nox-dev, egeni-trials, deployment-help
- **Bug tracking system:**
  - <http://www.openflowswitch.org/bugs/snac>,  
[/bugs/toroki](http://www.openflowswitch.org/bugs/toroki), [/bugs/flowvisor](http://www.openflowswitch.org/bugs/flowvisor), [/bugs/openflow](http://www.openflowswitch.org/bugs/openflow)
  - For bugs with HP, please mail [jean.tourrilhes@hp.com](mailto:jean.tourrilhes@hp.com)
  - For bugs with NEC, please mail [ofs-support@spf.jp.nec.com](mailto:ofs-support@spf.jp.nec.com)

# Support system (contd.)

- BBN/GPO information wiki:

- <http://groups.geni.net/geni/wiki/OFCLEM>,  
[wiki/OFGT](http://groups.geni.net/geni/wiki/OFGT),  
[wiki/OFRG](http://groups.geni.net/geni/wiki/OFRG),  
[wiki/OFNOX](http://groups.geni.net/geni/wiki/OFNOX),  
[wiki/CampusConnectivity](http://groups.geni.net/geni/wiki/CampusConnectivity)
- [wiki/OFIU](http://groups.geni.net/geni/wiki/OFIU),  
[wiki/OFUWA](http://groups.geni.net/geni/wiki/OFUWA),  
[wiki/OFBBN](http://groups.geni.net/geni/wiki/OFBBN),
- [wiki/OFPR](http://groups.geni.net/geni/wiki/OFPR),  
[wiki/OFUWI](http://groups.geni.net/geni/wiki/OFUWI),  
[wiki/EnterpriseGeni](http://groups.geni.net/geni/wiki/EnterpriseGeni),

- BBN/GPO mailing lists:

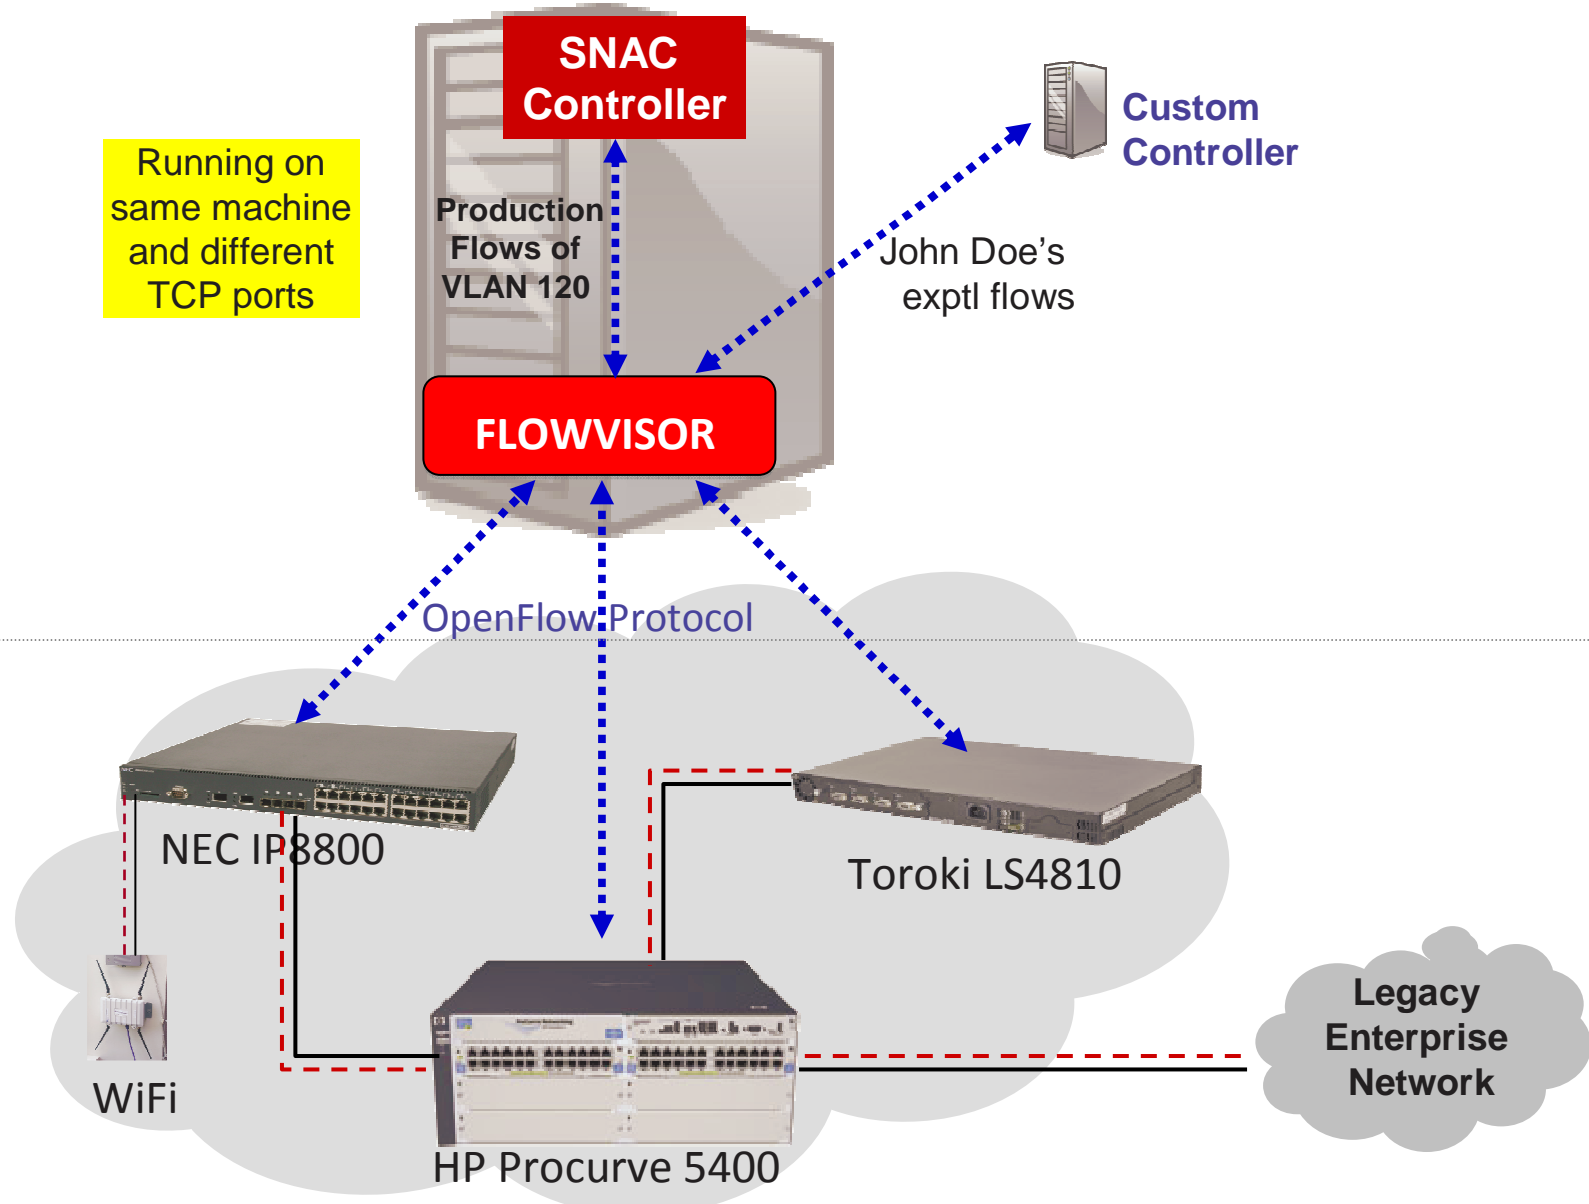
- [openflow@geni.net](mailto:openflow@geni.net), [backbone-integration@geni.net](mailto:backbone-integration@geni.net), [geni-node-ops@geni.net](mailto:geni-node-ops@geni.net), [response-team@geni.net](mailto:response-team@geni.net)

- One-on-one support from Josh Smift for

- Wide-area network GENI connection
- GENI API and integration

# Status of Components

# Different components in the Network



# Availability of OpenFlow components

Modules	Currently Available Version	Version used for GEC8	Version used for GEC9	When GEC9 demo version becomes available?
OpenFlow Switch	0.8.9 (1.0 for s/w ref design)	1.0(Stanford + ?), 0.8.9 (others)	1.0*	•HP & NEC: April 2010 (Alpha version available for HP)
NOX	0.6	0.6	1.0	Aug 2010
SNAC	0.4	0.4	1.0	TBD
FlowVisor	0.4	0.5	1.0	Aug 2010
FlowVisor console	-	0.5	1.0	
Aggregate Manager	SFA_0.9.5	0.5	1.0	
ENVI	Available online in the production deployment page			
LAVI				
Monitoring & Debugging Tools				

(\*) Ensures compatibility across campuses

# Summary of resolved issues

- Frequent stats request causing HP CPU spikes
  - Well understood issue that we pay attention to
  - Workaround: Reduce frequency of stats request or block it at FV
- HP switch dropping LLDP packets:
  - HP dropping LLDP packets with multicast source address
  - Resolved by fixing *discovery* module of SNAC
- Switches not allowing hot swap of ports
  - The controller ignores port status change during runtime
  - Resolved by fixing *discovery* module of SNAC
- Link timeout incorrect causing frequent churn
  - Resolved by increasing link timeout in SNAC module

# Summary of resolved issues (contd.)

- Packet\_out action of TABLE did not work for NEC switch
  - Caused first packet to be dropped
  - Resolved by firmware fix from NEC
- HP switch issues:
  - Poor browsing performance
  - Resolved by firmware fix from HP
- Wireless DHCP
  - Invalid packet forwarding
  - Resolved by erasing stale bindings in *authenticator* of SNAC
- Duplicate packets sent to OFPP\_LOCAL
  - For WiFi APs having of0 port, invalid action is sent by *FlowVisor*
  - Resolved by performing additional check in *FlowVisor*



# Summary of existing issues

Most issues are non-blockers in our deployment

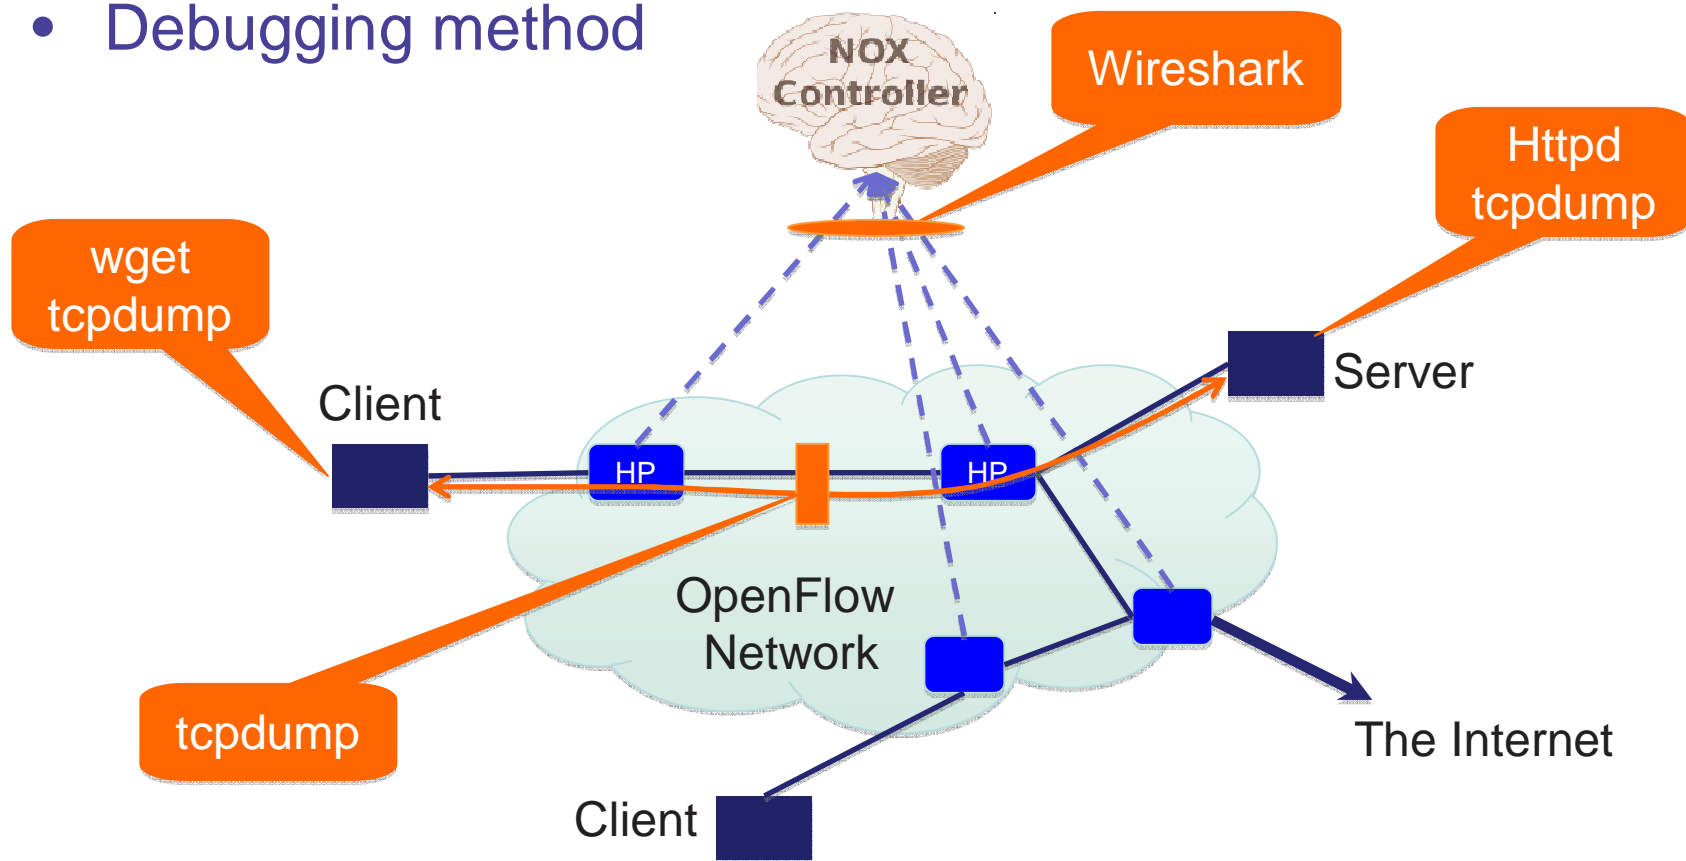
- Toroki switch issues:
  - Open issues:
    - MAC rewriting not working
    - Instability during power cycle
    - Flows not expiring when controller is stopped while traffic is running
  - Status: Vendor is working on a fix
- Invalid state storage in SNAC
  - Removing port during run time of SNAC is not supported
  - Status: Need to investigate performance impact
- Invalid bindings in SNAC following topology change
  - Status: Being discussed on nox-dev list

# Summary of existing issues (contd.)

- No spanning tree support in controller
  - Caused an outage in CIS/CISX, when operator installed a loop
  - Status: Developing a NOX/SNAC module
- No link bundling (LACP) support in OpenFlow switch
  - Status: Vendors are looking at fix
  - Workaround: Use dedicated OpenFlow links
- No redundancy or failover with ver0.8.9
- No IPv6, Multicast, or 802.1X support in controller

# Resolved #1: HP wget performance issue

- Symptom
  - Web browsing performance was poor if HP switch is on the path
- Debugging method



# Resolved #1: HP wget performance issue

## DATA PATH INDICATED SYN RETRANSMITS:

1266568067.414724 IP 172.24.74.121.44544 > 171.67.216.18.80: S 288018868:288018868(0) win 5840

**1266568070.412083 IP 172.24.74.121.44544 > 171.67.216.18.80: S 288018868:288018868(0) win 5840**

1266568070.412554 IP 171.67.216.18.80 > 172.24.74.121.44544: S 2119182178:2119182178(0) ack 288018869 v

We recommend using the *wireshark* dissector for debugging purposes

No. ↓	Time	Source	Destination	Protocol	Info
4	0.348571	171.67.75.2	128.208.4.156	TCP	8833 > 38309 [ACK] Seq=13 Ack=429 Win=1002 Len=0 TSV=1180389900 TSER=97037950
5	0.603936	Supermic_b0:54:ab	Broadcast	OFF+ARP	Packet Out (CSM) (108B) => Who has 0.0.0.0? Tell 49.48.46.55
6	0.604148	128.208.4.156	171.67.75.2	OFFP	Error (SM) (12B)
7	0.623652	171.67.75.2	128.208.4.156	TCP	8833 > 38309 [ACK] Seq=121 Ack=441 Win=1002 Len=0 TSV=1180389969 TSER=97038019
8	0.632078	10.79.1.41	10.79.1.29	OFF+TCP	Packet In (AM) (BufID=397822) (92B) => 42620 > http [SYN] Seq=0 Win=5840 Len=0 MSS=1460 TSV=51494144 TSER=0
9	0.651661	171.67.75.2	128.208.4.156	TCP	8833 > 38309 [ACK] Seq=121 Ack=533 Win=1002 Len=0 TSV=1180389976 TSER=97038026
10	0.666384	IntelCor_b0:04:b1	Broadcast	OFF+ARP	Packet Out (CSM) (108B) => Who has 0.0.0.0? Tell 49.48.46.55 (duplicate use of 49.48.46.55 detected!)
11	0.666541	128.208.4.156	171.67.75.2	OFFP	Error (SM) (12B)
12	0.724099	171.67.75.2	128.208.4.156	TCP	8833 > 38309 [ACK] Seq=229 Ack=545 Win=1002 Len=0 TSV=1180389994 TSER=97038035
13	0.726943	Supermic_b1:35:cf	Broadcast	OFF+ARP	Packet Out (CSM) (108B) => Who has 0.0.0.0? Tell 49.48.46.55 (duplicate use of 49.48.46.55 detected!)

▶ Frame 8 (158 bytes on wire, 158 bytes captured)

▶ Ethernet II, Src: Intel\_4c:ce:68 (00:0e:0c:4c:ce:68), Dst: Cisco\_15:44:80 (00:18:74:15:44:80)

▶ Internet Protocol, Src: 128.208.4.156 (128.208.4.156), Dst: 171.67.75.2 (171.67.75.2)

▶ Transmission Control Protocol, Src Port: 38309 (38309), Dst Port: 8833 (8833), Seq: 441, Len: 92

▼ OpenFlow Protocol

▶ Header

▼ Packet In

Buffer ID: 397822  
Frame Total Length: 74  
Frame Recv Port: 63  
Reason Sent: No matching flow (0)

▼ Frame Data: 003048B054AB080027A1081B08004500003C2D3840004006...

▶ Ethernet II, Src: CadmusCo\_a1:08:1b (08:00:27:a1:08:1b), Dst: Supermic\_b0:54:ab (00:30:48:b0:54:ab)

▶ Internet Protocol, Src: 10.79.1.41 (10.79.1.41), Dst: 10.79.1.29 (10.79.1.29)

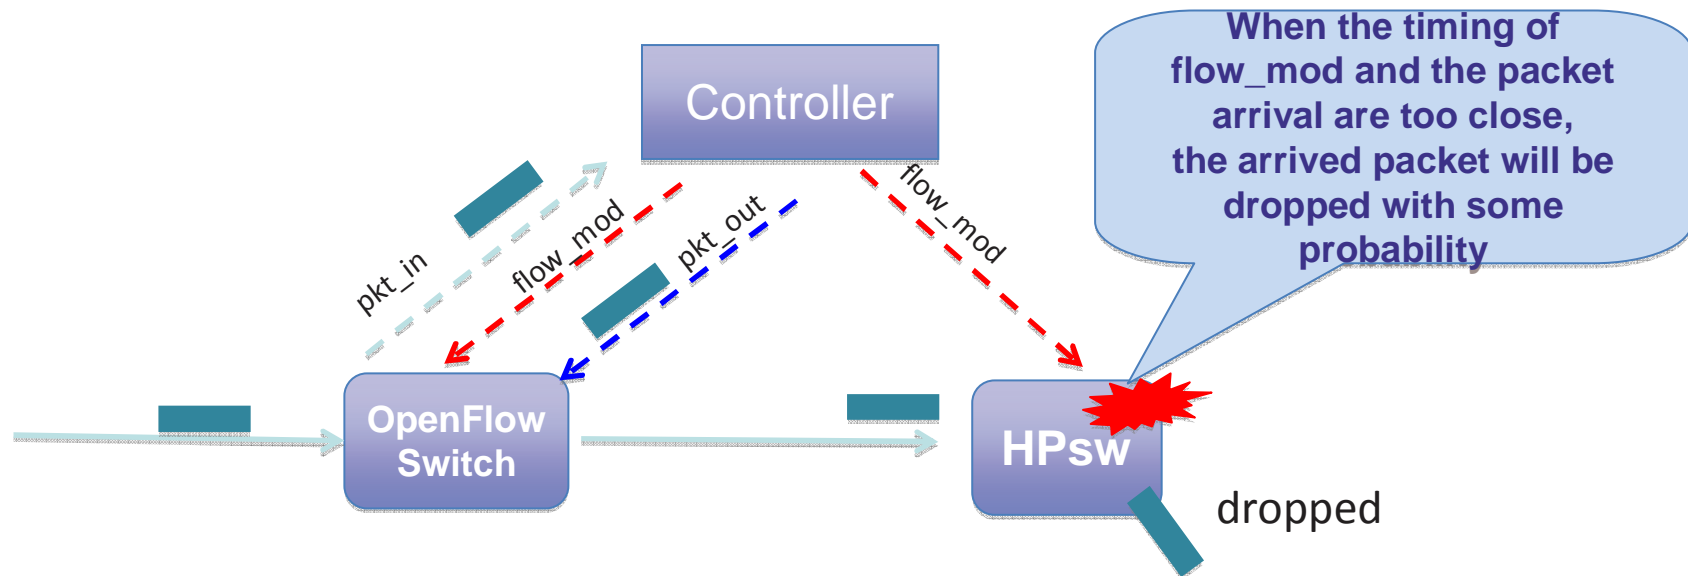
▶ Transmission Control Protocol, Src Port: 42620 (42620), Dst Port: http (80), Seq: 0, Len: 0

# Resolved #1: HP wget performance issue

**CONTROL TRAFFIC INDICATED PROPER OPENFLOW HANDSHAKE FOR FLOW**  
(MAC 0db916ef50->0d055d240, IPV4, 172.24.74.121 -> 171.67.216.18, TCP, 44544 -> 80, HTTP)

1266568066.254337, **PACKET\_IN**, necsw port 35, Buf id 30158480  
1266568066.254483, **FLOW\_MOD**, necsw port 35  
1266568066.254559, **PACKET\_OUT**, necsw port 35, Buf id 30158480  
1266568066.273144, **FLOW\_MOD**, hpsw1 port 47


- Behavior at microscopic level



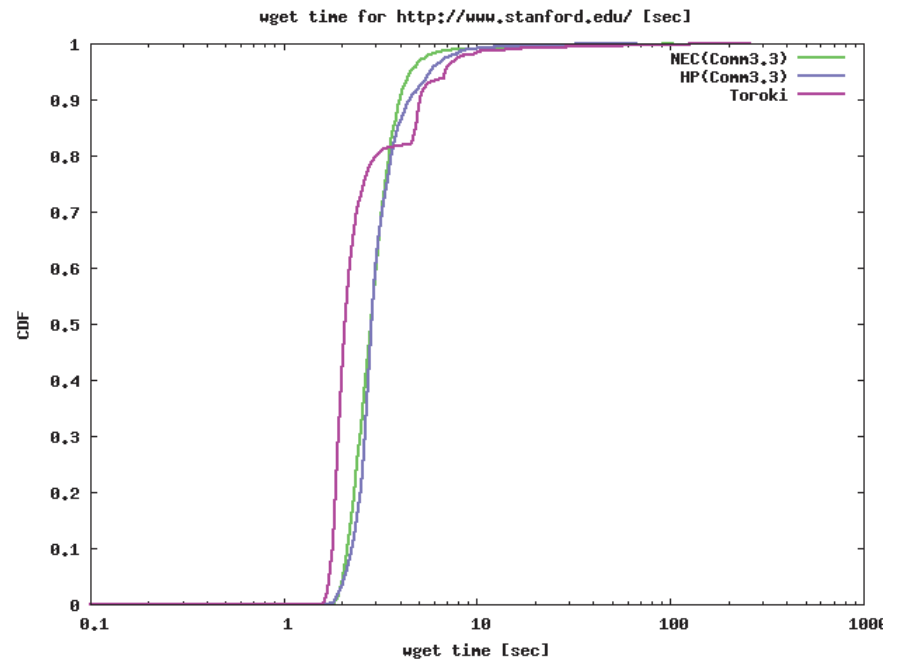
# Resolved #1: HP wget performance issue

- Status: fixed (firmware fix)

Before (Week 38)



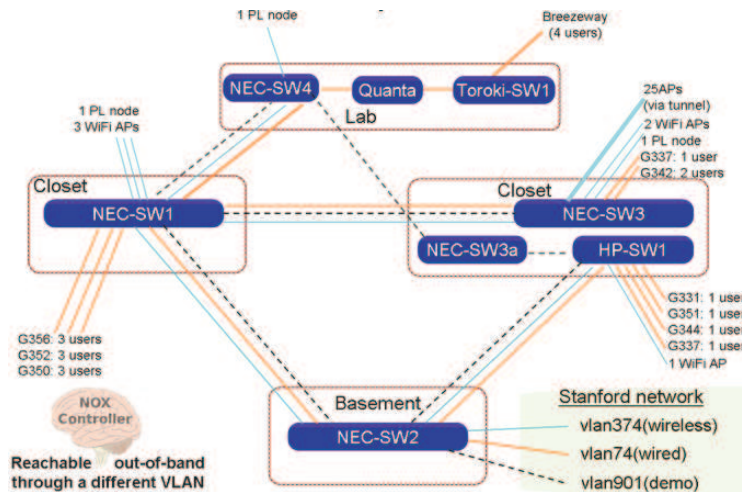
After (Week 41)



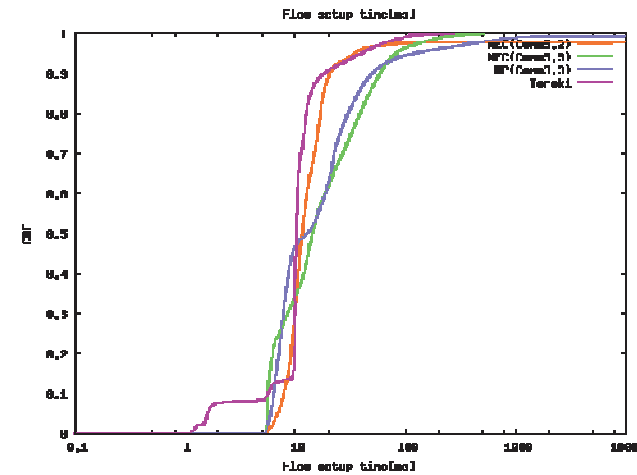
# Stanford OpenFlow deployment

# Status of Stanford deployment

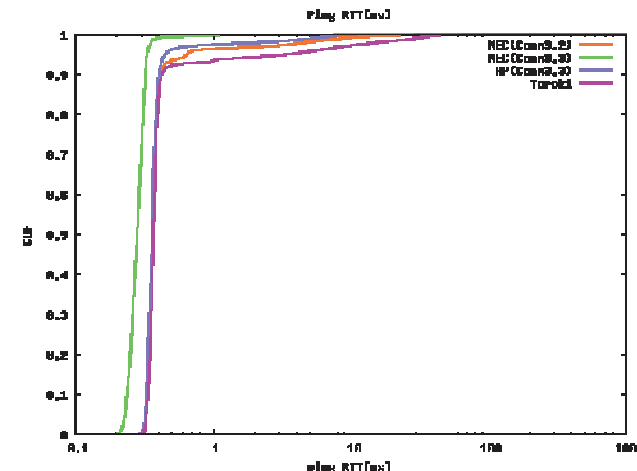
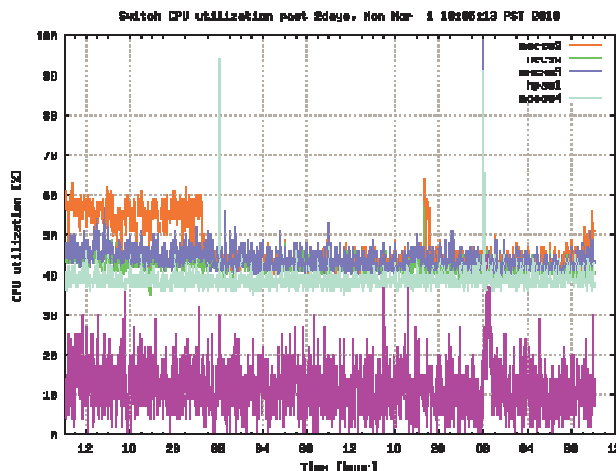
- Network is getting more stable



VLAN 74 in Last week of Feb

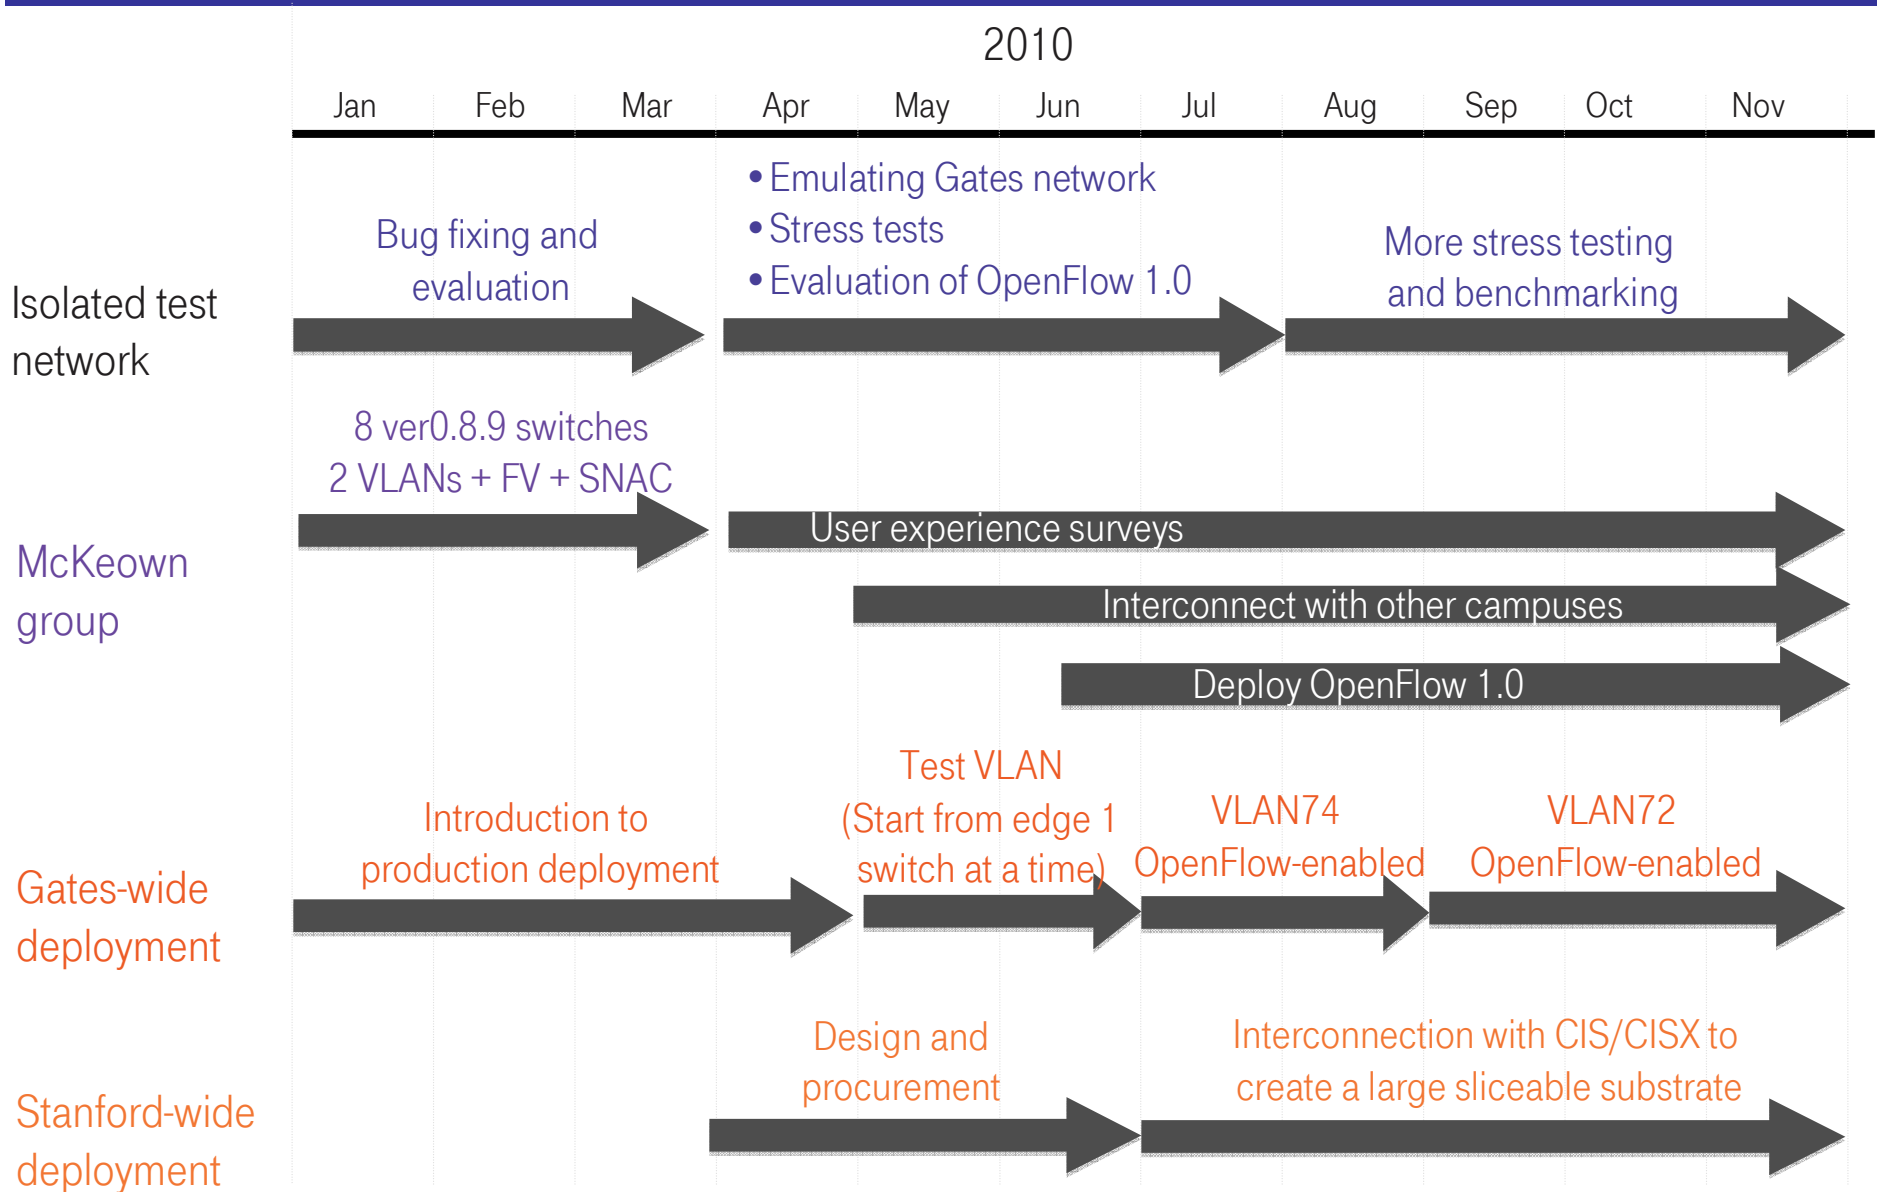


CPU early this month





# Next steps for Stanford deployment



# Summary

- OpenFlow is getting closer to production quality
- Carefully plan "production deployment" to ensure we don't lose trust of our users and campus networking folks
- How may we help you?
  - Are you ready to help other newcomers?