




# Introduction to OpenFlow

Presented by:  
Shangxin Du – TAC DC Solution

20 April, 2015

# *For the next 60 minutes ...*



Forget everything you know about:

- Transparent Bridging
- STP
- L2 Forwarding
- IP Routing
- L3 Forwarding
- And so on ...

# What is OpenFlow?



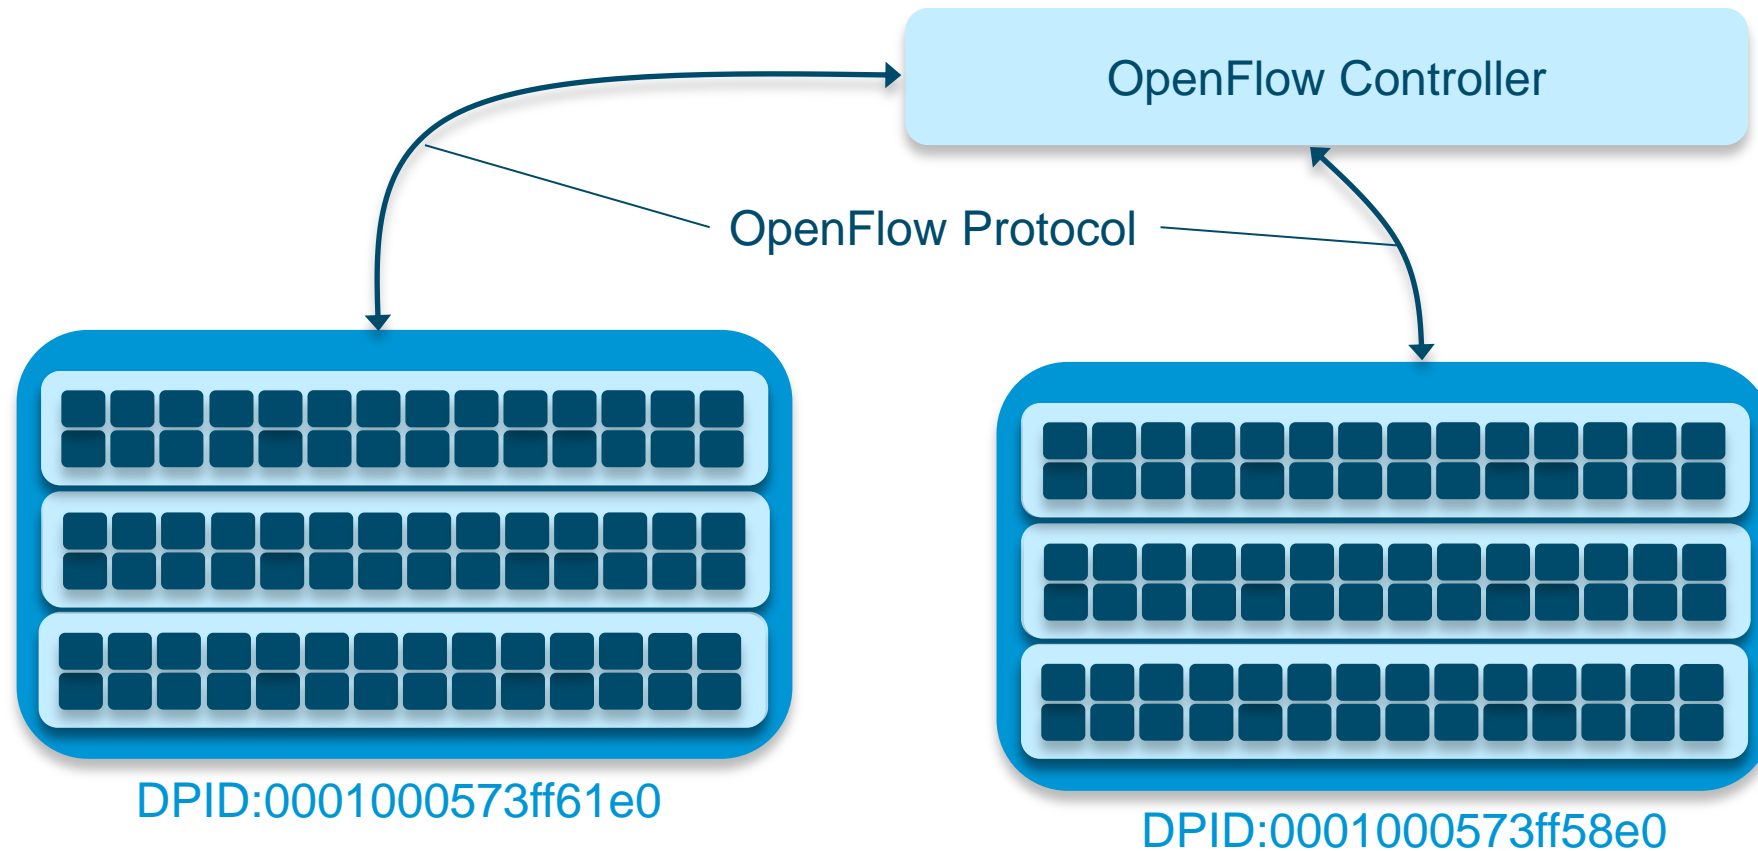
“standard communications interface  
between control and forwarding layers of an  
SDN architecture ...

allows direct manipulation of the forwarding  
plane of network devices such as physical  
and virtual switches and routers”

Open Networking Foundation

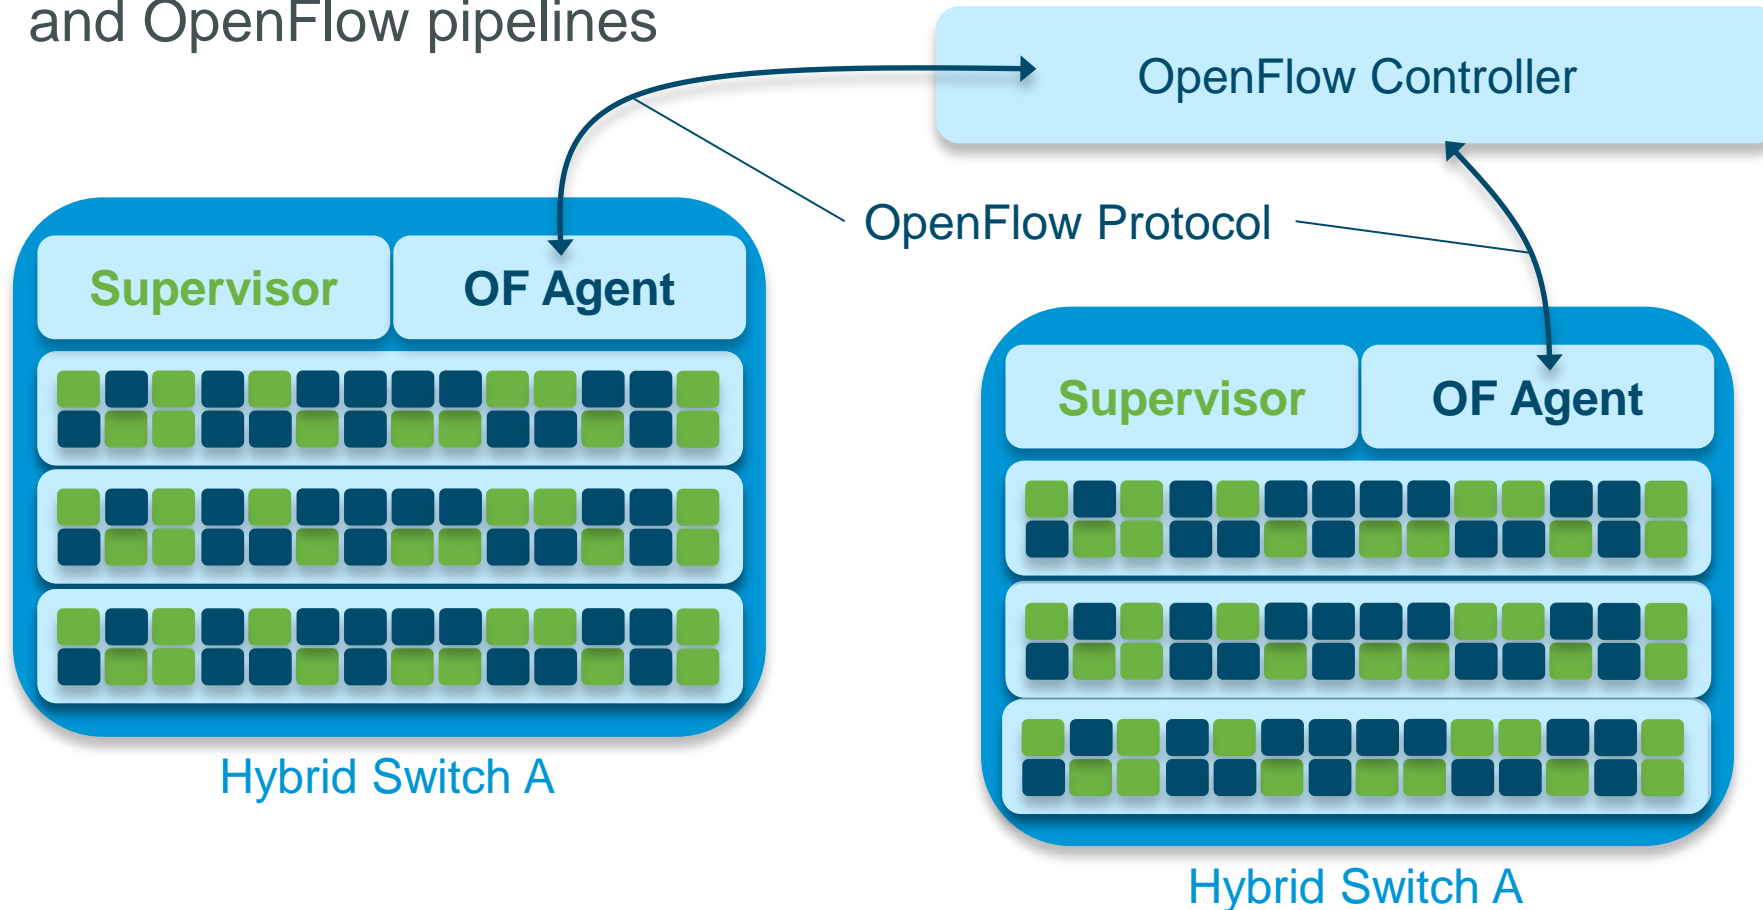
# OpenFlow-only Switch Architecture

- Control plane is decoupled from data plane; centralized and given span of control over multiple data plane switch elements
- Controller manages switch flow table using OpenFlow protocol
  - Add, update, delete flow entries, both reactively and proactively

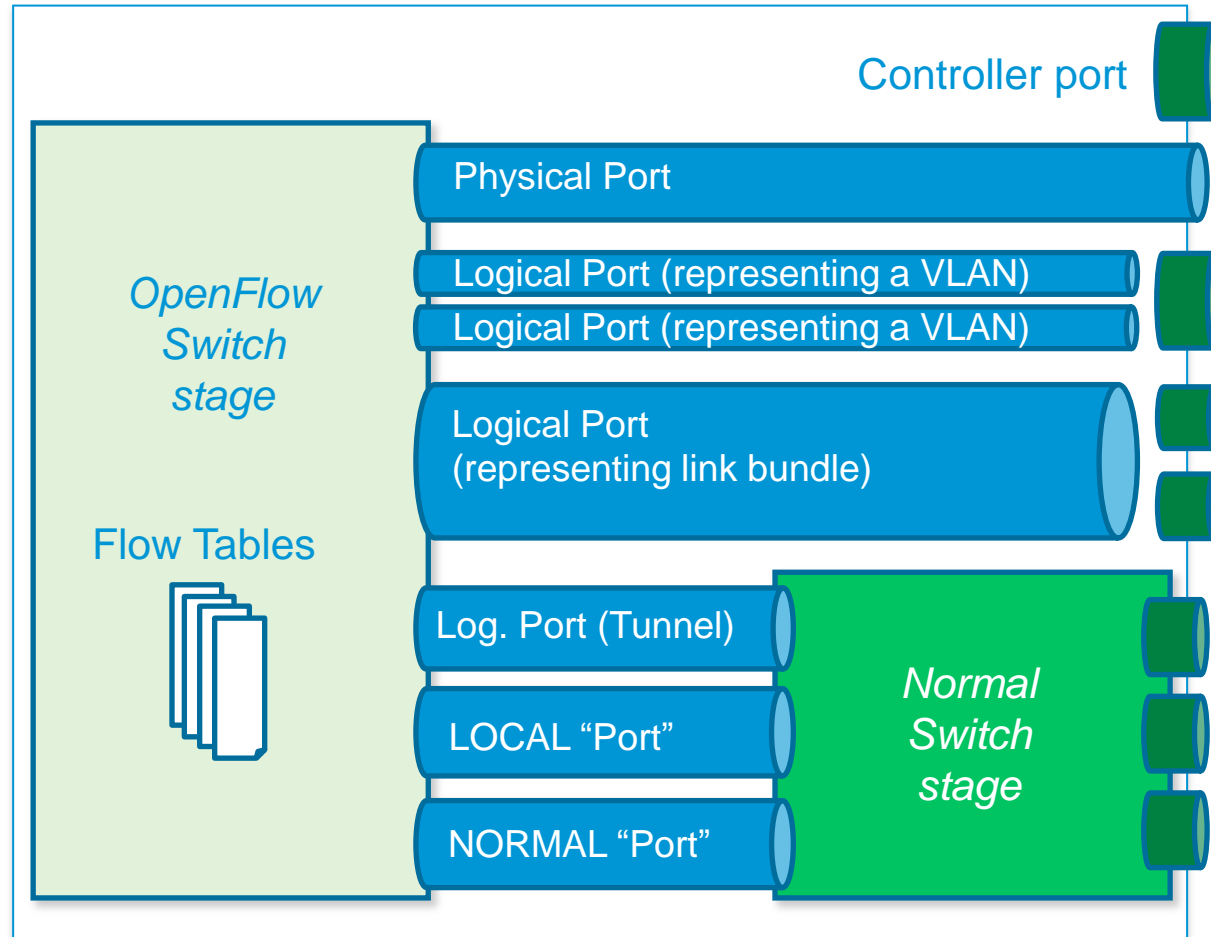


# OpenFlow-hybrid Switch Architecture

- Traditional captive control plane/data plane components co-exist with OpenFlow Switch capabilities
- OpenFlow defines logical “ports” for passing packets to traditional and OpenFlow pipelines

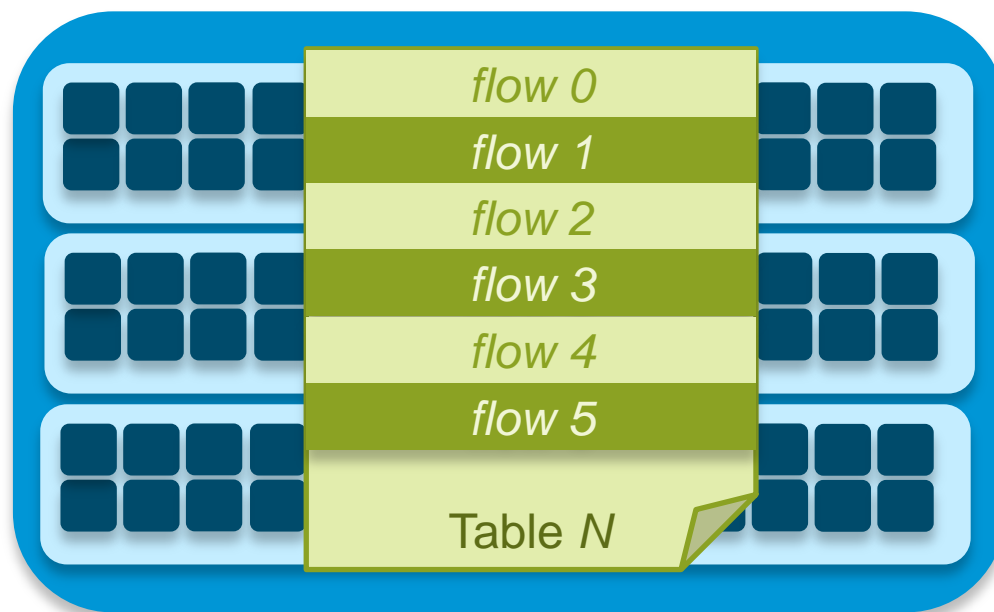


# OpenFlow Ports



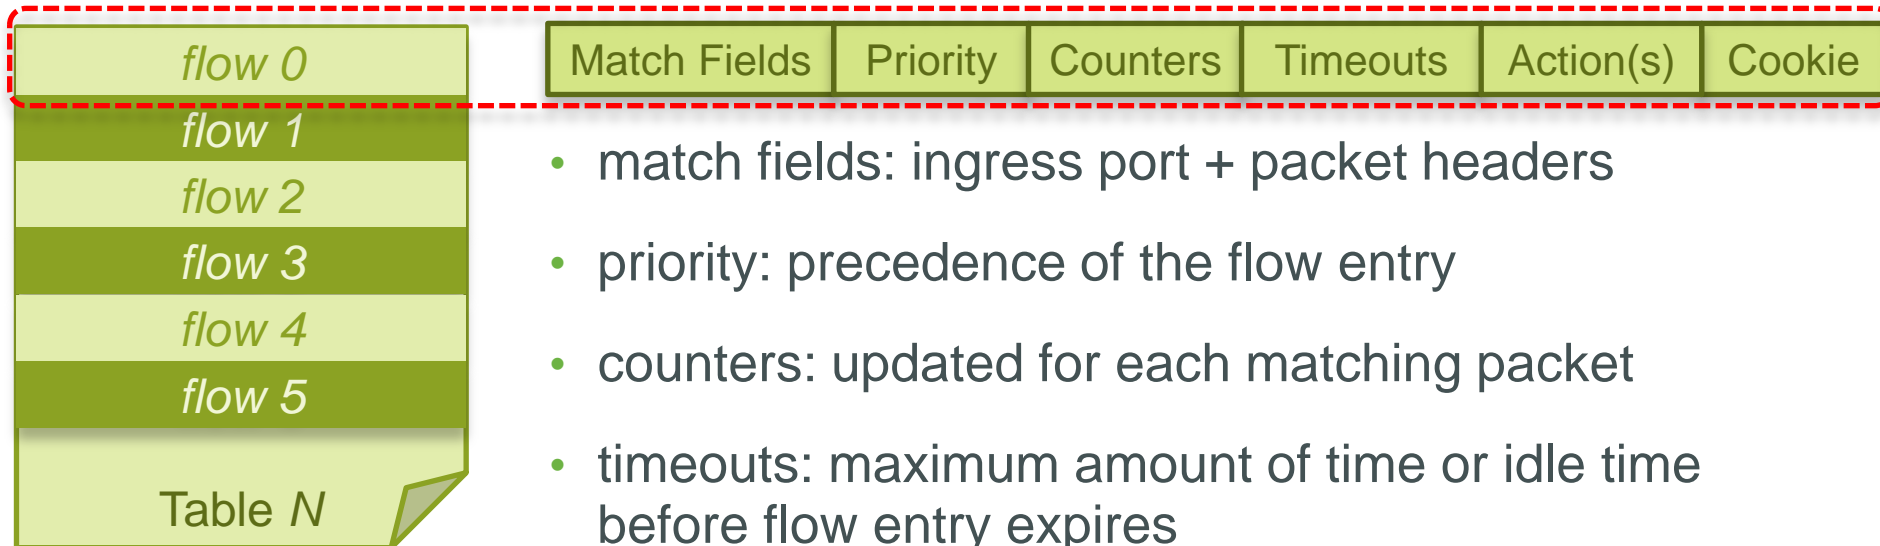
# Flow Tables (OpenFlow 1.0)

- Each OpenFlow Switch has a single flow table, which is used for packet lookup and forwarding
- Packets not matching an entry in the flow table are punted to the controller



# Flow Entries

- A flow table consists of one or more flow entries
- Each entry consists of one row in the flow table



- match fields: ingress port + packet headers
- priority: precedence of the flow entry
- counters: updated for each matching packet
- timeouts: maximum amount of time or idle time before flow entry expires
- action(s): one or more actions to take on match packets
- cookie: opaque data chosen by controller



# Flow Table Header Fields

FLOW TABLE		
HEADER FIELDS	COUNTERS	ACTIONS
...	...	...
...	...	...

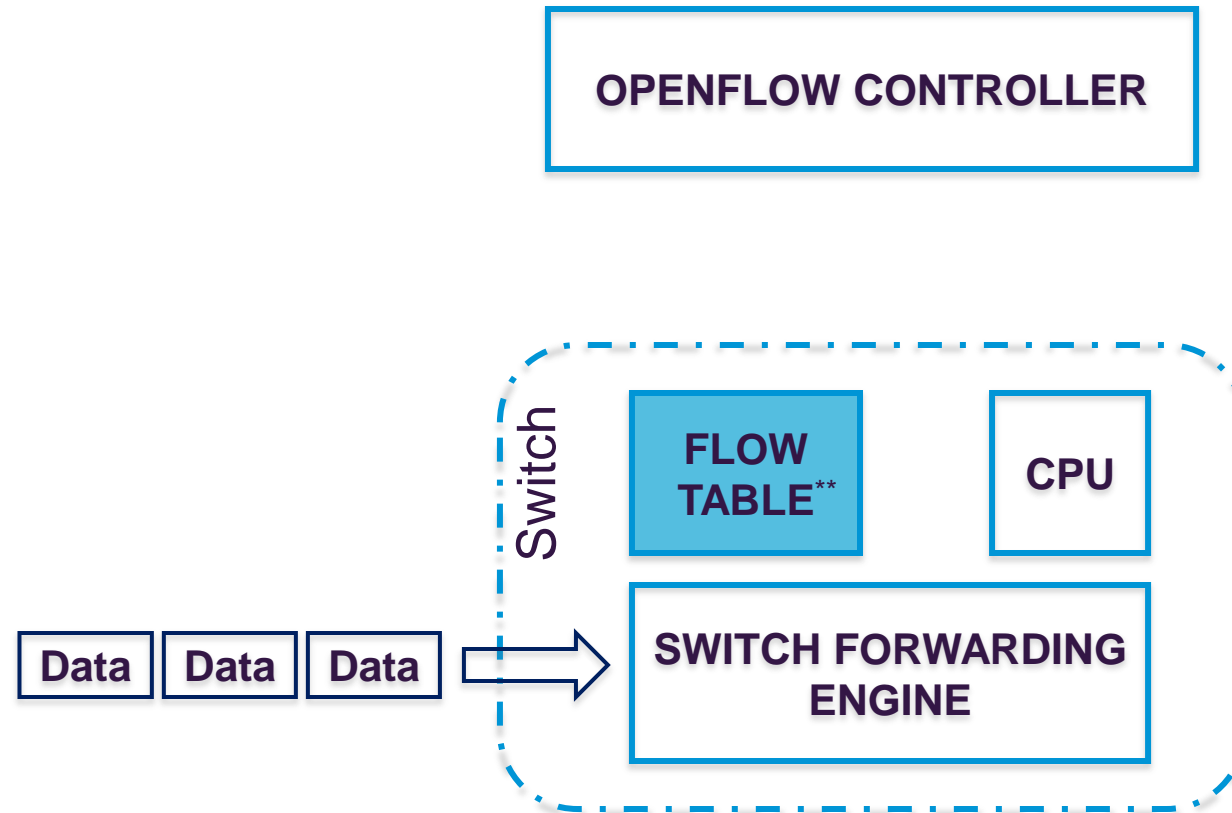


OF1.0 HEADER FIELDS

Ingress Port	Source MAC	Dest MAC	Ether Type	VLAN ID	VLAN Priority	IP SRC	IP DEST	IP Proto	IP TOS	TCP/UDP SRC	TCP/UDP DEST
1	2	3	4	5	6	7	8	9	10	11	12

# OpenFlow Forwarding Model

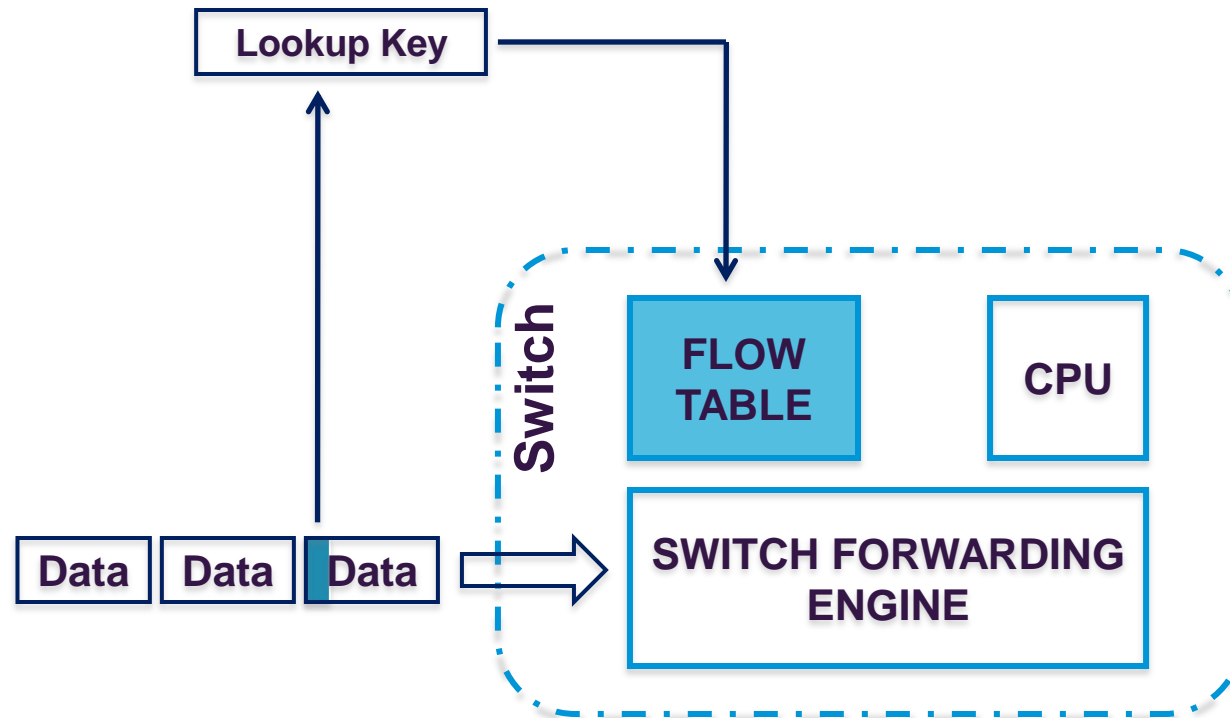
Incoming Packet Arrives at Switch



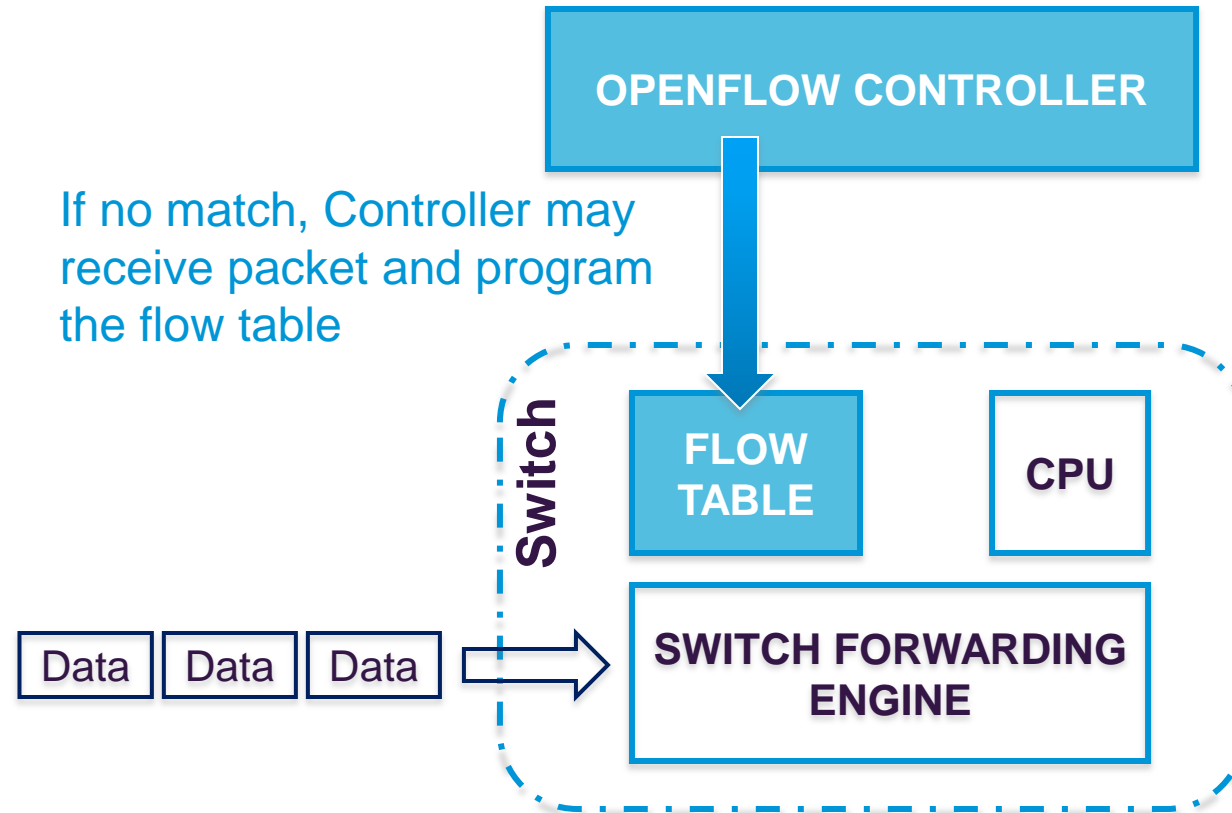
\*\*OpenFlow 1.0 supports a lookup into a single flow table

# Header Fields are Used for Lookup

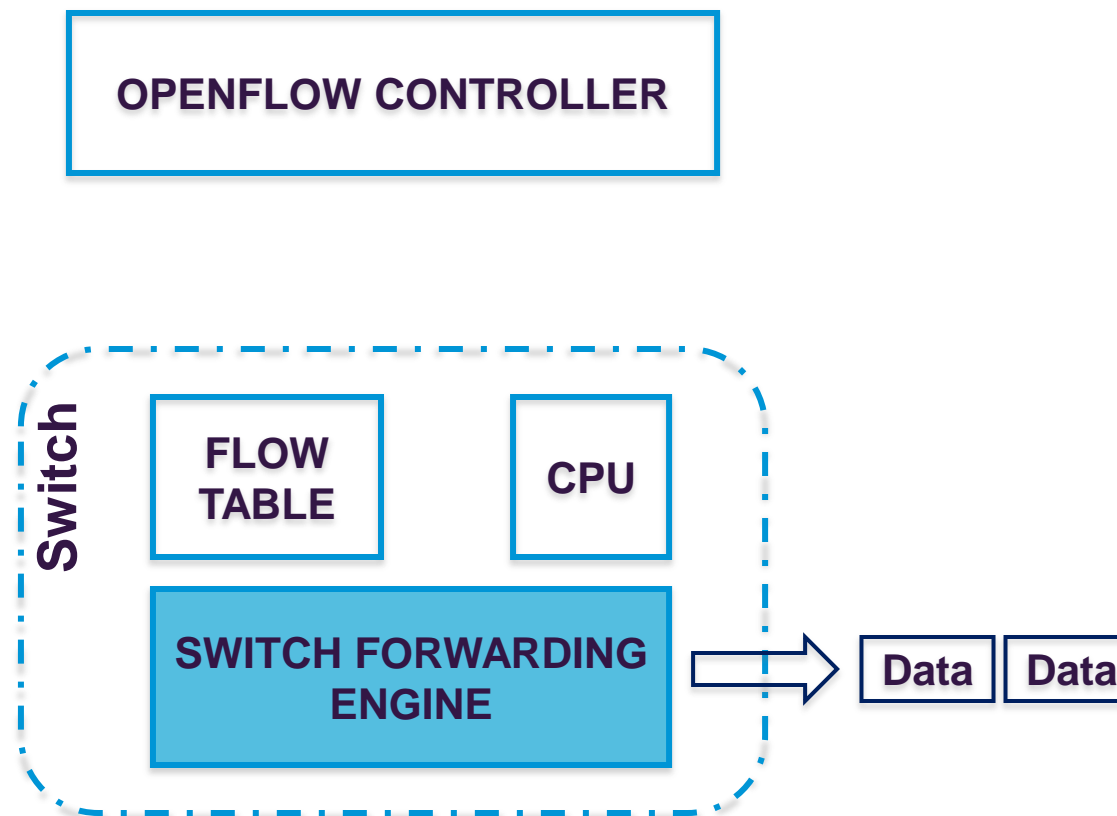
Header fields used to build lookup key



# Controller Writes to Switch Flow Table



# Forwarding Engine Forwards Packets




# Flow Table Counters and Actions

FLOW TABLE		
HEADER FIELDS	COUNTERS	ACTIONS
...	...	...
...	...	...




Several important OF1.0 Actions  
Let us explore in more detail...

# Action #1: Forward to all ports except Input




# Action #2: Redirect to Controller

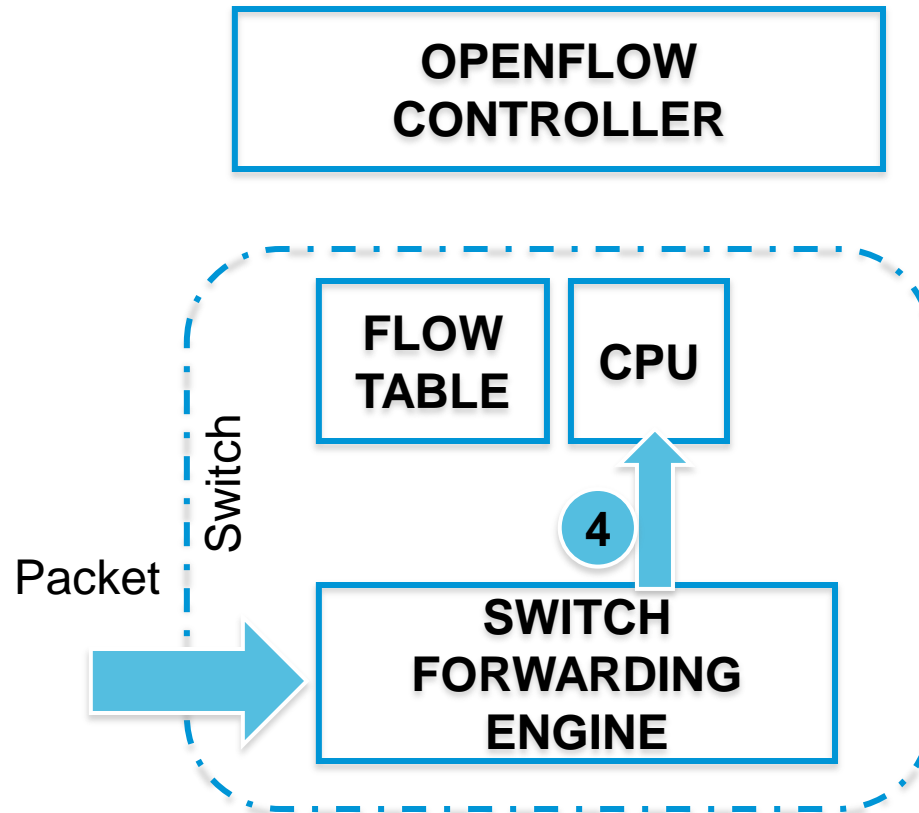




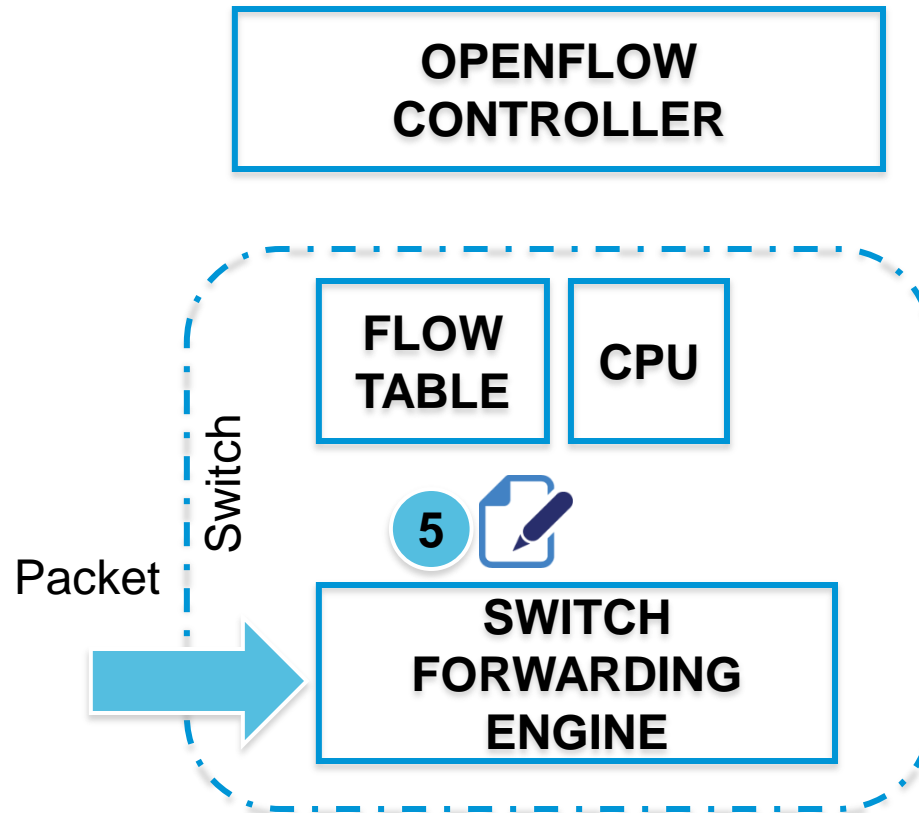
# Action #3: Output packet from Controller




# Action #4: Forward to Local CPU




# Action #5: Rewrite Packet Headers



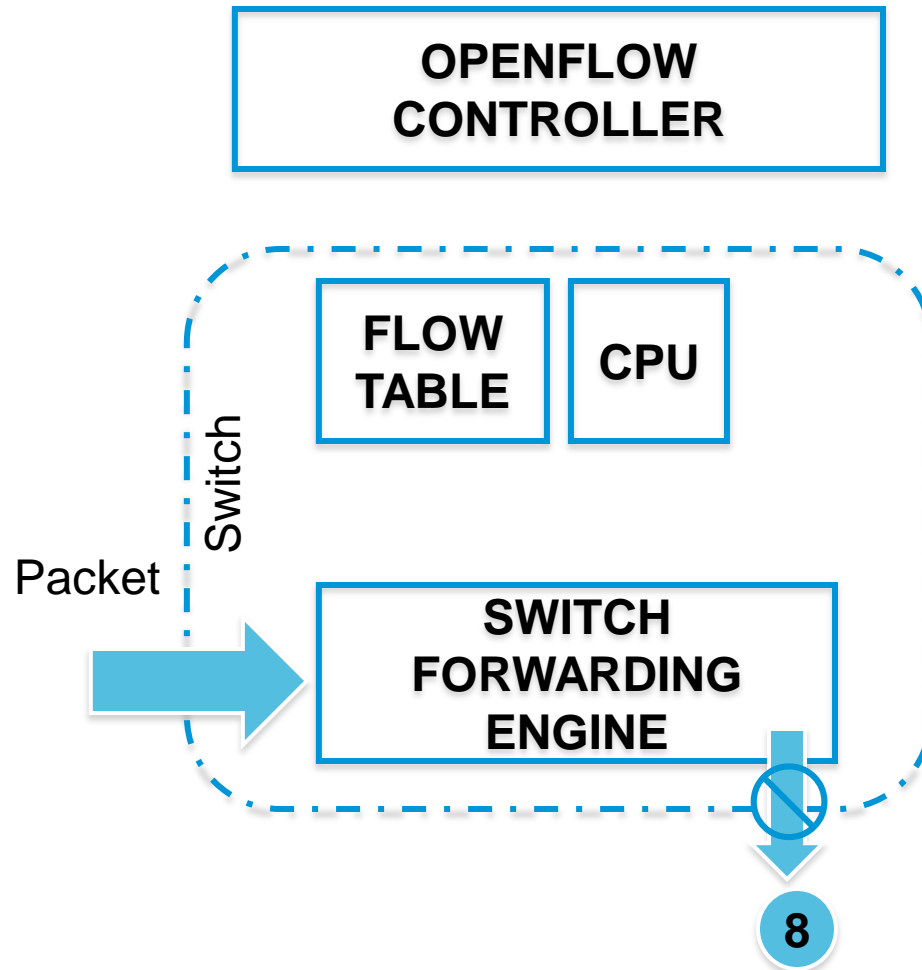
# Action #6: Forward to Input Port



# Action #7: Forward to Destination Port




# Action #8: Drop Packet





# OpenFlow 1.1



- |   |  |
|---|--|
| <ul style="list-style-type: none"><li>• Single flow Table</li><li>• L2, IPv4 focused matching</li><li>• Basic actions</li></ul> | <ul style="list-style-type: none"><li>• Multiple flow tables</li><li>• Group table</li><li>• Packet processing options</li><li>• MPLS</li><li>• VLAN</li></ul> |
|---|--|



# OpenFlow 1.2

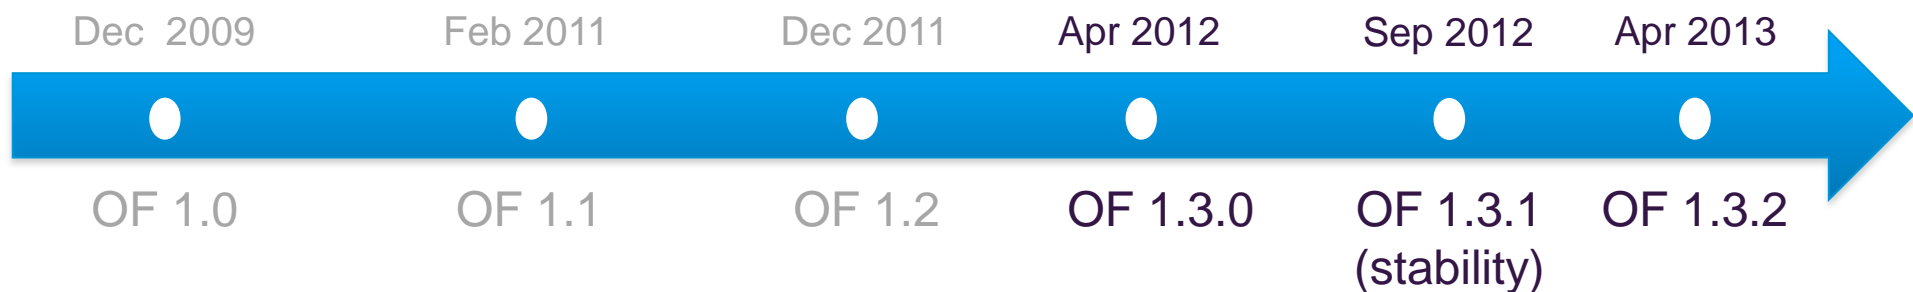


- Single flow Table
- L2, IPv4 focused matching
- Basic actions

- Group table
- Multiple flow tables
- Packet Processing Options
- MPLS, VLAN

- IPv6
- Extensible matching

# OpenFlow 1.3.x



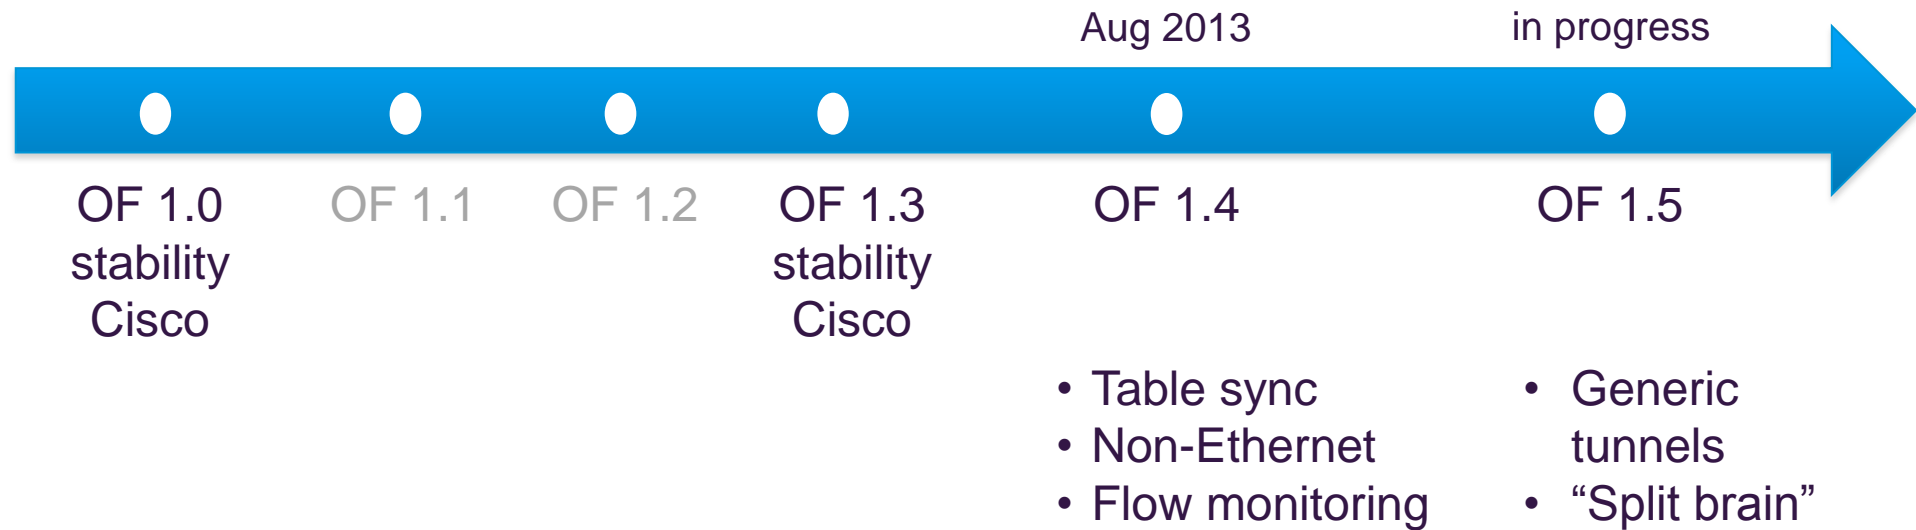
- Single flow Table
- L2, IPv4 focused matching
- Basic actions

- Group table
- Multiple flow tables
- Packet Processing Options
- MPLS, VLAN

- IPv6
- Extensible matching

- IPv6 Extended Headers
- Meter table
- Auxiliary Connections
- Advanced MPLS
- PBB
- Version negotiation
- Controller connections

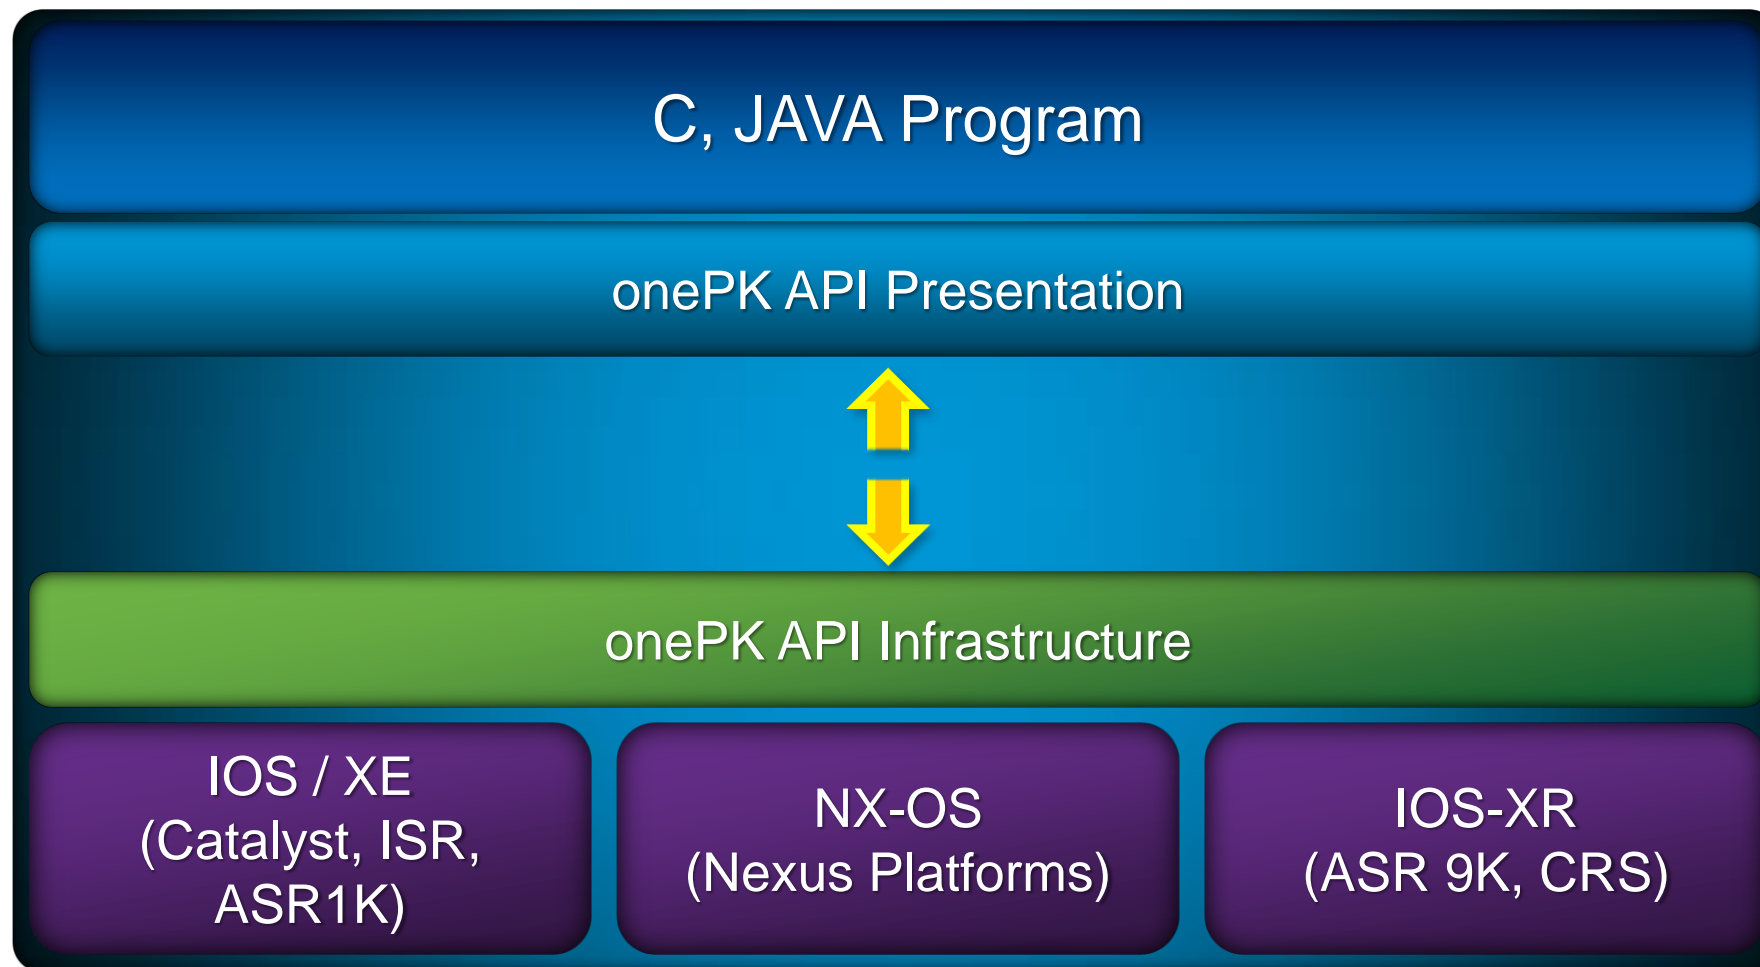
# OpenFlow stability, OpenFlow 1.4 and up



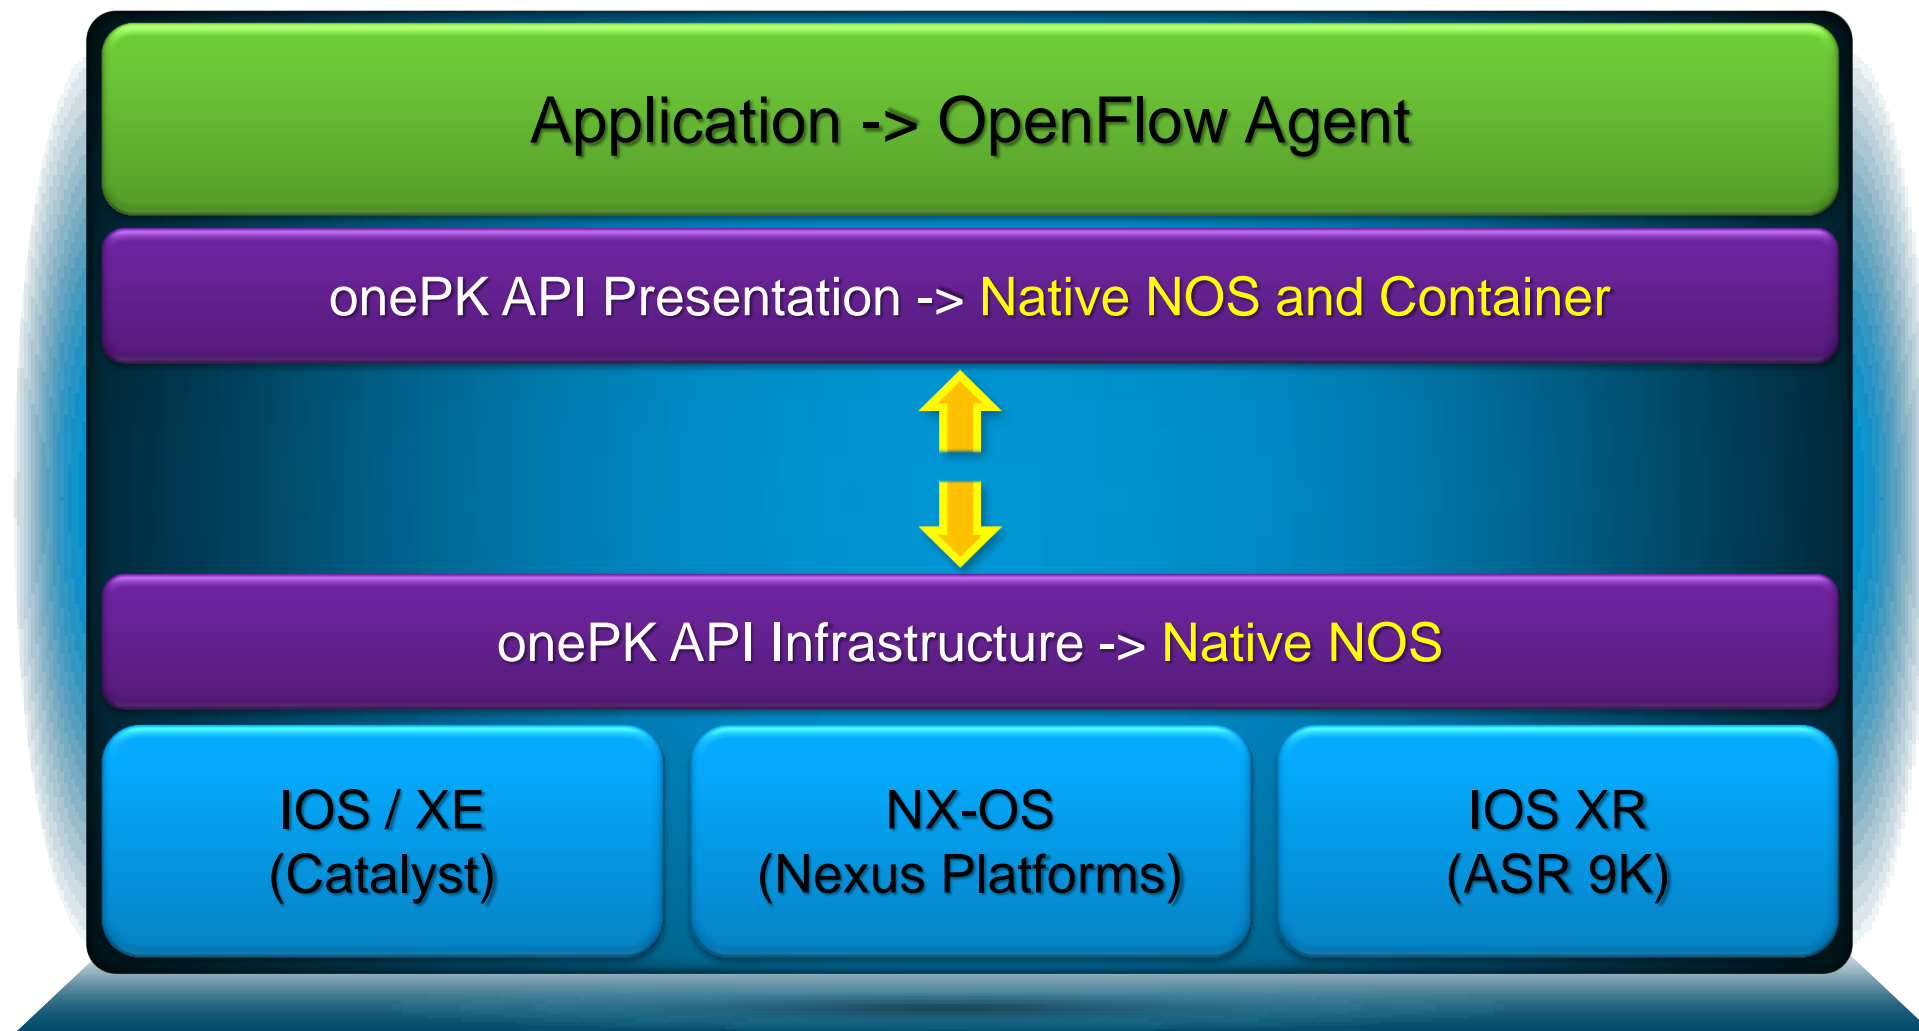
# OpenFlow @ Cisco



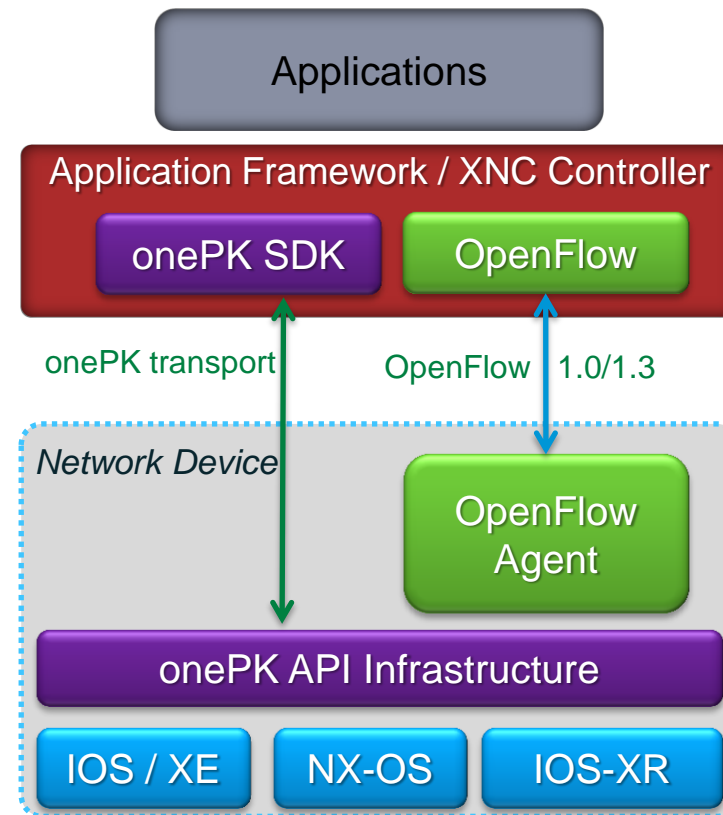

# OnePK Architecture



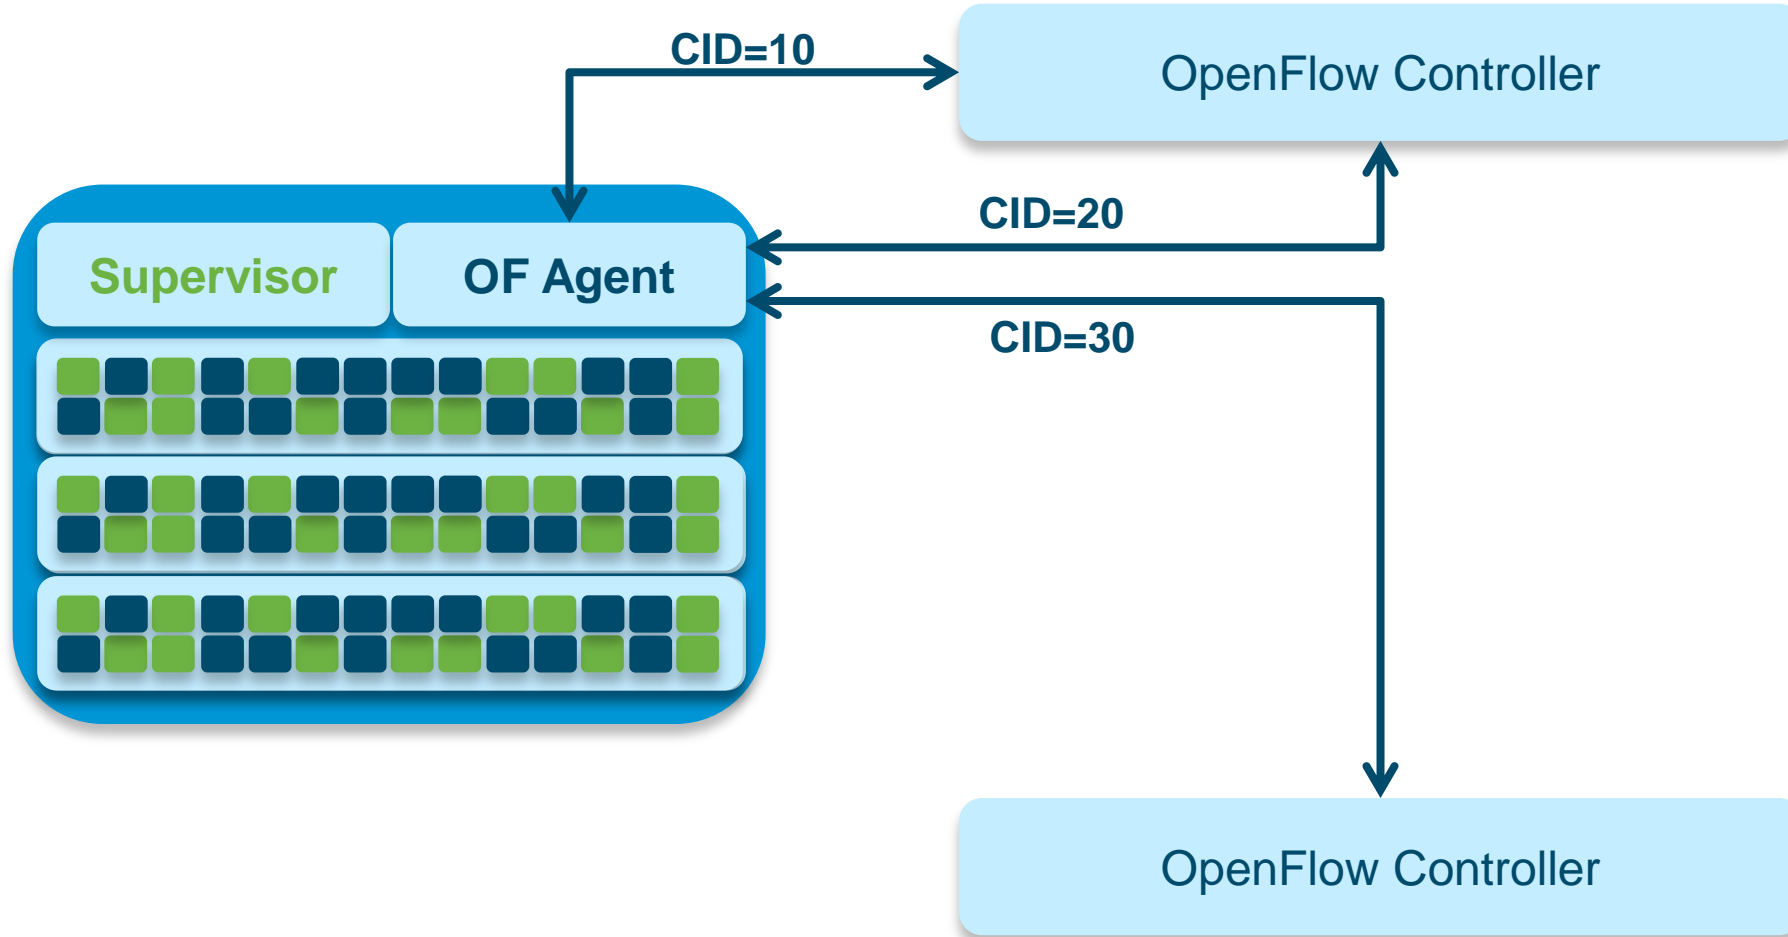
# OpenFlow Agent with onePK Architecture



# Application Flexibility



# OpenFlow Logical Switch(es)





# Hybrid OpenFlow on Cisco devices

## Pure OpenFlow\*

- All ports are OpenFlow only
- All forwarding decisions by Controller
- Example: Network Monitoring

\* There still may be some non OpenFlow ports for connecting to management device or initial setup

## Ships in the Night

- OF co-exists with normal forwarding
- Port segregation
- No traffic flow between domains
- Example: Network Slicing

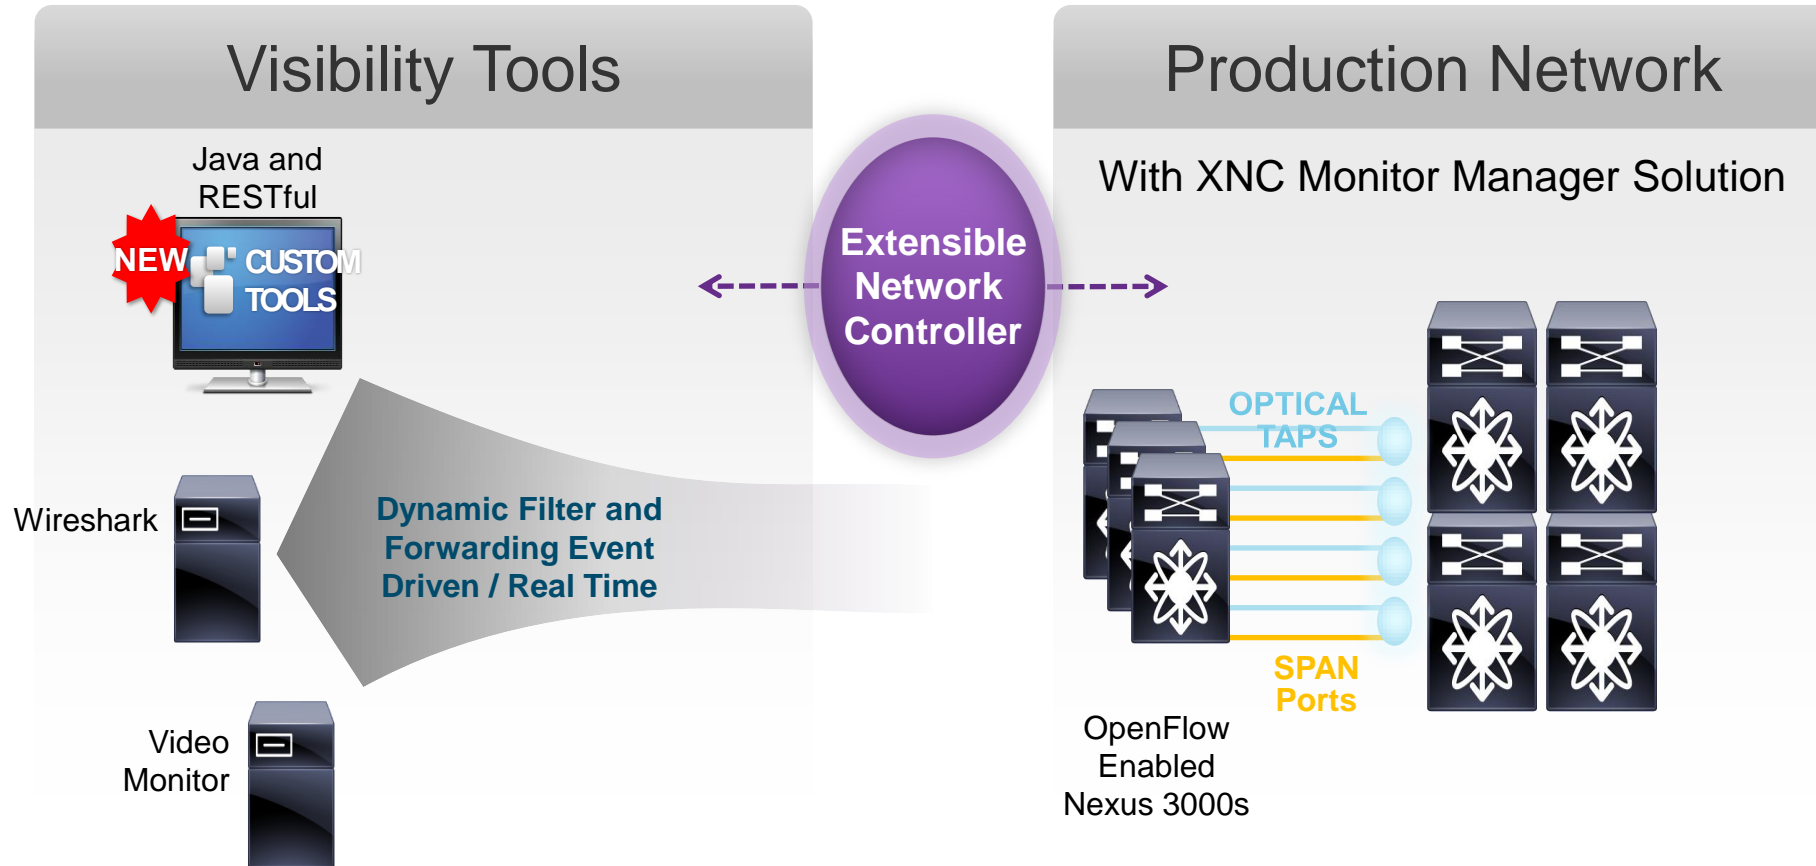
## Integrated

- Packet may traverse OpenFlow pipeline and normal pipeline
- Output to NORMAL and Logical ports
- Example: Traffic Steering

# Use Cases



# Use Case: Data Center Monitoring



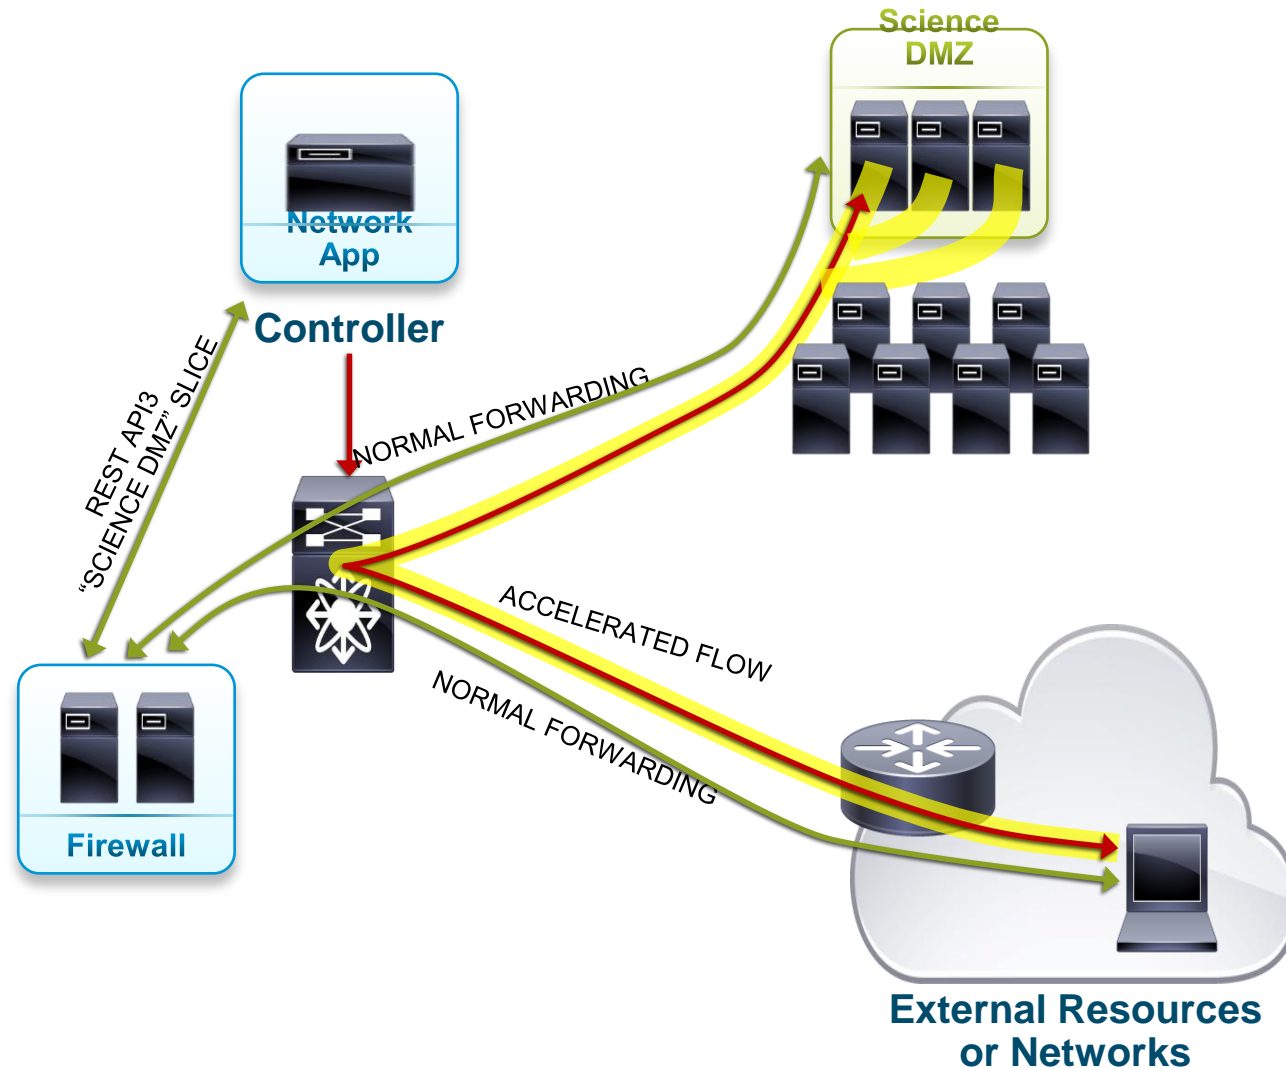
Replaces Matrix Network with  
Nexus 3000s, Controller, and Monitor Manager App

# Higher Ed/Enterprise: Service Insertion or Bypass

## Firewall Trusted Flow Acceleration

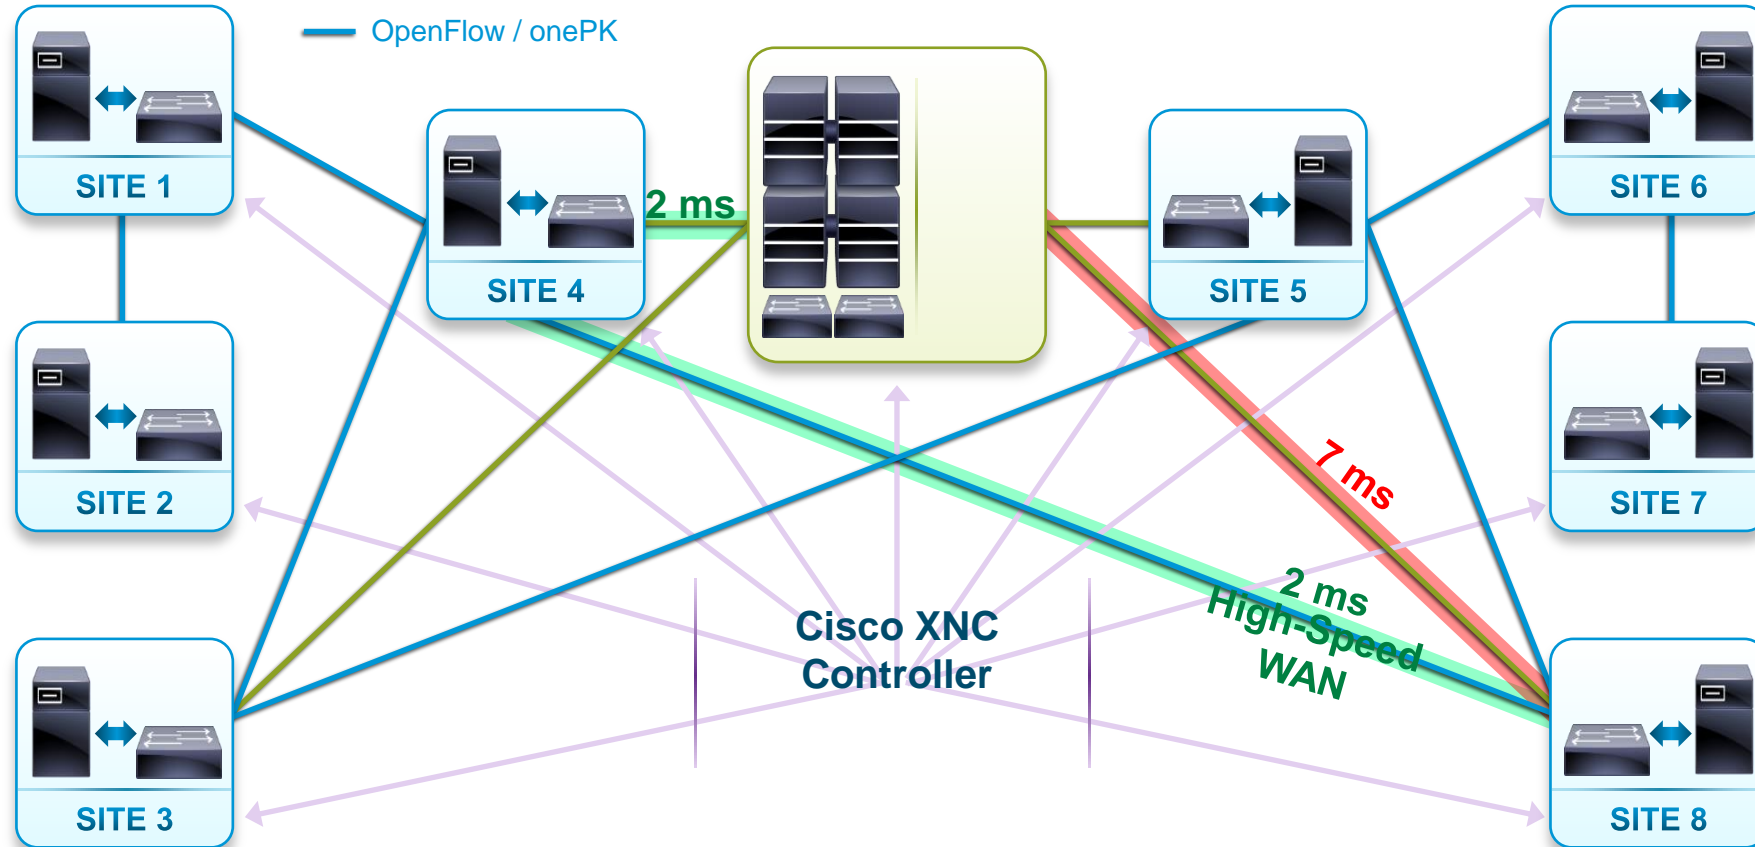
**Objective:**  
Establish a flow to bypass the firewall for trusted traffic to improve application performance

Accelerated Flows and "Science DMZ" Slice



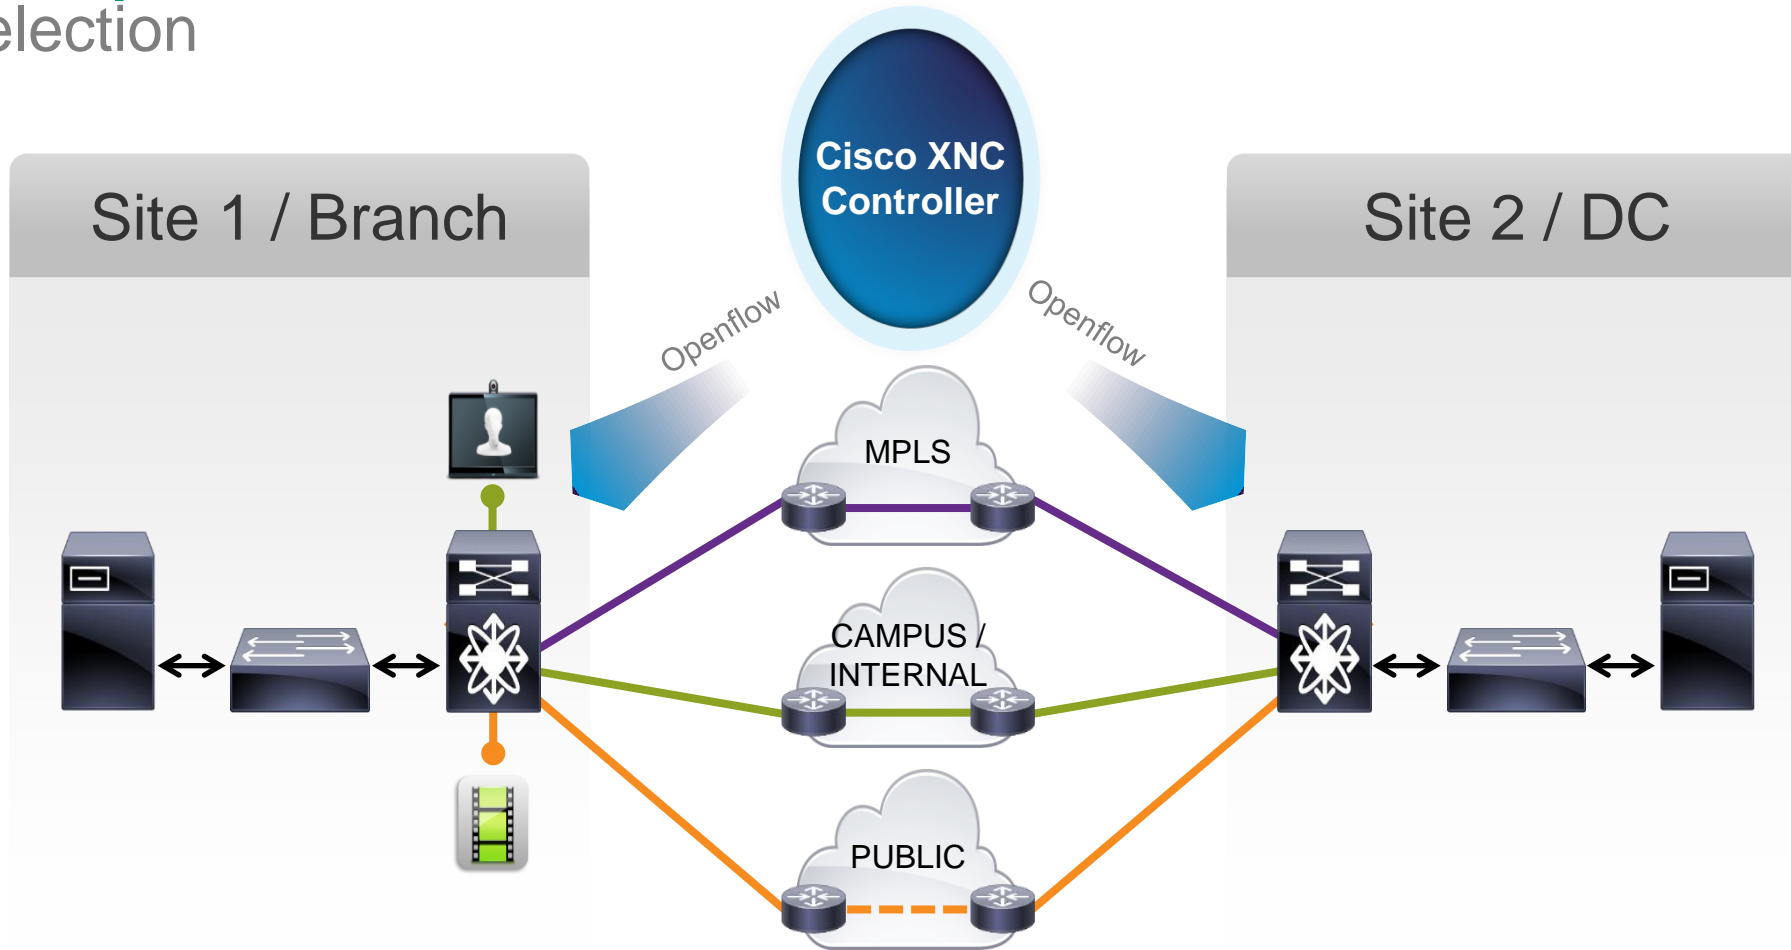
# Data Center Optimization

Transit Selection: Network Parameter Driven (Latency)



Other Transit Selection Parameters Also Possible

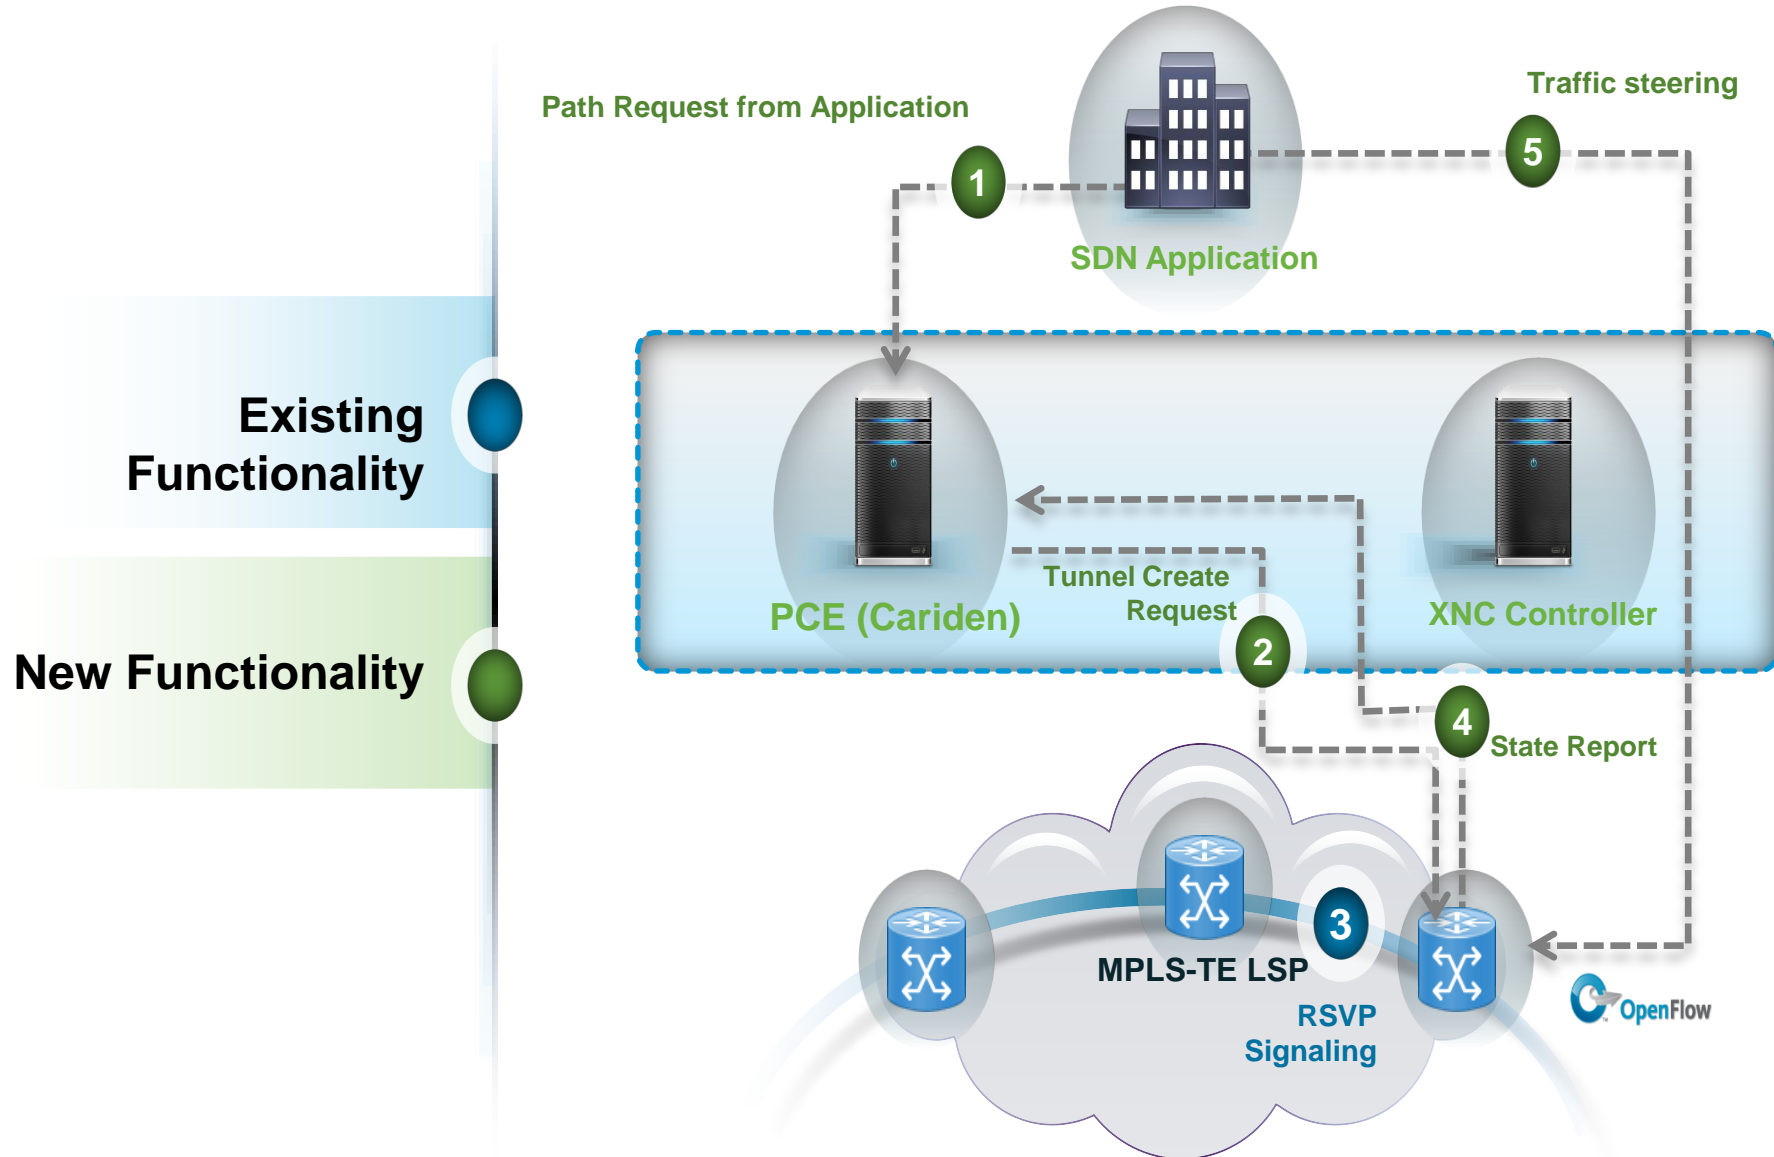
# SP/Enterprise: Transit Selection



Control egress path based on application requirement and priority  
Simple forwarding rules help improve WAN utilization

# Service Provider

## Traffic Engineering with PCE and OpenFlow



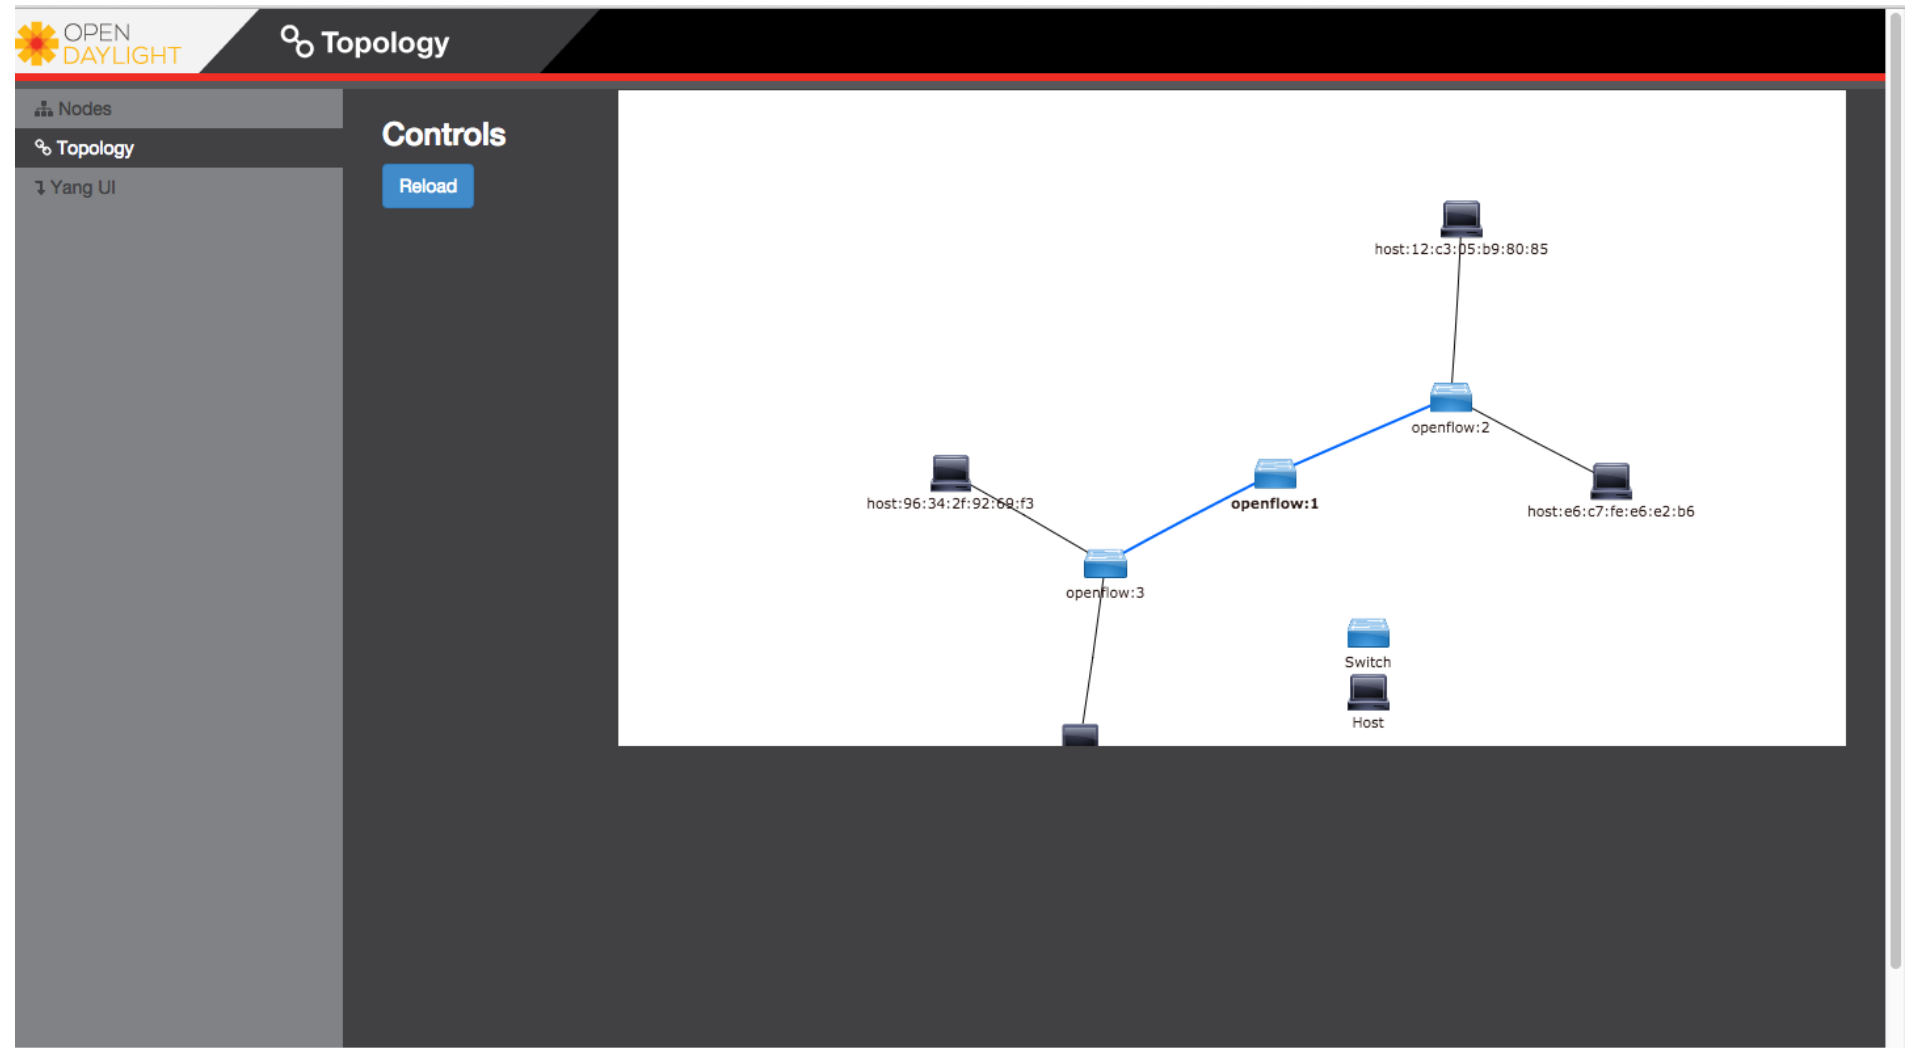
# OpenDaylight

- A Linux Foundation Project
- Industry-Supported





# OpenDayLight



# OpenFlow Packet

316 17.287798 10.75.45.138 10.75.45.137 OpenFlow 1186 Type: OFPT\_MULTIPART\_REPLY, OFPMP\_FLOW

▷ Frame 316: 1186 bytes on wire (9488 bits), 1186 bytes captured (9488 bits)

▷ Ethernet II, Src: Vmware\_b1:0e:85 (00:50:56:b1:0e:85), Dst: Vmware\_b1:6a:66 (00:50:56:b1:6a:66)

▷ Internet Protocol Version 4, Src: 10.75.45.138 (10.75.45.138), Dst: 10.75.45.137 (10.75.45.137)

▷ Transmission Control Protocol, Src Port: 33728 (33728), Dst Port: 6653 (6653), Seq: 76085, Ack: 60929, Len: 1120

▼ OpenFlow 1.3

- Version: 1.3 (0x04)
- Type: OFPT\_MULTIPART\_REPLY (19)
- Length: 1120
- Transaction ID: 1084
- Type: OFPMP\_FLOW (1)

▼ Match

- Type: OFPMT\_OXM (1)
- Length: 24
- ▼ OXM field
  - Class: OFPXMC\_OPENFLOW\_BASIC (0x8000)
  - 0000 100. = Field: OFPXMT\_OFB\_ETH\_SRC (4)
  - .... ..0 = Has mask: False
  - Length: 6
  - Value: c2:17:c7:ab:62:fe (c2:17:c7:ab:62:fe)
- ▼ OXM field
  - Class: OFPXMC\_OPENFLOW\_BASIC (0x8000)
  - 0000 011. = Field: OFPXMT\_OFB\_ETH\_DST (3)
  - .... ..0 = Has mask: False
  - Length: 6
  - Value: c2:83:78:31:d7:f8 (c2:83:78:31:d7:f8)

▼ Instruction

- Type: OFPIT\_APPLY\_ACTIONS (4)
- Length: 24
- Pad: 00000000
- ▼ Action
  - Type: OFPAT\_OUTPUT (0)
  - Length: 16
  - Port: 2
  - Max Length: OFPCML\_NO\_BUFFER (0xffff)
  - Pad: 000000000000

Thank you.

