

# SDN AND THE DATAPLANE



**CHI-NOG 3**  
**June 14<sup>th</sup>, 2014**

# So is the network really the problem?

- Elasticity and virtualization have moved the network square in the crosshairs as the delay of any deployment.
- Compute elements can be provisioned in minutes, networks still take time to provision.
- Application flows have become increasingly complex as migration of traffic from traditional north/south to east/west.
- Datacenters are no longer islands of traffic flows. Active/Active Datacenters are becoming more of the norm.
- The Trombone effect of network traffic from virtual to physical layers.
- Value added ( L4-L7 ) resources need to be placed either in path or virtually in path.

# Historical Approach

- Historically we designed the network with attempts at consistent throughput and latency. Applications had to live within the constraints of that design.
- Flows were mainly north/south.
- Vlans provided L2 isolation.
- Many attempts to mitigate spanning-tree from the network either with aggregation methods, control plane consolidation, or the use of proprietary methods.
- Oversubscription.
- Quality of Service
- Datacenter Interconnects were mainly for replication and disaster recovery.

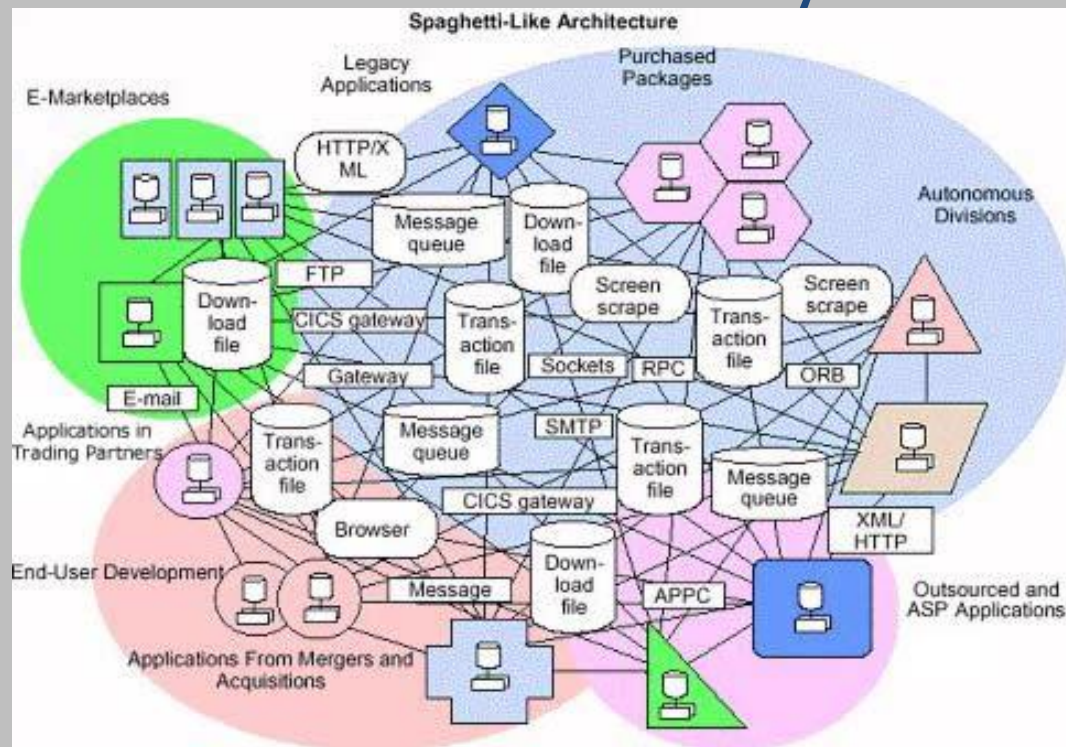
# Applications are driving things now.

## Business Expectations are Changing..


- Business want their environments to work more like AWS.
- L2 boundaries within a datacenter are roadblocks. VMs are mobile, elastic.
- Applications can be hosted on premises, at a colo, or in the cloud. Application flow can cross all boundaries.
- Application Rollout Time
  - Pre-Virtualization: Months
  - Post-Virtualization: Days
  - Converged Infrastructure: Hours
  - SDN: Minutes ?

# Applications

- In the end its about the applications
- Distributed Complex flows and then they blame our networks? Really.



# Traditional North/South Design



# The Dataplane

- As with any emerging standard manufacturers are lining up to solve this in their own way.
- Some adhere to new standards, some build their own, some are hedging their bets.
- We will focus on two Dataplane methodologies.
  - Fabric Programming and Overlay Networks
  - Their impact to your network

# The Fabric





# Well its not an SDN Preso without OpenFlow


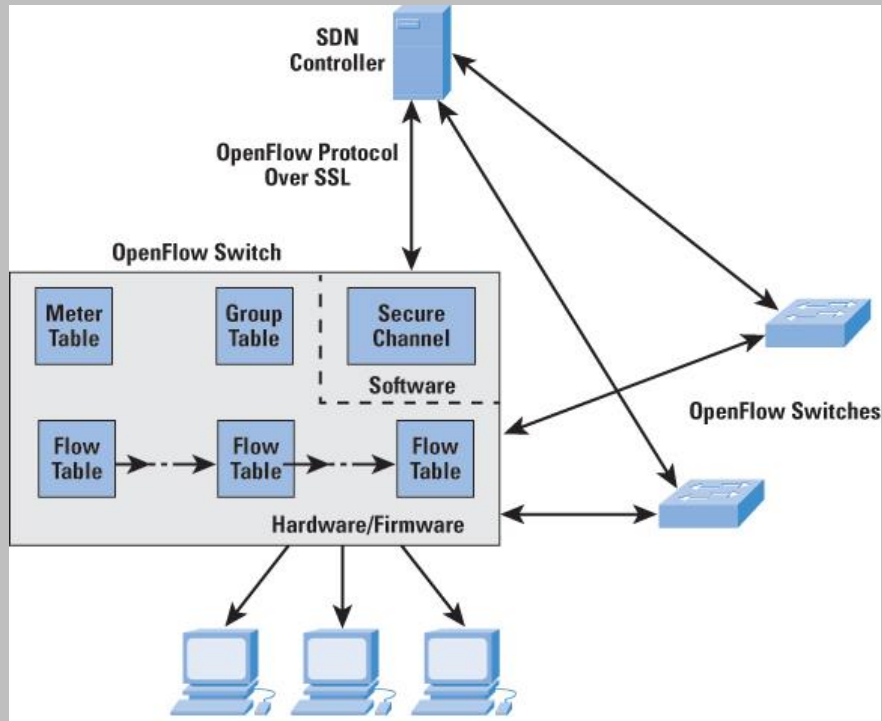
- SDN does not equal Openflow.
- Starts to address FIB Programming from a central control plane.
- Openflow is not the only way to program the dataplane on devices.
- Scale is still an issue.
- Hybrid Openflow allows for the programmability of flows with keeping local forwarding capability. Ships in the Night model.
- Openflow is not a true policy construct. Other protocols solve policy deployment.
- Pure Openflow environments are still a ways off.

# Openflow 1.3

## Updates from 1.0

- Openflow 1.0/1.1/1.2 Features
- Support for Extended IPv6 headers
- Per-flow metering
- QoS and Bandwidth Management
- Provider Backbone Bridging tagging
- Tunnel-ID metadata ( specify encap and decap of packets )
- Success Stories of Openflow. Marist College with their DWDM on demand network.

# OpenFlow



Switch-1 Flow Table	
Match	Action
Src.Port=10	Fwd to Port 20
Src.Port=20	Fwd to Port 10

# So lets talk about Tunnels Overlay



# Revolutionary vs evolutionary

- Tunnels are nothing new.
- Good/Bad we have stitched segments together with MPLS/GRE/IP
- Sometimes these tunnels were used to hide design flaws.
- Sometimes these were used to get around controls.
- Tunnels add complexity instead of reducing complexity. ( Visibility/Troubleshooting )

# Tunnels with regards to SDN is a means to an end

- Most traffic is east/west based.
- L2 stretch is not always available.
- Everyone hates Spanning Tree.
- Tunnels allow to abstract the underlay network transport and to stitch together layer 2 networks.
- The end is to interconnect applications and vms.
- To simplify the inter vm communication.
- And more to the point, SO WE CAN VMOTION!!!!!!

# Overlay versus Underlay

- Don't get wrapped up in Overlay tunnel technology. Tunnels are Tunnels.
- Overlay can abstract visibility into underlay network conditions.
- Overlays do not mean that the underlay is less complex.
- Underlay network should allow for all paths and should be non-blocking.
- Long flows versus short flows still an issue with Overlay.

# Tunnels are not Magic

- Underlay network can still be an issue.
- Physics still applies. So earth is round, Einstein is still correct about the speed of light. Latency still is an issue.
- Lack of visibility between the Overlay and Underlay networks. Tools need to marry analytics to help make better decisions on traffic.
- Traffic Engineering is still needed.



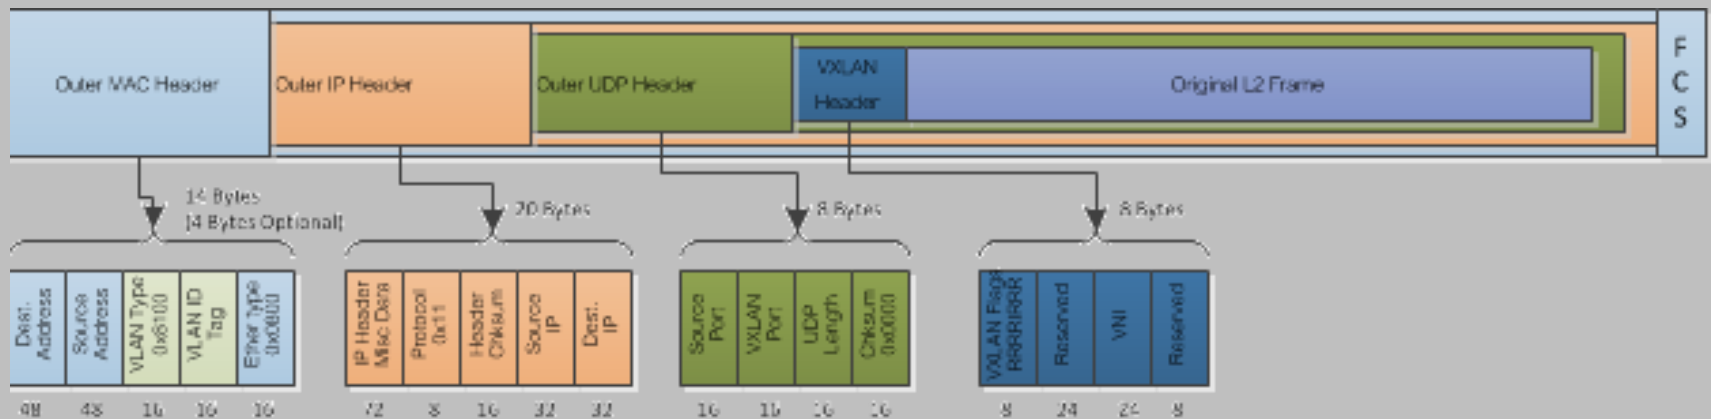
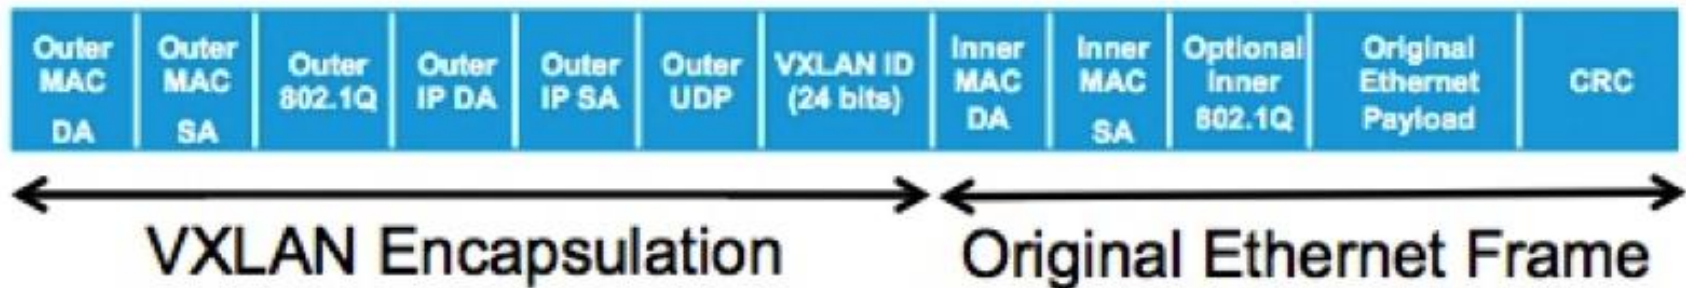
# Several Tunneling Technologies

- NVGRE ( Microsoft )
- STT (Stateless Transport Tunneling )
- VxLAN ( VMWare)
- GENEVE ( Draft ) Microsoft/VMWare

**VXLAN GIVES US**

**BILLIONS AND BILLIONS OF  
VLANS**

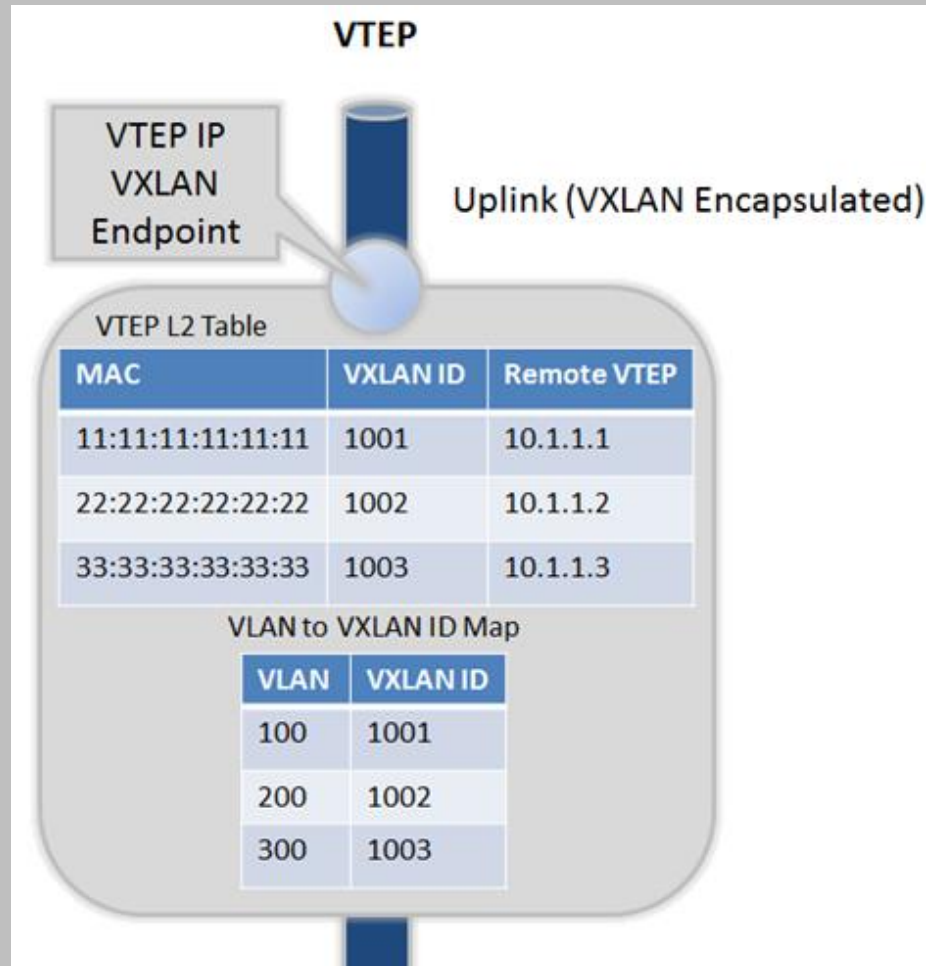
# VXLAN




# VXLAN Integration with Traditional DC

- Even the most ardent SDN advocate realizes that integration will take time.
- Islands of SDN will appear in your DC.
- So how to we provide connectivity between the old and the new.
- Hardware VTEPs are the key. Hardware manufactures are using features in merchant silicon to terminate VTEPs. Make sure for scale that this is an ASIC process and not run solely in software.
- We then map these terminated VXLAN flows with traditional vlans and L3 routing.
- Hardware VTEPs might be needed beyond migration as certain processes still scale better with hardware.

# Hardware VTEP



# Islands of SDN



# Abstraction with Tunnels

- Traditionally we have used DWDM/L2/VPLS to provide Data Center Interconnect.
- Tunnels abstract the underlay network and allow us to not worry about L2.
- Another opportunity to rip out more L2 in the DC.


# Clos: The underlay

- Can use L3 Clos to avoid running spanning tree.
- Remember all of those magical tunnels, they abstract the underlay.
- Head of Line blocking not as much of an issue.
- Multistage Clos fabrics can allow for scale.
- Pod architecture to help scale in your datacenter.




# Underlay Possibilities

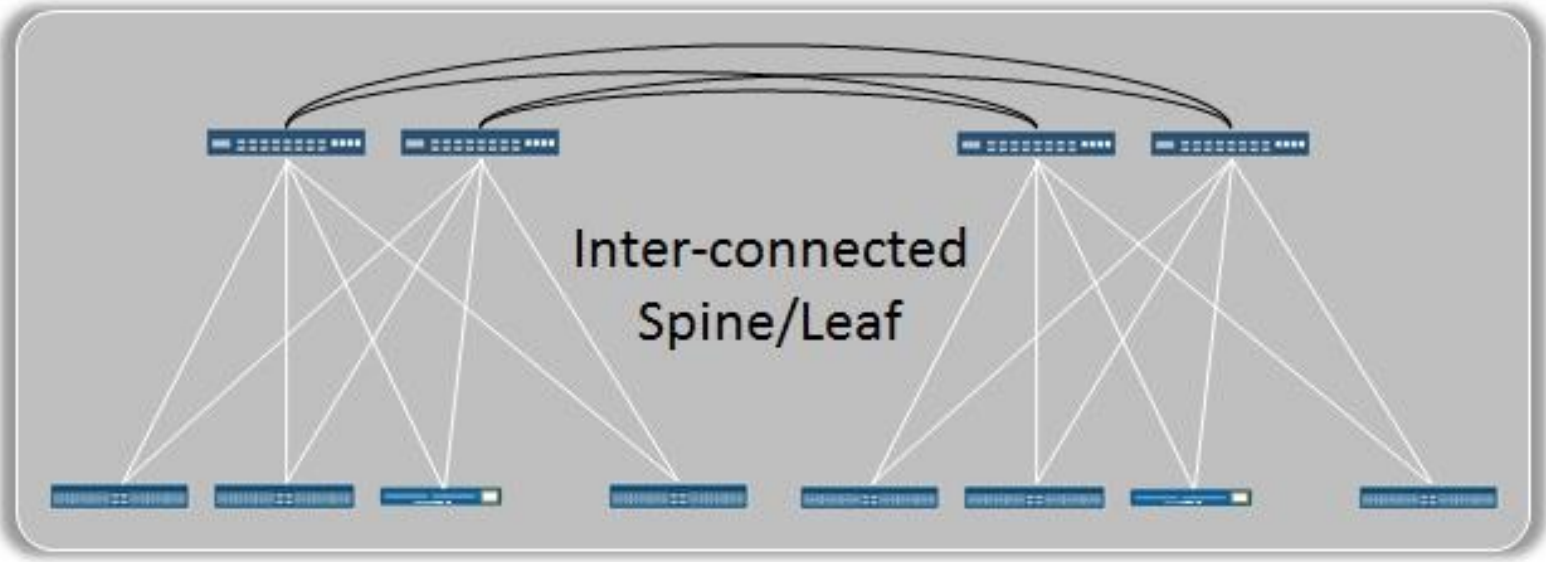
Full Mesh



Spine Leaf




Inter-connected  
Spine/Leaf



# Stitching the Virtual Flow together

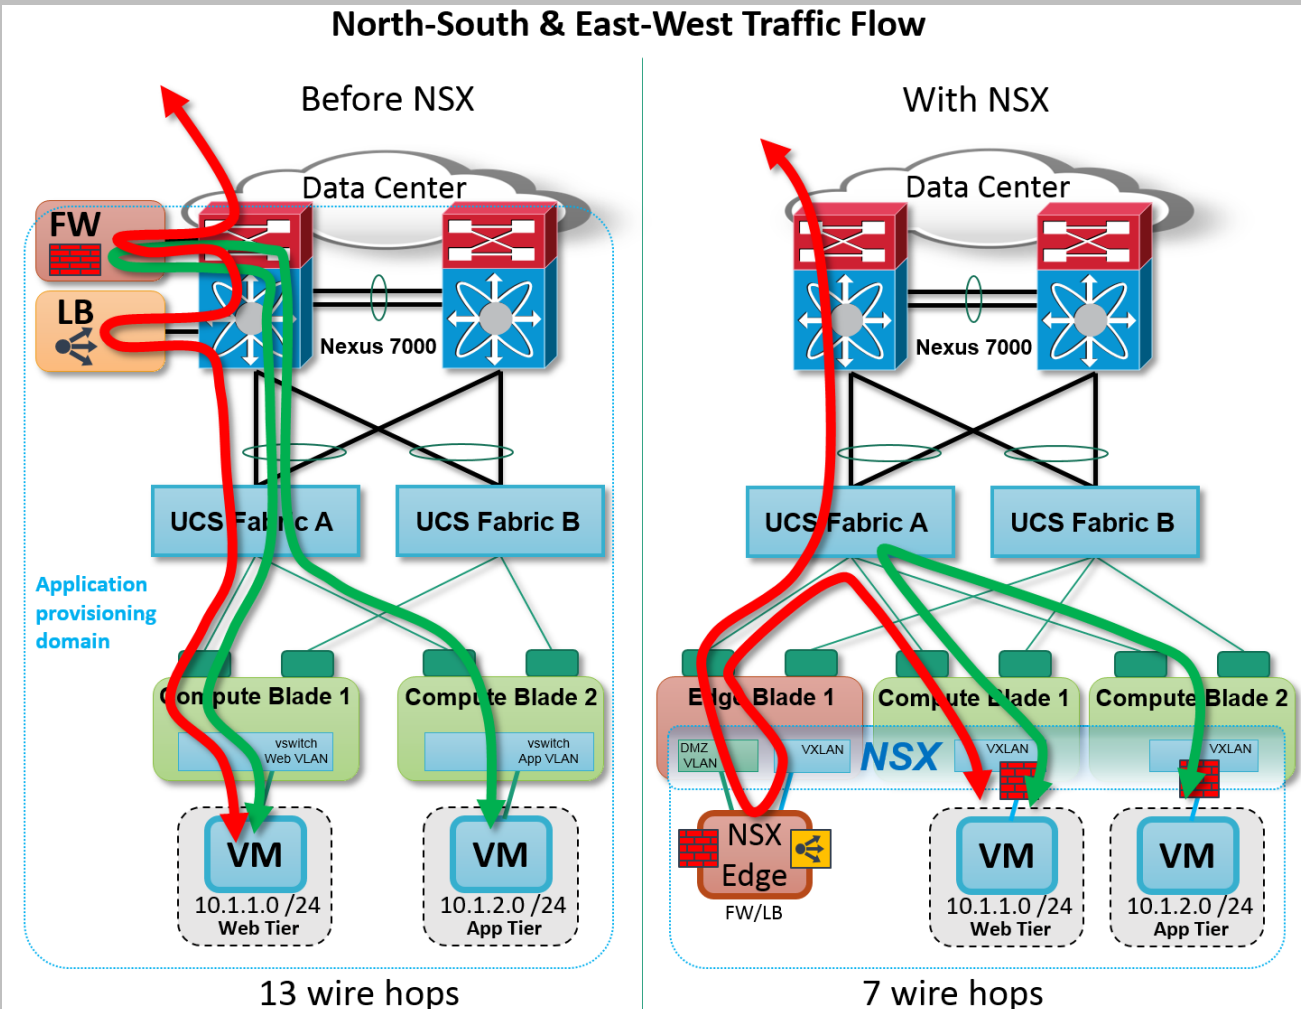
## SDN Service Chain Example 1

Datacenter Cloud Application  
Internet to Web Server



# Efficient Flows

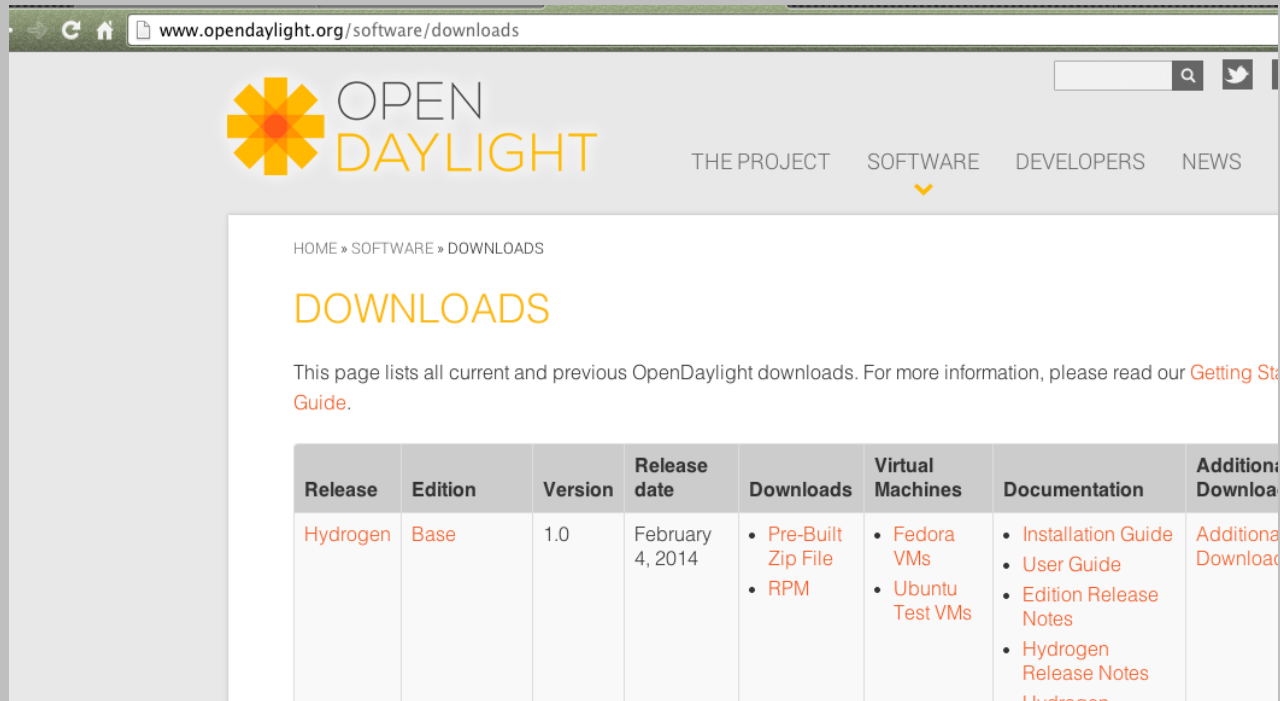
## North-South & East-West Traffic Flow



# How do I get started?

- Start with Openflow. Build a lab.
- Can run it as a self contained VM with Opendaylight and a virtual switch.
- Education: Free online education for Opendaylight, Openflow, and SDN.
- [SDNCentral.com](http://SDNCentral.com)

# Opendaylight download



The screenshot shows the OpenDaylight website's download page. The browser address bar displays 'www.opendaylight.org/software/downloads'. The page features the OpenDaylight logo (a stylized orange flower) and the text 'OPEN DAYLIGHT'. A navigation menu includes 'THE PROJECT', 'SOFTWARE' (which is highlighted with a yellow arrow), 'DEVELOPERS', and 'NEWS'. Below the navigation, a breadcrumb trail reads 'HOME » SOFTWARE » DOWNLOADS'. The main heading is 'DOWNLOADS' in large orange letters. A paragraph states: 'This page lists all current and previous OpenDaylight downloads. For more information, please read our [Getting Started Guide](#).' Below this is a table with the following data:

Release	Edition	Version	Release date	Downloads	Virtual Machines	Documentation	Additional Downloads
Hydrogen	Base	1.0	February 4, 2014	<ul style="list-style-type: none"><li>• <a href="#">Pre-Built Zip File</a></li><li>• <a href="#">RPM</a></li></ul>	<ul style="list-style-type: none"><li>• <a href="#">Fedora VMs</a></li><li>• <a href="#">Ubuntu Test VMs</a></li></ul>	<ul style="list-style-type: none"><li>• <a href="#">Installation Guide</a></li><li>• <a href="#">User Guide</a></li><li>• <a href="#">Edition Release Notes</a></li><li>• <a href="#">Hydrogen Release Notes</a></li><li>• <a href="#">Hydrogen</a></li></ul>	<a href="#">Additional Downloads</a>

# Questions?

