

OPENFLOW & SOFTWARE DEFINED NETWORKING

Greg Ferro
EtherealMind.com and PacketPushers.net

HUH ? OPENFLOW.


- **What is OpenFlow ? From the bottom up. With big words.**
- **How OpenFlow does stuff.**
- **Then WHY we want OpenFlow to do that**


OPENFLOW.
WHAT IS IT ?

Nerd Glasses on please


PLANES OF OPERATION

CLI/SSH/SNMP/XML






FORWARDING TABLES




Forwarding Tables =
Forwarding Information Base
= FIB

FORWARDING TABLES → RIBS



ROUTING PROTOCOLS → FIBS



ROUTING PROTOCOLS GOOD ?

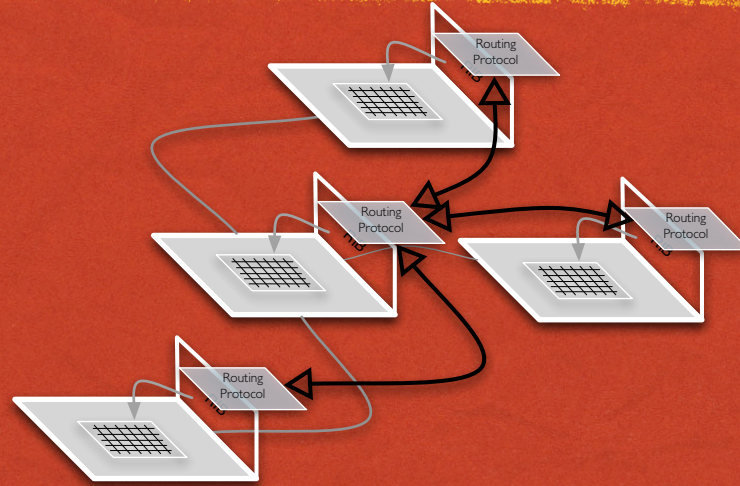
- **reliable**
- **proven**
- **deterministic (knowable)**
- **self healing**
- **autonomous**
- **scalable ?**

ROUTING PROTOCOLS

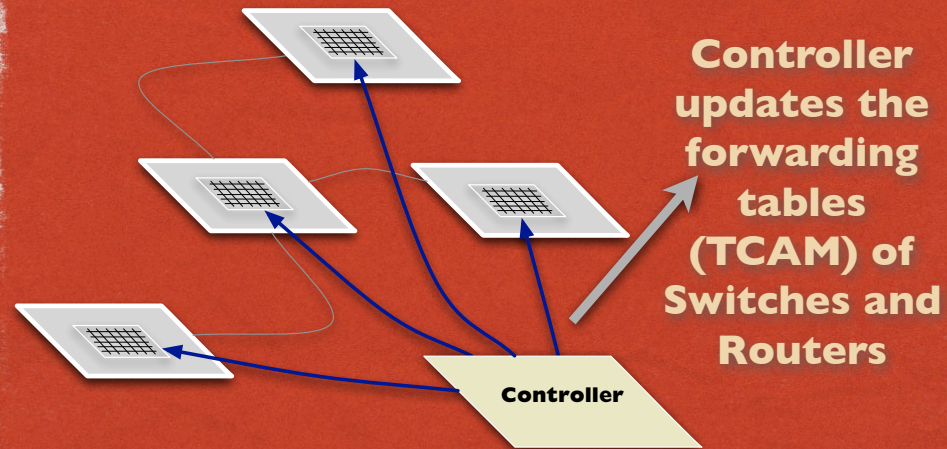
..... OR BAD?

- **loosely coupled**
- **hop by hop**
- **homogenous system**
- **change resistant**
- **limited external configuration**
- **not load based**
- **destination only**

WHAT IF, INSTEAD OF THIS:



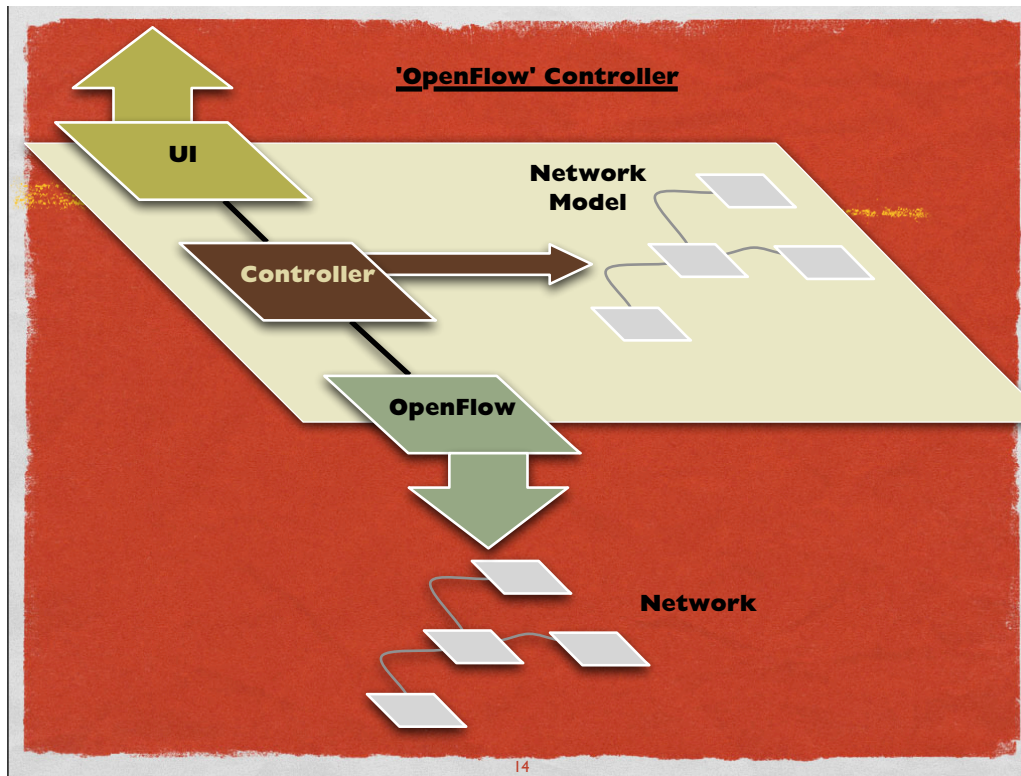
WE USED A CENTRAL CONTROLLER



CONTROLLER ?

- **it's software, it's a program***
- **connects to all devices**
- **builds a network topology**
- **runs an algorithm**
- **then updates the Forwarding table by an API.**

* (& we don't know much about it yet)



!HOW IT DOES STUFF

Keep your Nerd Glasses on please

FLOW TABLES

OpenFlow Draft Specification 1.1

4.1 Flow Table

A flow table consists of flow entries.

Match Fields	Counters	Instructions
--------------	----------	--------------

Table 1: Main components of a flow entry in a flow table.

Each flow table entry (see Table 1) contains:

- **match fields:** to match against packets. These consist of the ingress port and packet headers, and optionally metadata specified by a previous table.
- **counters:** to update for matching packets
- **instructions** to modify the action set or pipeline processing

OPENFLOW

- controller-to-switch
- asynchronous
- symmetric

MATCH FIELDS

- Ingress Port
- Metadata
- Ether src
- Ether dst
- Ether type
- VLAN id
- VLAN priority
- MPLS label , MPLS traffic class
- IPv4 SRC, IPv4 DST
- IPv4 proto (ARP , ICMP , TCP , UDP , SCTP , ICMP code, IPv4 ToS bits)
- TCP/ UDP / SCTP src port
- ICMP Type
- TCP/ UDP / SCTP dst port
- ICMP Code

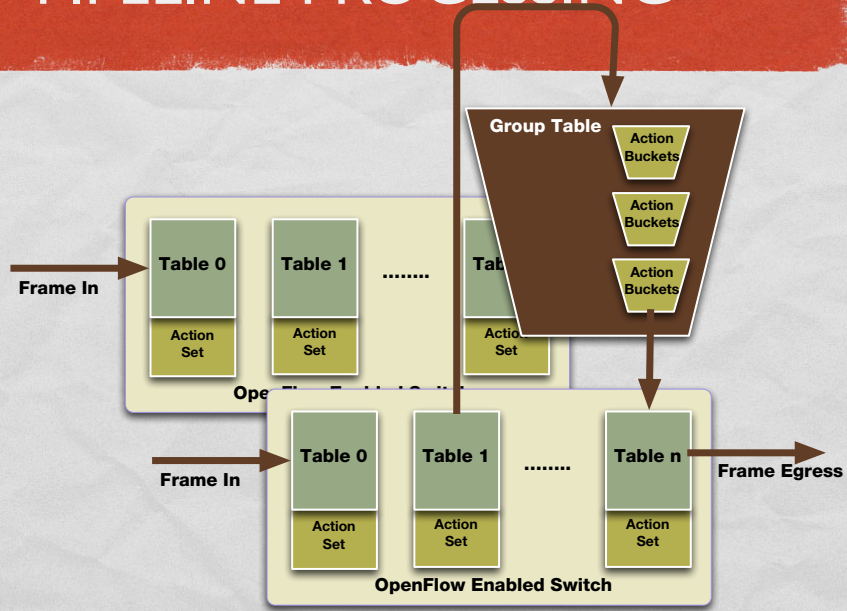
Pretty Good Selection

MATCH FLEXIBILITY

Field	Bits	When applicable	Notes
Ingress Port	32	All packets	Numerical representation of incoming port, starting at 1. (physical or virtual port)
Metadata	64	Table 1 and above	
Ethernet source address	48	All packets on enabled ports	Can use arbitrary bitmask
Ethernet destination address	48	All packets on enabled ports	Can use arbitrary bitmask
Ethernet type	16	All packets on enabled ports	Ethernet type of the OpenFlow packet payload, after VLAN tags. 802.3 frames have special VLAN identifier of outermost VLAN tag.
VLAN id	12	All packets with VLAN tags	VLAN identifier of outermost VLAN tag.
VLAN priority	3	All packets with VLAN tags	VLAN priority of outermost VLAN tag.
MPLS label	20	All packets with MPLS tags	Match on outermost MPLS tag.
MPLS traffic class	3	All packets with MPLS tags	Match on outermost MPLS tag.
IPv4 source address	32	All IPv4 and ARP packets	Can use subnet mask or arbitrary bitmask
IPv4 destination address	32	All IPv4 and ARP packets	Can use subnet mask or arbitrary bitmask
IPv4 protocol / ARP opcode	8	All IPv4 and IPv4 over Ethernet, ARP packets	Only the lower 8 bits of the ARP op-code are used.
IPv4 ToS bits	6	All IPv4 packets	Specify as 8-bit value and place ToS in upper 6 bits.
Transport source port / ICMP Type	16	All TCP, UDP, SCTP, and ICMP packets	Only lower 8 bits used for ICMP Type
Transport destination port / ICMP Code	16	All TCP, UDP, SCTP, and ICMP packets	Only lower 8 bits used for ICMP Code

Shortcomings Exist

PIPELINE PROCESSING



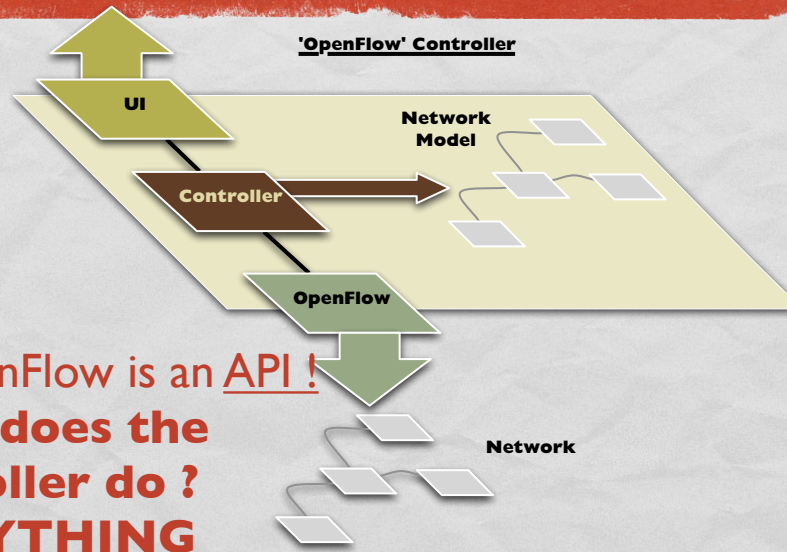
INSTRUCTIONS

- **Apply-Actions actions** :Applies the specific actions immediately.
- **Clear-Actions:** Clears all the actions in the action set immediately.
- **Write-Actions actions:** Merges the specified actions into the current action set
- **Write-Metadata metadata / mask:** Writes the masked metadata value into the metadata field.

WILDCARDS

MAC SRC	MAC DST	SRC IP	IP DST	TCP Dport	TCP SPort	Action	Count
*	00:02:.	*	*	*		Port 1	250
*	*	*	10.2.2.1	80	*	Port 3	320
*	*	192.*	*	*	*	drop	890
*	*	192.*	*	*	*	local	100
*	*	*	*	*	*	Controller	11

IT'S ALL ABOUT THE CONTROLLER



- So OpenFlow is an API !
- **What does the controller do ?**
- **EVERYTHING**

REALITY CHECK

- “[OpenFlow] doesn’t let you do anything you couldn’t do on a network before” - Scott Shenker
- OpenFlow You don’t need OpenFlow to solve every age-old problem. - Ivan Pepelnjak

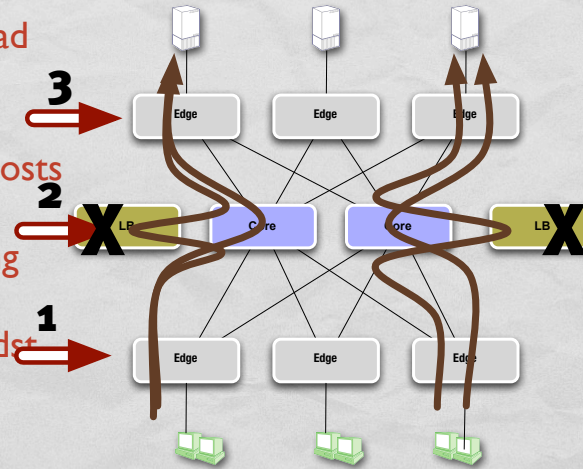
SO WHY DO IT ALL ?

ACTION POSSIBILITIES

set output port	unequal path load balancing, multipath routing,
modify IPv4 fields	NAT, QoS munging,
modify Ethernet fields	set VLAN, set TOS,
push/pop MPLS tags	interoperate existing MPLS networks
	blah blah
	blah
	<u>CLOUD</u>

EG, LOAD BALANCING

- if every switch is a load balancer
- then distributed processing means it costs nothing
- load balance by setting forwarding tables to balance different src/dst IP to different paths.



WHY WE WANT OPENFLOW

Nerd Glasses off please

GOOD PARENTS

- Martin Casado, Nick McKeown
- Big Vendors: Juniper, Brocade, NEC,
- Small Vendors: OpenGear, et al
- Flotilla of startups - Big Switch, Nicira etc
- Big name companies Google, Yahoo, Facebook
- Indiana University setup “official” interoperability lab

ADD OPENFLOW *REPLACE NOTHING*

“Virtual tenant network
plane in the controller”

Keep your existing Control
plane OSPF, BGP, ISIS,
MPLS TE.
ADD OpenFlow, don't
replace anything

TOPOLOGY LITE

Topology independent
virtual networks

Forwarding tables don't
dictate topology.

Cabling paths do.
Routing protocols do.
Business partnerships do.

DYNAMIC

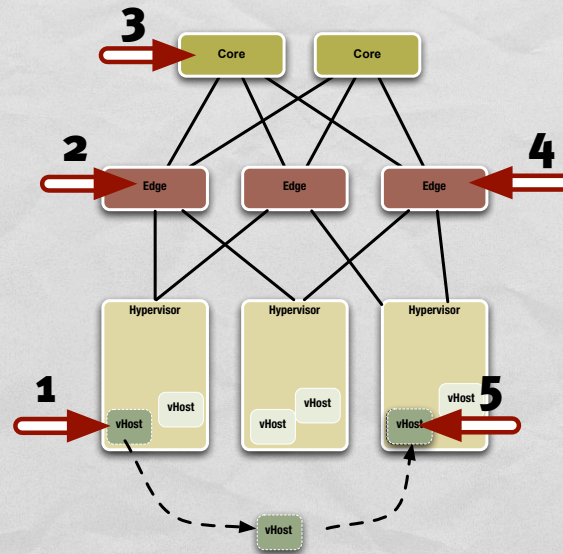
Controller can repeatedly update forwarding tables

- Power saving (shutdown unneeded nodes)
- change per flow topology on a “time” basis
-

SECURITY

- Divert flow to service
eg. IPS, Proxy, NAC,
- perform security tasks
or check
- change forwarding table
to forward

DATA CENTRE



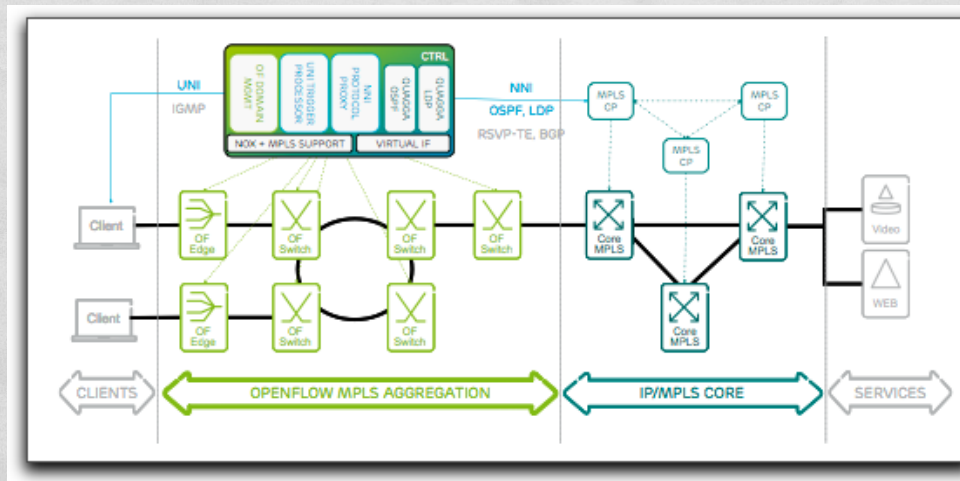
VM Migration

- no arp
- no routing
- no ip mobility
- works today

CAMPUS

- Indiana University
- 3000 Access Points
- Students, Teachers, Researchers, Admin
- Dormitory, Classrooms, Offices,
- **Mobile forwarding privileges according to user type, user location, application**
- Student in dormitory / classroom
- Researchers in Lans

INTEGRATED MPLS/OPENFLOW



3G NETWORKS

- User control
- DPI redirection
- Counters

WRAP

- OpenFlow still early
- It's all about the controller giving us a better control plane
- the controller is a another way of defining the flow forwarding in the network
- software defined networking
- you can do both, at once
- Many Open Source projects
- Many more commercial projects



PacketPushers
&
TechFieldDay

presents

**OpenFlow
Symposium**

**Practical and Applied
OpenFlow Networks**

LIVE VIDEO STREAM!

26TH OCTOBER, 2011

<http://bit.ly/AppliedOpenFlow>

LINKS

- <http://www.openflow.org>
- <http://www.opennetworkingfoundation.org/>
- <http://www.openvswitch.org>
- <http://noxrepo.org/wp/>
- <http://www.openflowhub.org/display/Beacon>
-

Thankyou



Packet Pushers

<http://packetpushers.net>



Ethereal Mind

Data Networking: design, thinking, worklife, study

<http://etherealmind.com>