

# ESTRATÉGIAS DE SDN PARA DESAFIOS REAIS DE CLOUD COMPUTING

Marcelo Molinari  
mmolinar@brocade.com  
SE Manager – Brasil and SSA



# Software-defined Networking in the Data Center

## ENABLING NEW LEVELS OF INNOVATION

### NETWORK VIRTUALIZATION



### PROGRAMMATIC CONTROL

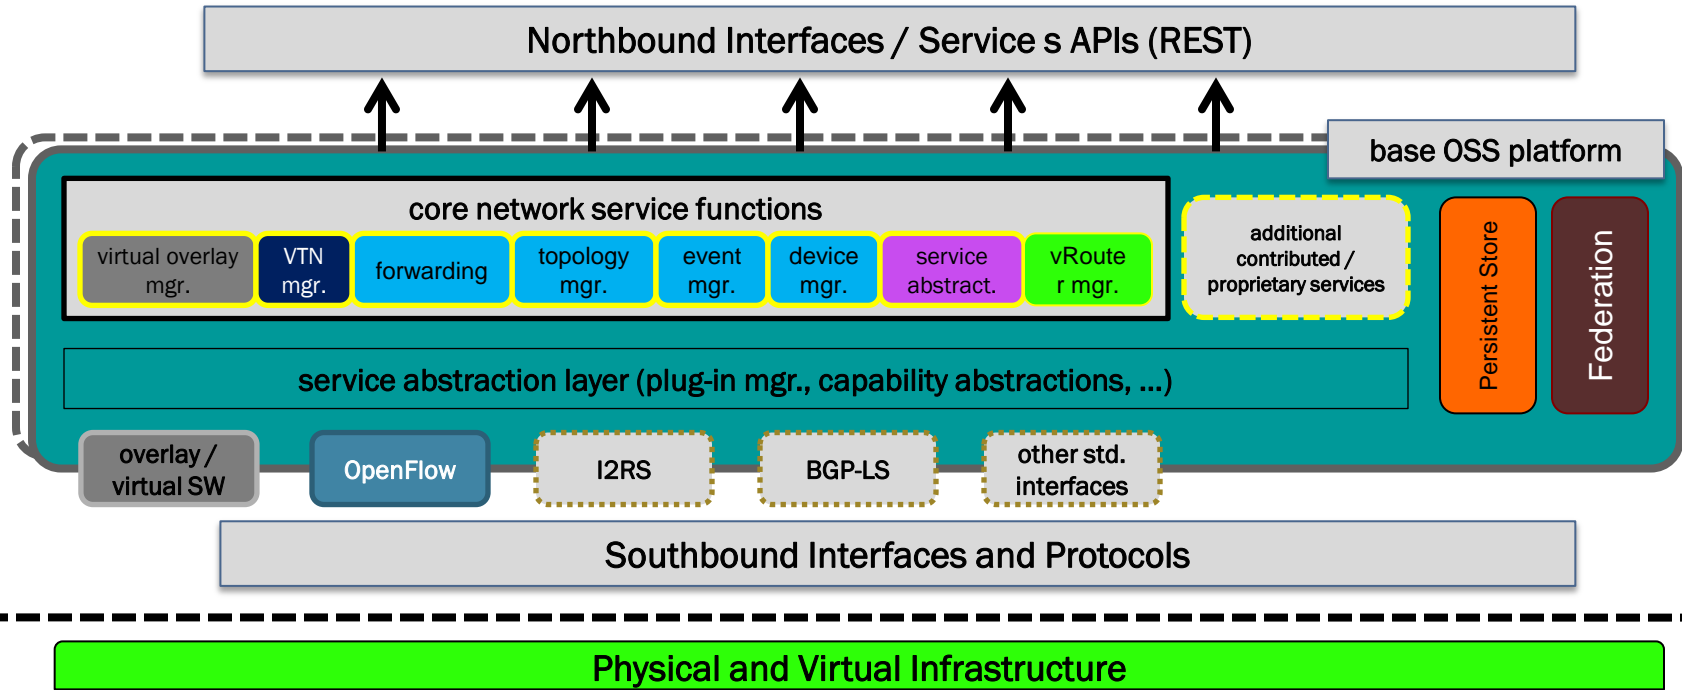


### CLOUD ORCHESTRATION



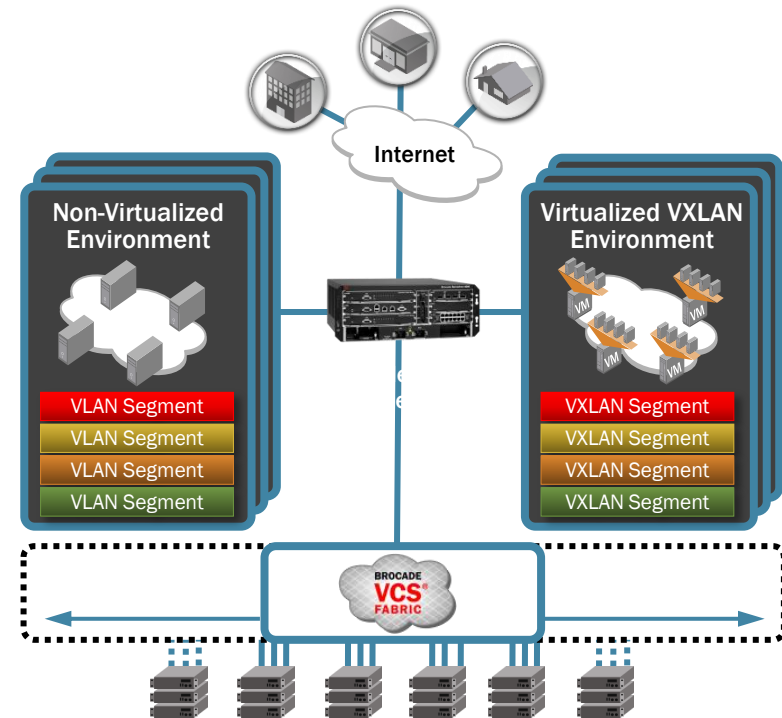
# Brocade Network Platform Will Enable Rich Network Services

FOUNDING MEMBER OF OPENDAYLIGHT OPEN SOURCE SDN CONSORTIUM

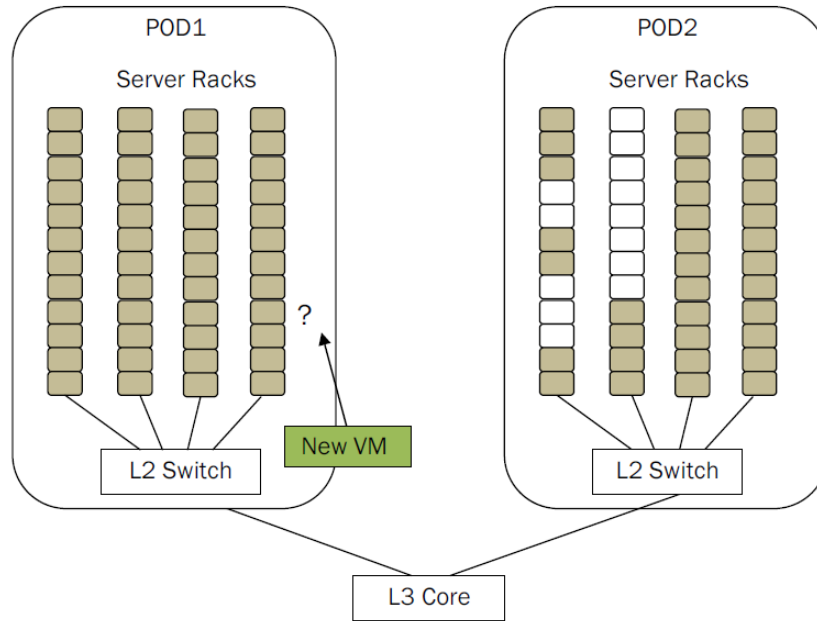


# Network Virtualization—Logical Networks

- Logical networks eliminate physical network limitations (e.g., MAC address and VLAN limits) to facilitate any-to-any connectivity and workload mobility
- Overlay/tunnel technologies: VXLAN
- Brocade VCS Ethernet fabric is a superior transport for network virtualization—leading automation, efficiency and simplicity
- Brocade VDX data center switch ASICs are VXLAN-ready with software support in 2013; Brocade ADX VXLAN gateway demo now

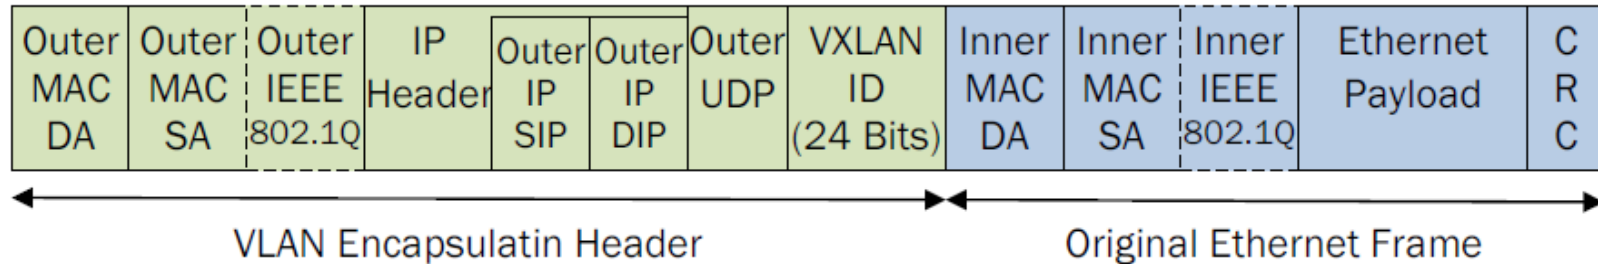


# Limitations of Current Network Virtualization



- POD1 servers are running at full capacity while POD2 has ample idle capacity.
- A new VM has to be instantiated. Due to limitations of the network virtualization, the new VM can be instantiated only in POD1.
- This may be due to reasons such as:
  - VMnew is being spun-up due to increased demand in a vApp. The vApp resides in POD1 and this requires that VMnew also be spun-up in POD1.
  - VMnew has an IP address in one of subnets (say C1) for which the router directs the traffic to POD1.

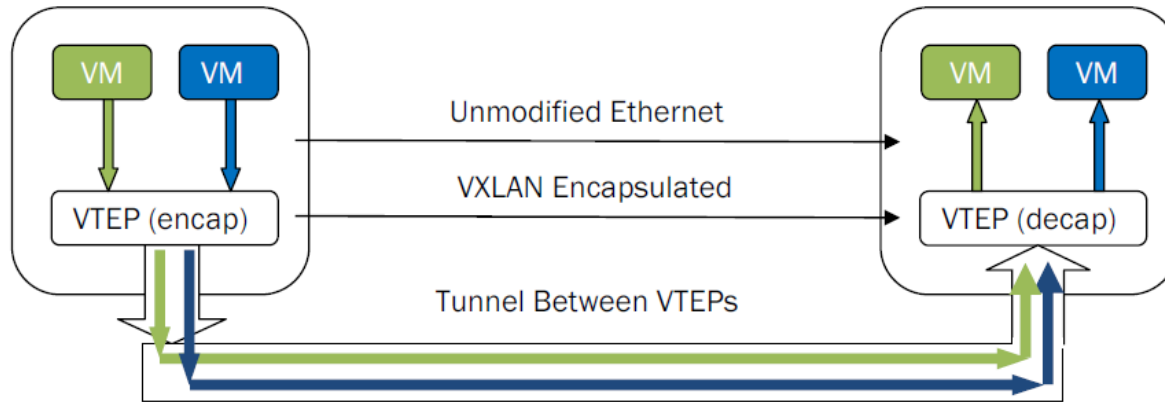
# Overcoming the limitations with VXLAN



- VXLAN overcomes the limitations of VLAN based network virtualization through the use of tunneling technology.
- Tunneling technology has been well proven in the WAN environment as a means to stretch Layer 2 Networks (VPLS) and Layer 3 Networks (GRE) across data centers.
- VXLAN uses a VXLAN/UDP/IP header to encapsulate the original unchanged Ethernet packet.



# VXLAN Tunnels



- The VXLAN encapsulation/de-capsulation, also referred to as VLAN tunnel termination, is performed by an entity in the Hypervisor known as the VXLAN Tunnel End-Point.
- This slide shows two tenants, the green tenant and the blue tenant, with VMs on the two Hypervisors.
- VXLAN provides complete isolation between the two tenants, enabling them to have overlapping MAC and IP addresses, and even overlapping VLAN tags.

# Network Virtualization—vRouting

## BROCADE VYATTA TECHNOLOGY

- Brocade acquired Vyatta, who pioneered the Virtual Router category and a software networking platform
- 1M plus downloads globally and 250,000 Active User Community
- vRouting technology to address east-west traffic in a highly virtualized and scale-out, multi-tenant environment
- An enterprise branch router platform for managed services





# Vyatta Delivers L3 Routing with Integrated Security


## RUN AS A VM OR STANDALONE SOFTWARE

### Router



OSPF, BGP

### Firewall



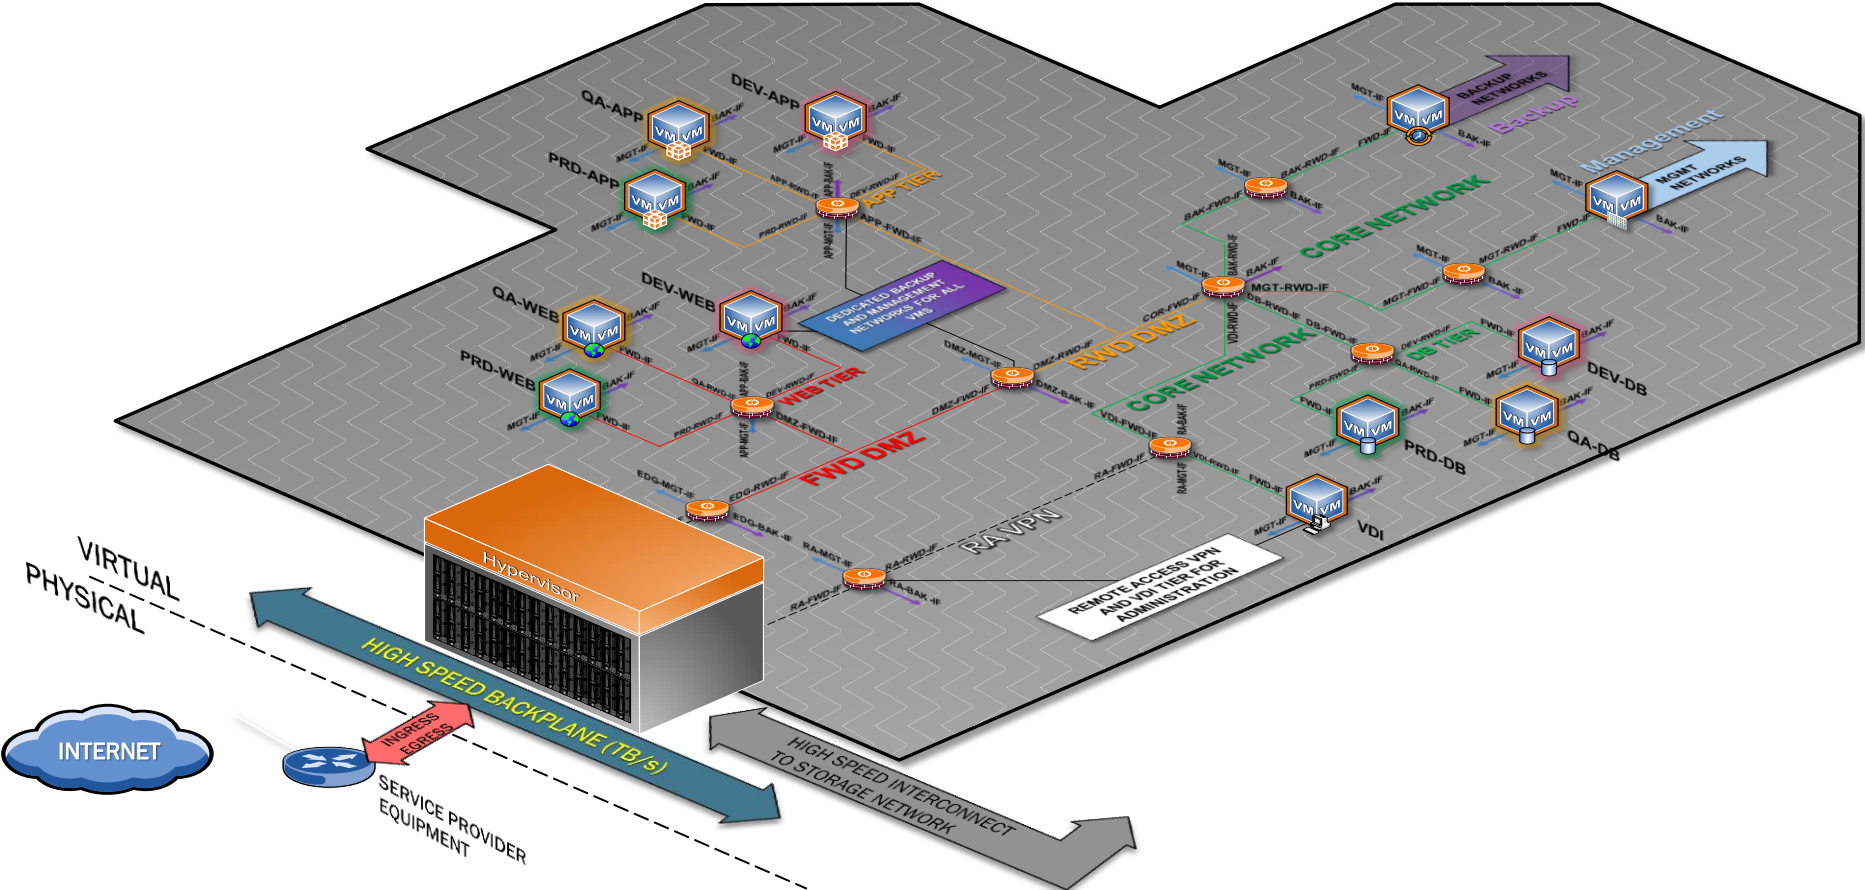
Stateful, NAT

### VPN



IPSec, SSL


# Data Center : Software Template Model



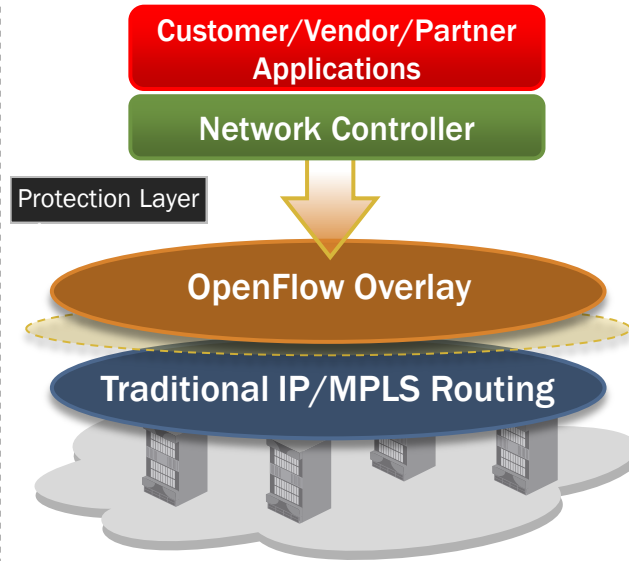
# OpenFlow—Programmatic Network Control



## A NEW POINT OF INNOVATION FOR NETWORK FUNCTIONALITY



**Traditional Network**



**OpenFlow-enabled Network**

- OpenFlow 1.0 shipping on Brocade MLXe router up to 100 GbE
- OpenFlow overlay runs in “hybrid mode” - concurrently with traditional IP/MPLS routing
- Benefits
  - Seamless interoperability with traditional networking
  - Separate slice of network for OpenFlow innovation
  - Isolation enforced in hardware

# OpenFlow Protocol Building Blocks

- Explicitly defined by OF Standard
  - Flow(s) - how network traffic is defined
  - Flow Table(s) - list of Flows used for processing network traffic
  - Open Flow Channel - communication protocol used between controller and network element.
- Pieces need to make an OF solution operate.
  - Application(s) - Software that determines network forwarding behavior
  - Controller(s) - Software that communicates with network element and application
  - Network Element(s) - Switches/routers (physical or virtual)



# OpenFlow Building Blocks

## The Flow

OF Flow



- Basic building block in OpenFlow is the Flow.
- A Flow represents the matching fields, actions and corresponding statistics.
  - Matching Fields (OF v1.0) include ingress port and L2/L3 packet header.
  - Actions include drop, forward to port(s), forward to controller, modify fields (and then forward).
  - The Flow also includes statistical information (counters).

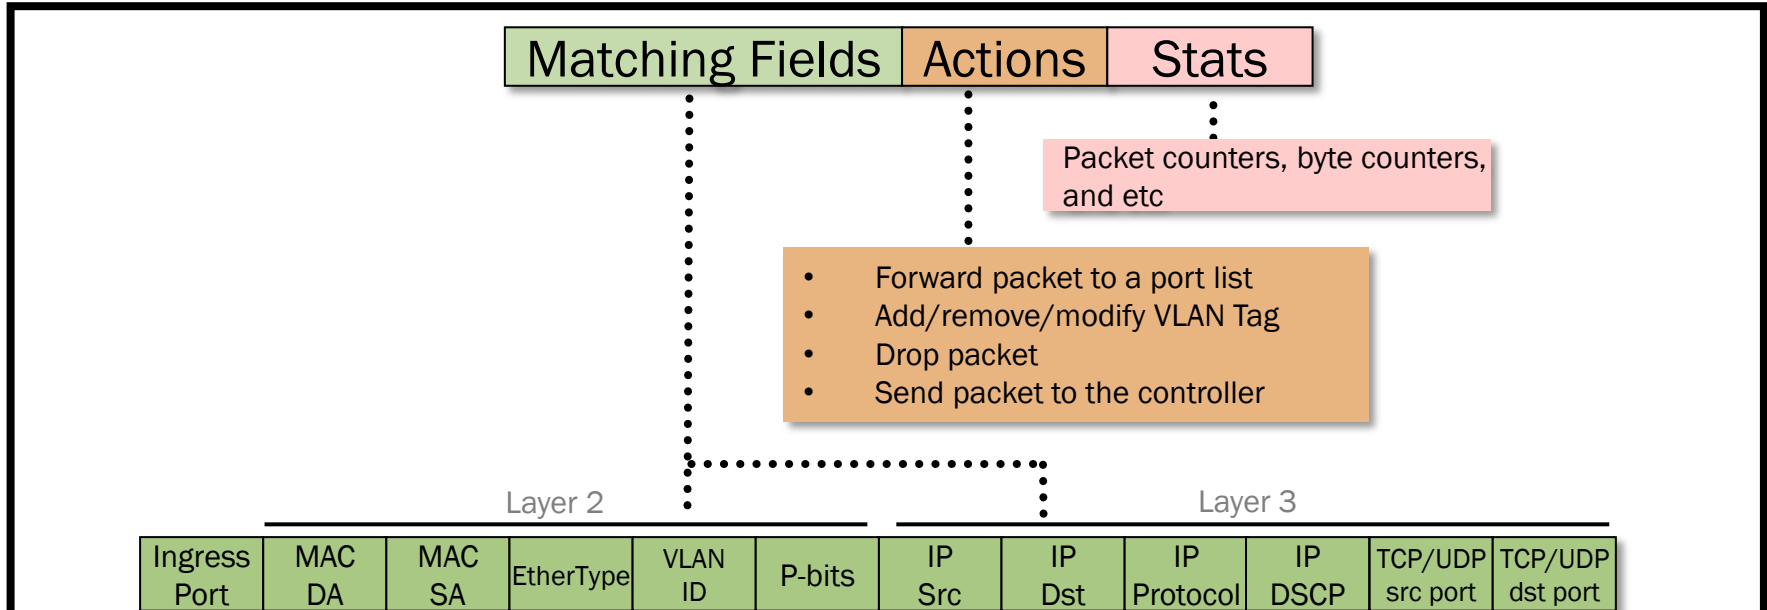


# OpenFlow Building Blocks

## Flow – OpenFlow v1.0

While the OF specification may support a particular match field (or combinations) and action (s), the underlying hardware may not. Always check hardware support for the spec.

### Flow Entry



- Each flow table entry contains a set of rules to match (e.g., IP src) and an action list to be executed in case of a match (e.g., forward to port list).



# OpenFlow Building Blocks

## Flow Table

- A flow table is the “blue print” that a switch uses to process packets through the data plane
- At a high level it operates very much like an access control list. The table contains flow entries and is used by the switch to process packets.
- Flow entries are ordered by priority

Flow Table


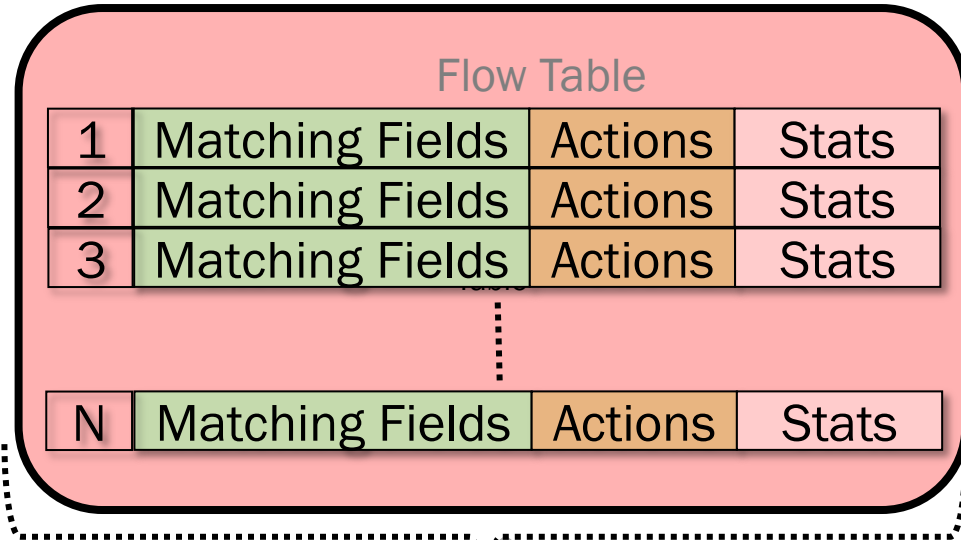
1	Matching Fields	Actions	Stats
2	Matching Fields	Actions	Stats
3	Matching Fields	Actions	Stats
⋮			
N	Matching Fields	Actions	Stats



# OpenFlow Building Blocks

## The Flow Table in action

- Packet Ingress to switch.
- By priority (order of entries in table), compare Packet Header to flow table.
- When match found, perform action (drop, forward out port(s), forward to controller) and update stats.
- If no entry is found, default behavior is to drop packet.





# OpenFlow Architecture

## Switch, Controller, Application

### Network Element (Switch)

- Communicates with Controller via the OpenFlow Channel (to/from)
- Forwards packets based on Flow Table\*
- *Note – can also operate in Hybrid Mode etc.*

### Controller

- Provides mechanism for application to push flow information to/from switches.
- Typically Server Based
- May include tool kit to write scripts (simple apps)

### Application

- Determines packet forwarding behavior and pushes the information (flow table) through controller to switch.
- Custom network behavior



# OpenFlow Architecture

## Application

Determine Network Forwarding Behavior Based on inputs from users, other software packages and information received from networking hardware

Application

```
.....  
TABLEINIT();//Forward all to flows to controller  
FOR(){  
FLOW = LISTENTOCONTROLLER();//Listen for new flows  
ADDFLOW(FLOW);//add source mac rule with port.  
PUSHFLOW(FLOWTABLE);//push flow to controller  
.....
```

NOT REAL Code.  
Simply illustrating that  
network configuration is  
now "programmatic"

Controller

## Controller and Switch

Communicate flow entry information to/from switch via OpenFlow Channel (Protocol)

## Switch

Forward Traffic based on entries in flow table (to controller, drop, out port etc).

Flow Table

Flow Table

Flow Table

Highest Priority

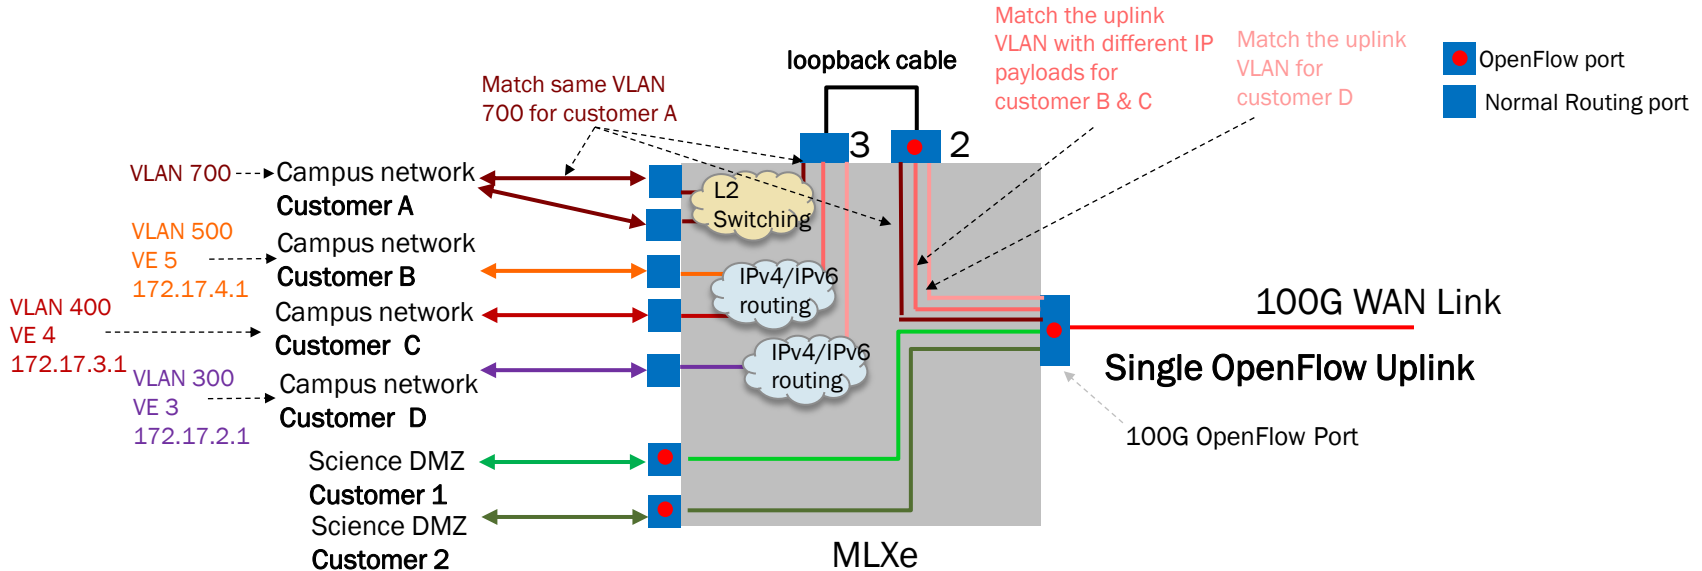
Lowest Priority

Flow Entries



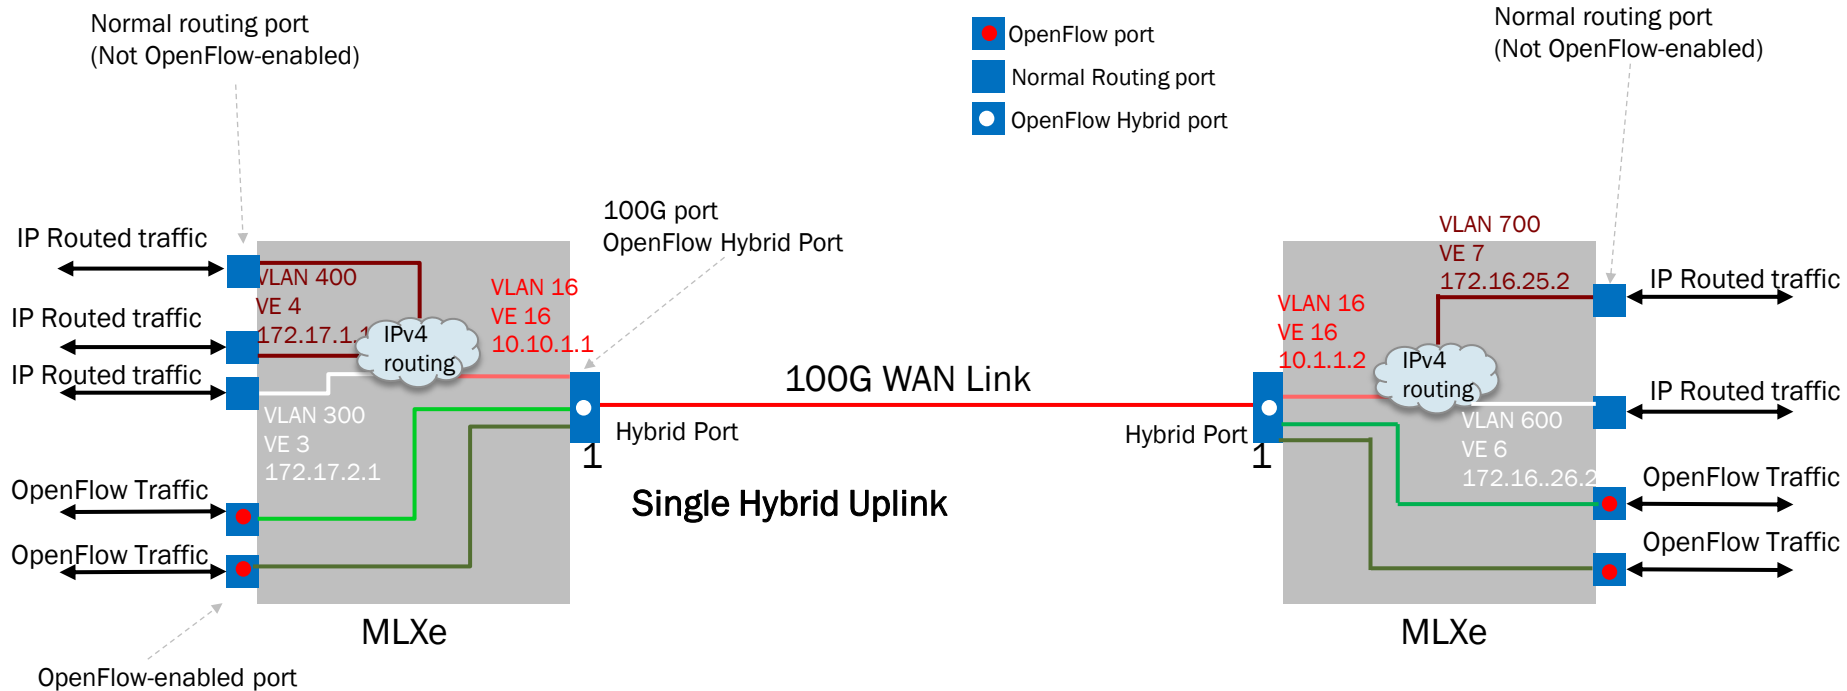
# Hybrid Switch Mode – What all vendors do

## Best Practice for using “Single OpenFlow Uplink” port



# Hybrid Port Mode – What Brocade does

## Best Practice for using “Single Hybrid Uplink” port



# Brocade MLXe High-Performance Core Router

## INTERCONNECTING BROCADE VCS FABRICS IN THE DATA CENTER CORE

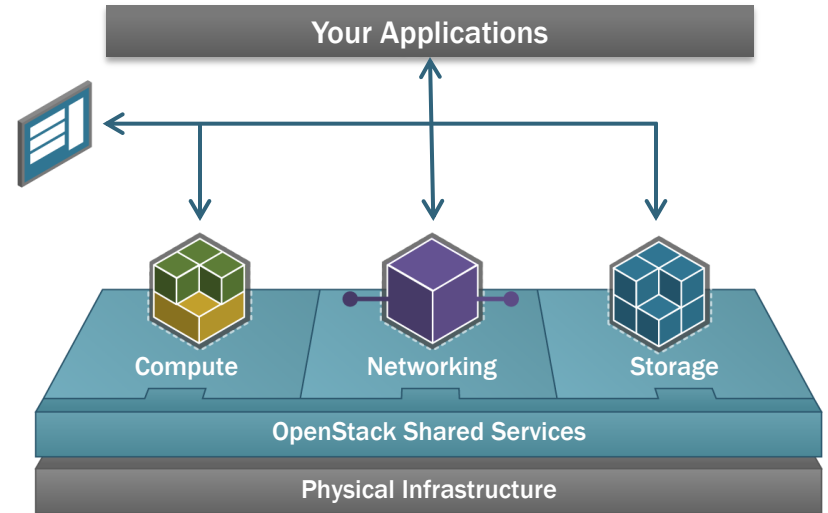


- **Massive scalability:** Industry-leading 10G and 100G router density
  - Supports up to 15.36 Tbps fabric capacity, 768 10G ports, and 32 100G ports in single chassis
  - On-demand capacity increase using high-capacity link aggregation
- **Advanced, scalable software features**
  - Hardware support for up to 1M MPLS labels
  - Routing over VPLS enables VM mobility and service flexibility between data centers
  - Multi-tenancy at scale with VRFs, VLANs, and QinQ
- **Compelling economics**
  - Dramatically consolidates network devices
  - Efficient, green design enables power and space savings
- **Seamless migration to SDN:** Industry's first true Hybrid mode OpenFlow
- **High availability:** Resilient Multi-Chassis Trunking with active-active links

# Cloud Orchestration—OpenStack



- Open source cloud management framework for private and public clouds
- Created by Rackspace and NASA in July 2010
- Capturing the hearts and minds of the industry
  - Rapidly becoming the de facto open source standard for cloud computing
  - ~200 participants and 6,000+ developers
  - Allows any organization to create and offer cloud computing capabilities using open source software, rapidly and at a low cost

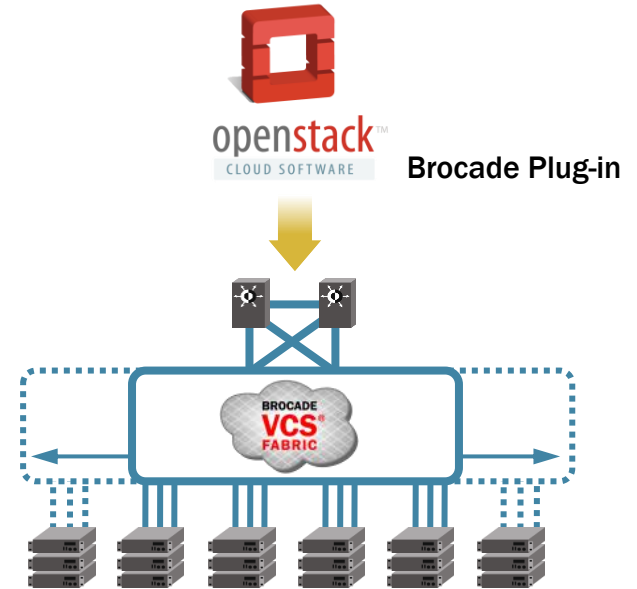


**OPENSTACK**  
CLOUD OPERATING SYSTEM

# Brocade Fabrics and OpenStack

## SELF-SERVICE, ON-DEMAND FABRIC PROVISIONING

- Brocade VCS fabric automation and OpenStack orchestration dramatically decrease time-to-deploy network capacity
- Brocade VCS plug-in contributed to OpenStack “Grizzly” release
- Brocade leading industry efforts to champion OpenStack support of Fibre Channel SANs
- Partnering with Red Hat and Piston Cloud for commercial versions of OpenStack that include Brocade VCS and FC fabrics by 2H13



# OpenStack Taxonomy

## Five Major Components

### OpenStack core services:

- Virtual Machines (compute)
- Object Store (object, data blurb)
- Block Storage aka Virtual Block Devices (hard drives, volume)
- Virtual Networks (networking)
- Dashboard (user portal)

There are two other components that serve “middleware” functionality:

- The disk image registry (Glance)
- The authorization and authentication framework (Keystone)

Cloud Core Services	OpenStack Project	Amazon Web Services Equivalent	Rackspace Equivalent
Virtual Machines	Nova	EC2	Cloud Servers
Object Store	Swift	S3	Cloud Files
Block Storage	Cinder	EBS	Cloud Block Storage
Virtual Networks	Quantum, Melange	VPC	Cloud Networks
Dashboard	Horizon	AWS Mgmt Console	Cloud Control Panel

Source: <http://www.pistoncloud.com/cloud-technology/what-is-openstack/>

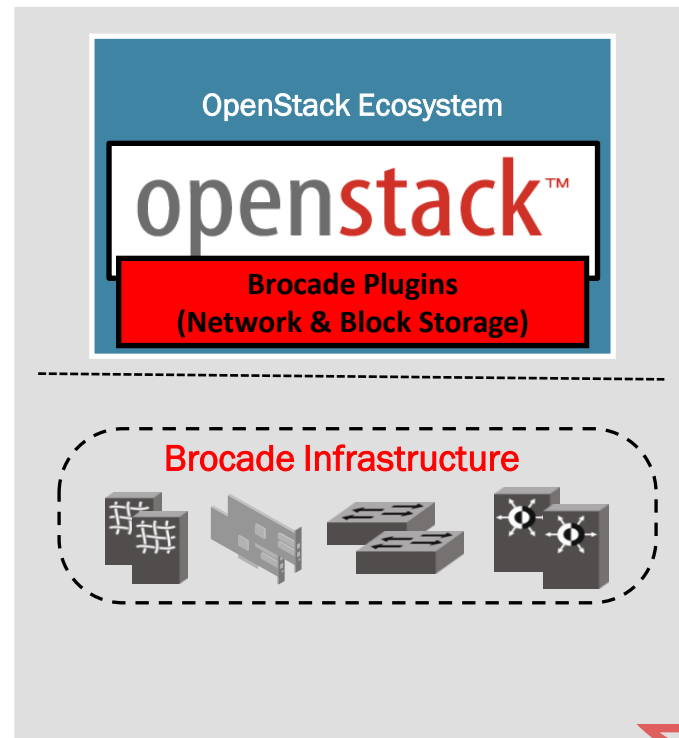
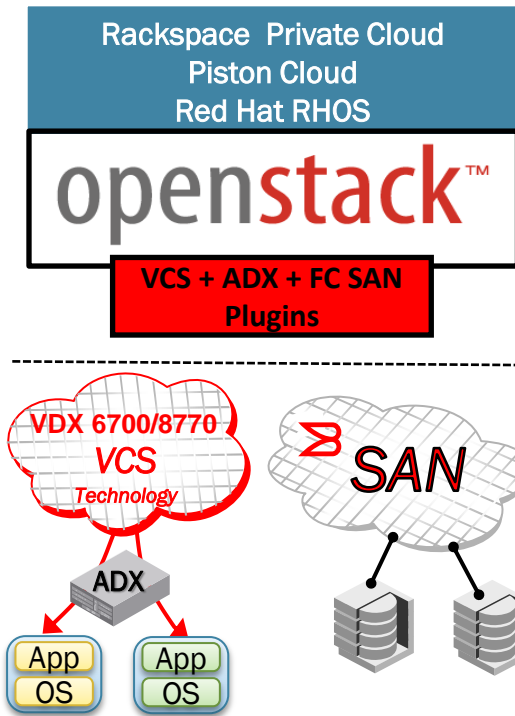
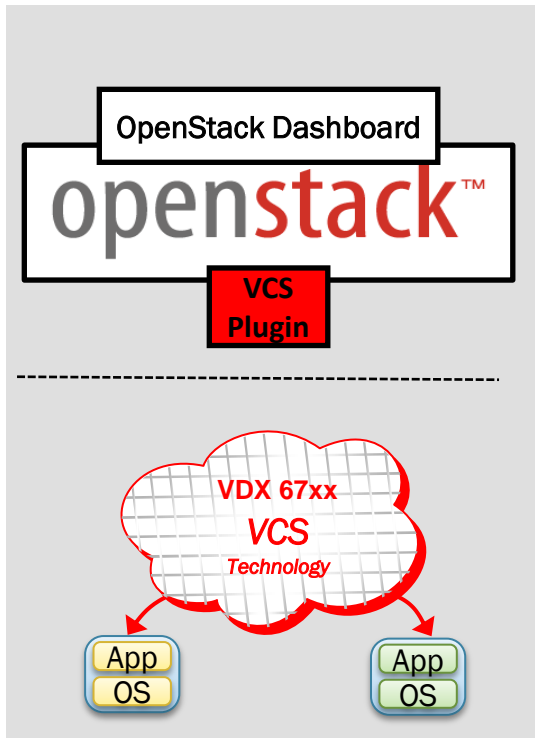




# OpenStack Integration Roadmap

## Execution Plan (Roadmap)

Phase 1 (Essex Release)	Phase 2 (Grizzly Release)	Phase n (Ultimate State)
POC delivered	H2 2013	TBD



# SDN Use Cases

SDN Will Evolve Through Value-Added Applications



# Brocade SDN Target Use Cases

## WAN Network Virtualization


WAN Virtualization  
App & SDN Controller



1

## Services Creation & Insertion


Services Insertion  
App & SDN Controller



2

## Network Analytics


Network Analytics  
App & SDN Controller



3

## DC Network Virtualization


DC Virtualization  
App & SDN Controller



4

## Application Delivery


ADP APP & SDN Controller



5

## SDN Cloud Gateway


SDN Orchestration &  
SDN Controller



6

## SDN Packet-Optical Integration

Packet-Optical Integration  
APP & SDN Controller



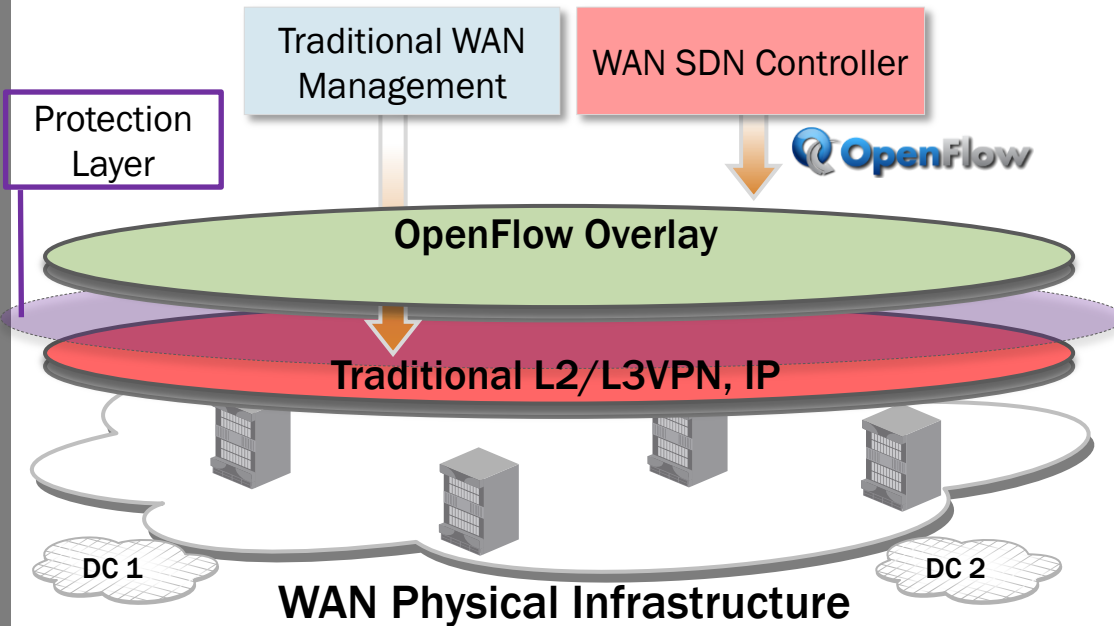
7

8



# WAN Network Virtualization

## Traditional L2/L3VPN-IP Network with OpenFlow Overlay



- OpenFlow as an overlay to existing network
  - Allows for new revenue-generating features on top of existing production network
- Enabled by Brocade's **“Hybrid port mode”**
  - OpenFlow and traditional features enabled concurrently on same router ports
- **Protected Hybrid Port Mode**
  - OpenFlow does not affect Traditional traffic
  - Protection in hardware
  - Allows for initial OpenFlow overlay service development without risk

# BROCADE OPENFLOW ENABLED 100G NATIONWIDE BACKBONE

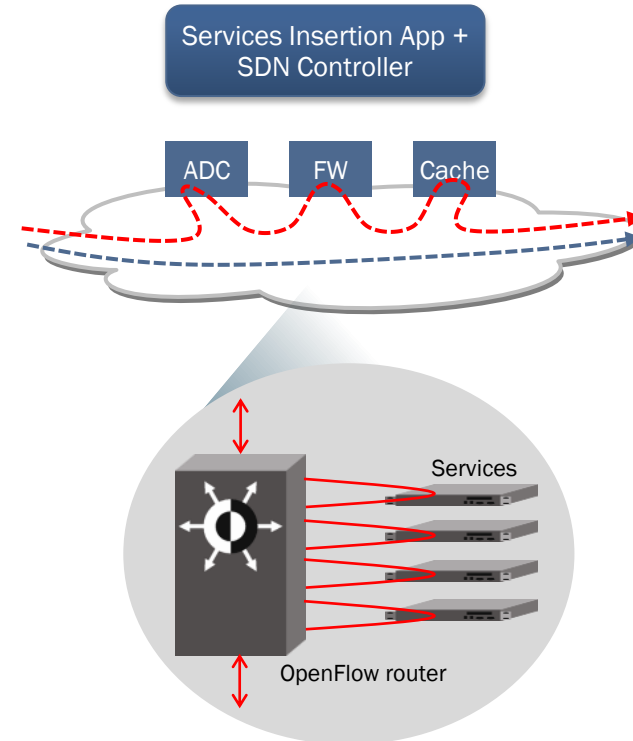


## Internet 2

- 49 Custom Location Facilities
- 15,500 miles of dark Fiber
- 8.8 Tbps of Optical Capacity
- Hybrid Mode with protected OpenFlow traffic

# Services Creation & Insertion

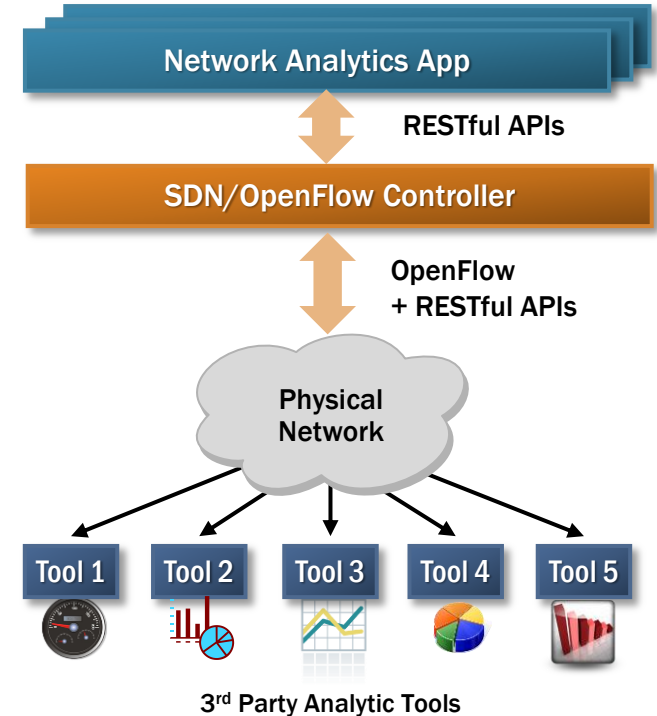
- SDN automates
  - Traffic steering to achieve desired pipeline of services
  - Customization of services according to customer needs
- Optimizes use of network resources
  - No need to steer traffic through traffic steering appliances



# SDN Approach to Network Analytics

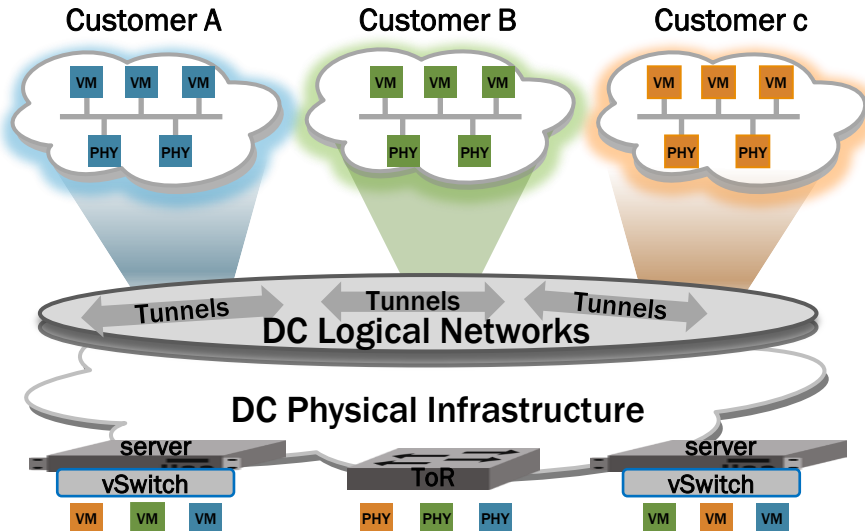
## Unlocking Advanced Operational Intelligence

- Why network analytics is important
  - Real-time network statistics collection & alerting
  - Summarization of normal and abnormal traffic
  - Detect network performance issues in advance of customer complaints
- Use cases
  - Internet/Mobile traffic analysis: Facebook, Youtube, Email, ...
  - Big Data analysis
  - Detection of unlawful content
  - ...



# Data Center Network Virtualization

## Scalable Cloud Services



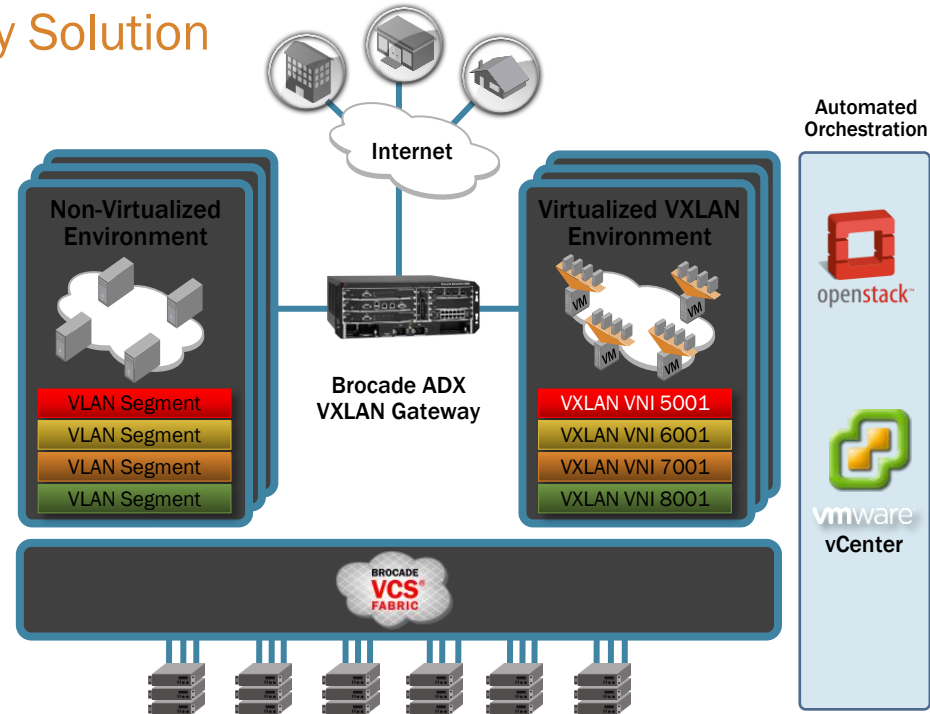
- Tunnels enable physical network abstraction (logical network)
  - VxLAN, NVGRE, STT
- Software Switches (vSwitches) connect virtual machines
- ToRs connect physical machines
- SDN Gateways enable scalable connectivity into the logical network



# Brocade ADX VXLAN gateway solution

## Brocade ADX Physical to Logical Gateway Solution

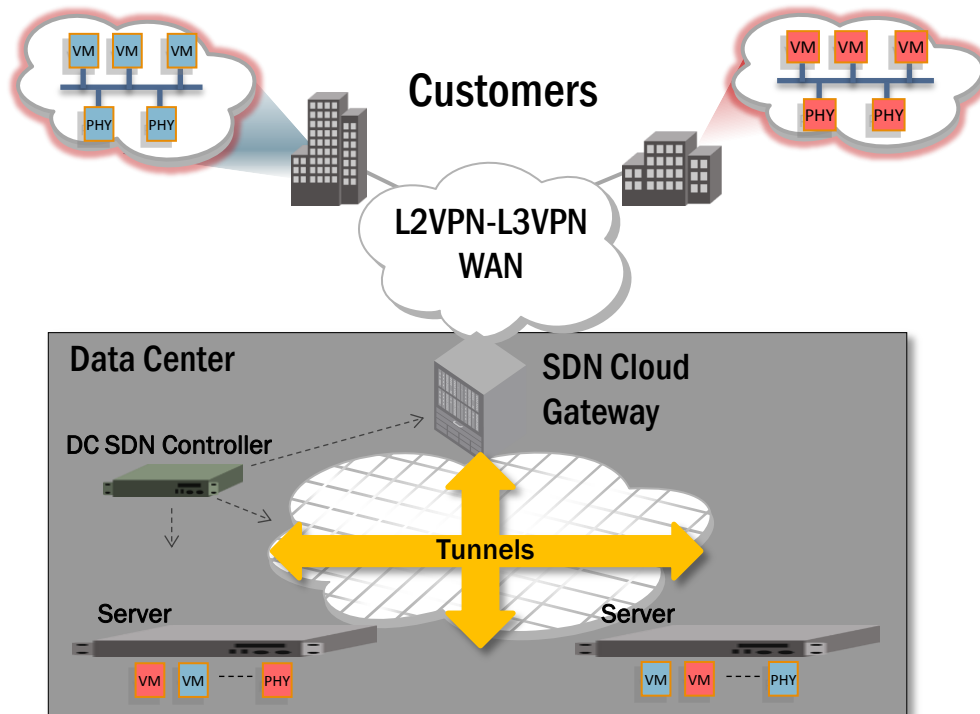
- Industry's first integrated VXLAN gateway and application delivery controller (ADC)
  - Connects VXLAN environments to the Internet
  - Spans VXLAN and non-VXLAN environment
  - Enables transparent access to both VXLAN connected VMs and physical servers
- VMware vShield Manager and vCenter support
- Demonstrated at VMWorld San Francisco, September 2012



**VXLAN GATEWAY ON THE BROCADE ADX  
APPLICATION DELIVERY SWITCHES**

# SDN Cloud Gateway

## Interconnecting DC Logical Networks to WAN L2/L3VPNs



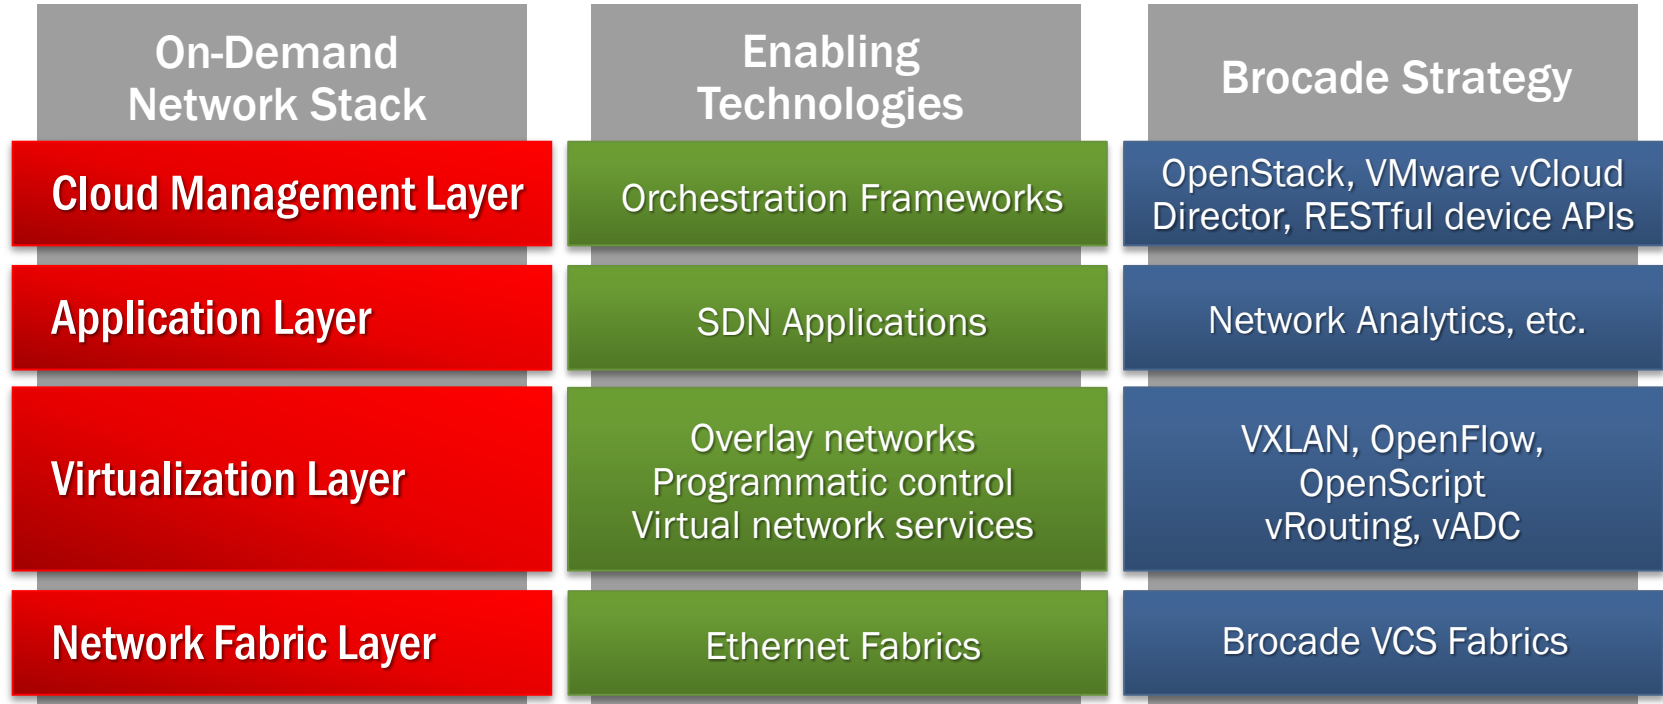
### SDN Cloud Gateway

- Direct mapping from customers' DC Logical Networks to WAN L2/L3VPNs
- SDN point of control between WAN and DC network

### Cloud Scale

- Optimized to large DC multi-tenancy requirements
- Inter-DC connectivity
- Termination of tunnels: VxLAN, NVGRE, etc

# Summary: Brocade SDN in the Data Center



# Leadership through Innovation

- Empowering the On-Demand Data Center
- Leadership and rapid innovation in Fabrics, IP routing and Software-defined Networking
- Delivering the world's most automated, efficient and agile networks

**BROCADE**  <sup>®</sup>