

Opt-In for Enterprise GENI

Goals, architecture and challenges




July 2009

Guido Appenzeller
Stanford University

Disclaimer: I haven't read the Opt-In Document
and this is my first Opt-In Working Group Meeting



Ethernet Switch



Control Path (Software)

The diagram consists of a large gray rounded rectangle containing two smaller light blue rounded rectangles. The top rectangle is labeled 'Control Path (Software)' and the bottom rectangle is labeled 'Data Path (Hardware)'. A horizontal dotted line separates the two rectangles. The entire structure is centered within a white rectangular frame.

Data Path (Hardware)

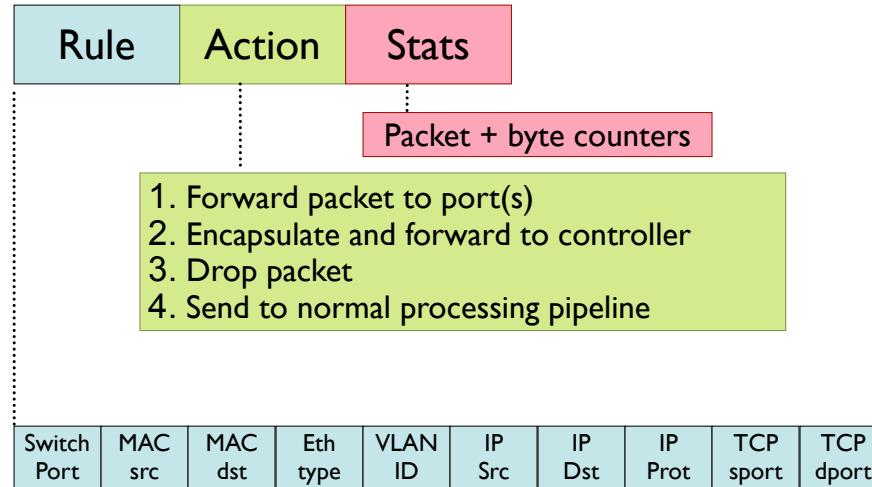
OpenFlow Controller

OpenFlow Protocol (SSL) 



FLOW TABLE ENTRY

Allows to make forwarding decisions based on Layers 1-4

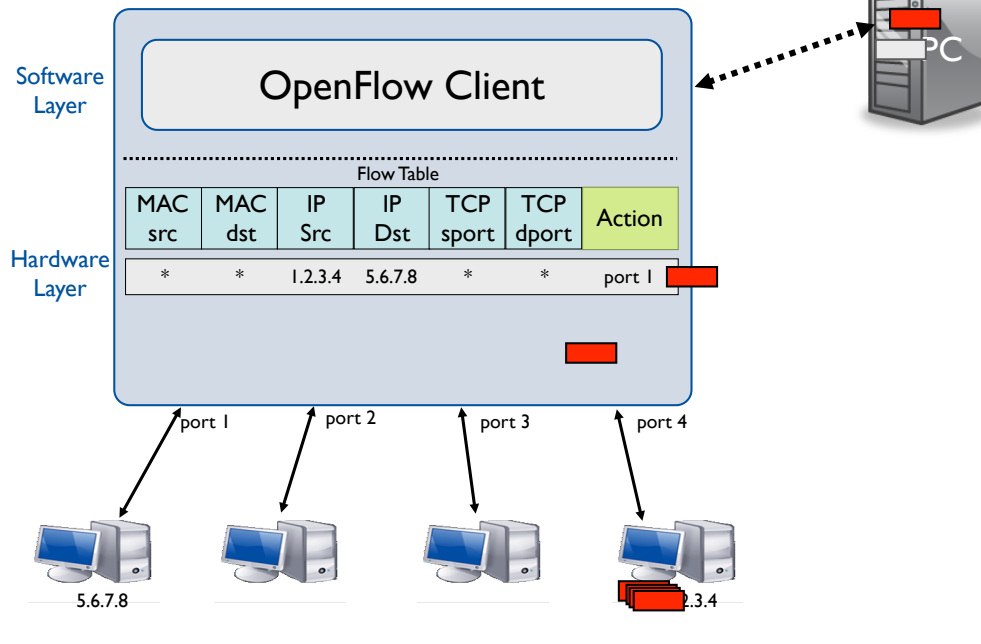


+ mask what fields to match

OPENFLOW EXAMPLE


Forwarding happens at line rates

Controller

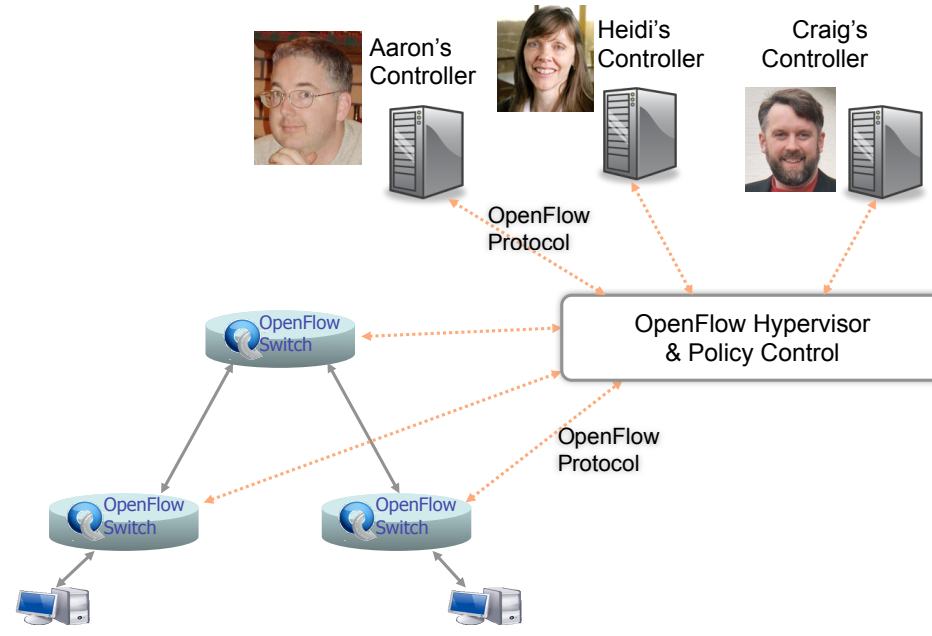


SWITCH BASED VIRTUALIZATION

Exists for NEC, HP switches but not flexible enough for GENI

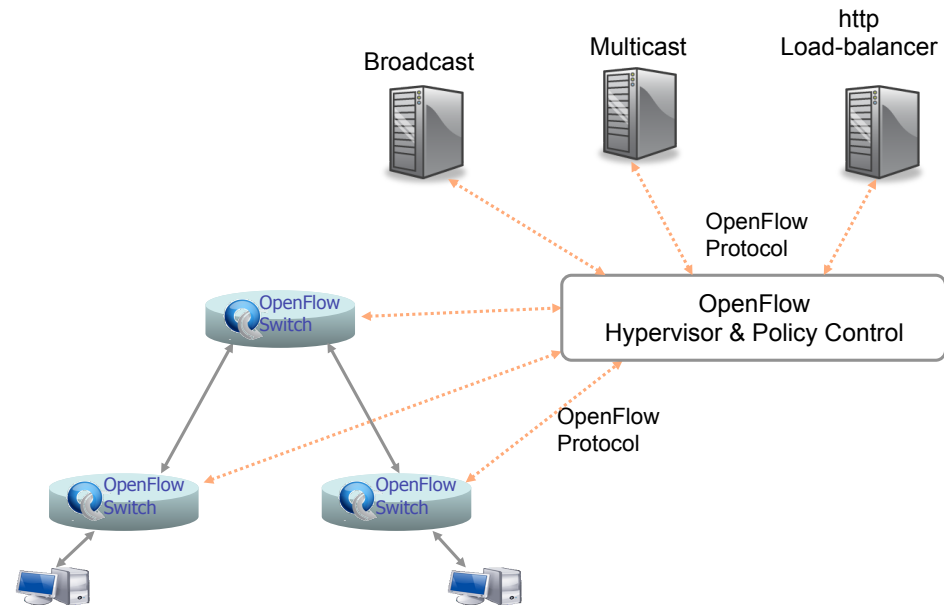


HYPERVISOR BASED VIRTUALIZATION



HYPERVISOR BASED VIRTUALIZATION




Separation not only by VLANs, but any L1-L4 pattern



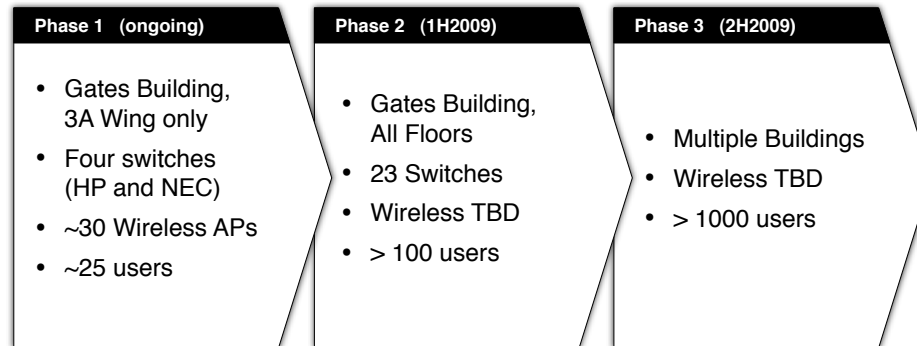
USE CASE: VLAN BASED PARTITIONING

Basic Idea: Partition Flows based on Ports and VLAN Tags

- Traffic entering system (e.g. from end hosts) is tagged
- VLAN tags consistent throughout substrate

	Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport
	*	*	*	*	1,2,3	*	*	*	*	*
	*	*	*	*	4,5,6	*	*	*	*	*
	*	*	*	*	7,8,9	*	*	*	*	*

STANFORD DEPLOYMENT HARDWARE





Our Thoughts on Opt-In

TARGET GROUP

Who do we target for opt in?

- 1000+ local users on the Enterprise GENI Network
 - Can't always assume they are technical
 - Might have locked down PCs, can't install software
 - Constrained Devices - e.g. iPhones
 - User may want to opt-in to multiple experiments or none
- Local Servers
 - Planet Lab Nodes, Production Servers
- Access Points (?)

Why not remote systems via tunnel?

- Route from my home to Stanford is ~10 routed hops
- If majority of time and hops is in tunnel, is experiment relevant?

USER EXPERIENCE

Simplicity of opt-in drives participation

- User gets email about new experiment
 - ▶ “Join the Coral++ experiment today and win an iPod”
 - ▶ User clicks on link in email and browses to web site
 - ▶ User reads description of the experiment
 - ▶ User clicks on “Opt In” button on web site
- Opt-In is about getting users to participate
 - ▶ Most users have little incentive to use experimental network
 - ▶ Each click will lower number of users participating.
 - Corollary: Clicks are bad
- We may need an authentication step for audit trail

HOW DO YOU DO THIS?

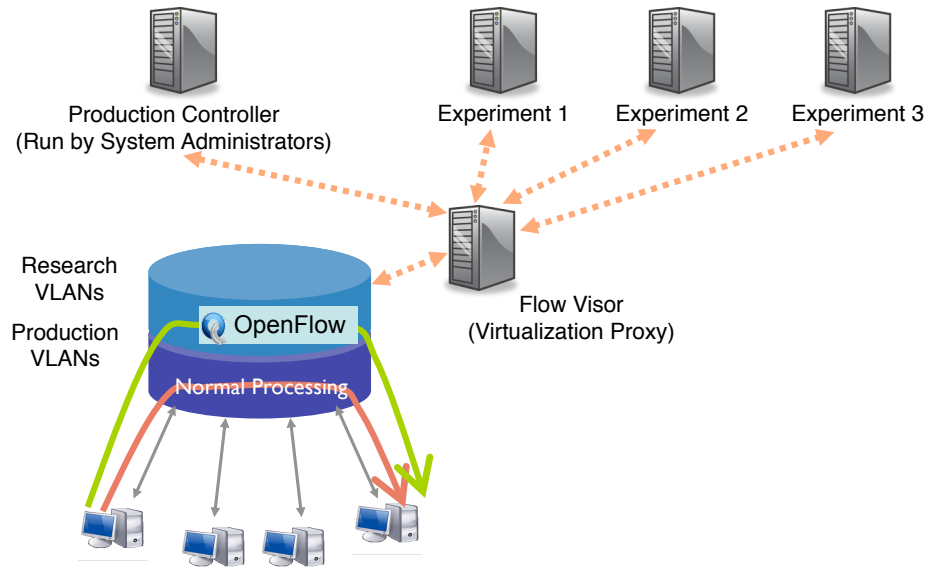
Disclaimer: This is not fully implemented yet, we are still tweaking the architecture. Things may change.

Basic Idea:

- If a user moves into any experiment he is switched from the production VLAN to an OpenFlow controller VLAN
- The Openlow Network has multiple controllers:
 - ▶ A production controller that controls all traffic that is not part of experiments
 - ▶ A FlowVisor that multiplex

SWITCH BASED VIRTUALIZATION

Exists for NEC, HP switches but not flexible enough for GENI



THE OPT-IN MANAGER

The Opt-In Manager is a server in the local substrate. It has the capability to:

- Move users from production VLAN to OpenFlow VLAN
 - Probably via Radius, backup is via console scripts
- Configure the FlowVisor to have a subset of a user's traffic controlled by an experimental controller
- Information about the experiment and what classes of traffic to opt-in is done via the clearing house

OPEN QUESTIONS

- Can experiments that overlap in Flow Space (i.e. **ALL** fields in the packet header could be the same) share the same link?
- What if user A has opted in all IP traffic into Experiment A, and user B has opted in all IP traffic into Experiment B, and they want to talk to each other?
- Is scope of the opt-in by Port, MAC, IP or a Mix?
 - May be constrained by hardware capabilities
- How do we do this across different OpenFlow slivers?
- How do we do this across different slivers with different architectures?
- Will this solution work for other campuses?

Thank You!

