BROCADE

Brocade Ethernet Fabrics

Łukasz Kozłowski

25 Nov 2015

Brocade
@Brocade

What Is an Ethernet Fabric?

- Optimized for server and network virtualization
- More efficient, higher throughput, lower latency
- Scale-out vs. scale-up to increase flexibility and protect investment
- Automated deployment and ongoing administration

Ethernet Fabrics vs. Legacy Networks

CLASSIC HIERARCHICAL ARCHITECTURE

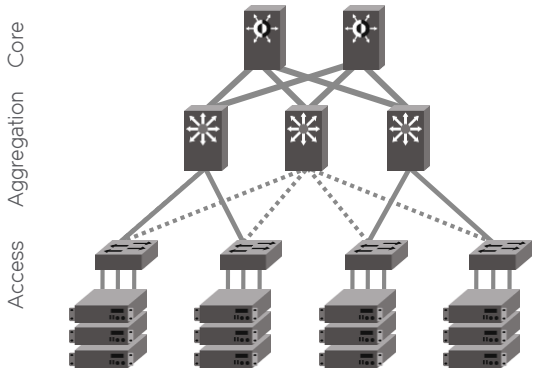

- Rigid architecture, north-south optimized
- Inefficient link utilization
- Individually managed switches
- VM-ignorant
- No network virtualization

ETHERNET FABRIC ARCHITECTURE

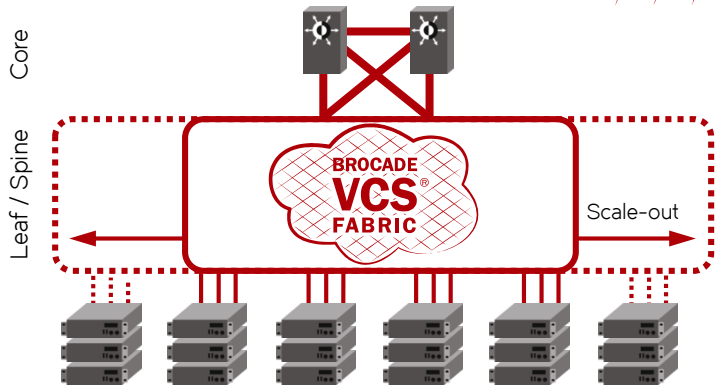

- Flat topology, east-west optimized
- All links active, L1/2/3 multipathing
- Fabric managed as one logical switch
- VM-aware
- Native and overlay network virtualization

TRILL – Transparent Interconnect of Lots of Links

Terminology

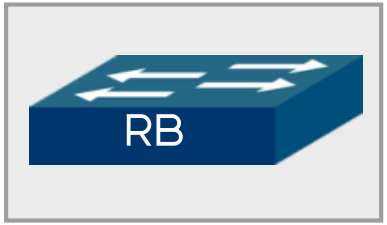

Devices are Routing Bridges (RBridges or Rbridges)

Data Plane is TRILL protocol

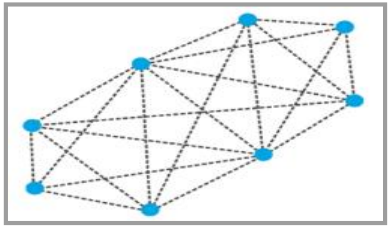

Control Plane is a L2 link state routing protocol (FSPF+OSFP-like)

...and the challenges are

Challenge 1 : Operational Complexity

Solution : Logical Chassis Management

Traditional Management

N points of management

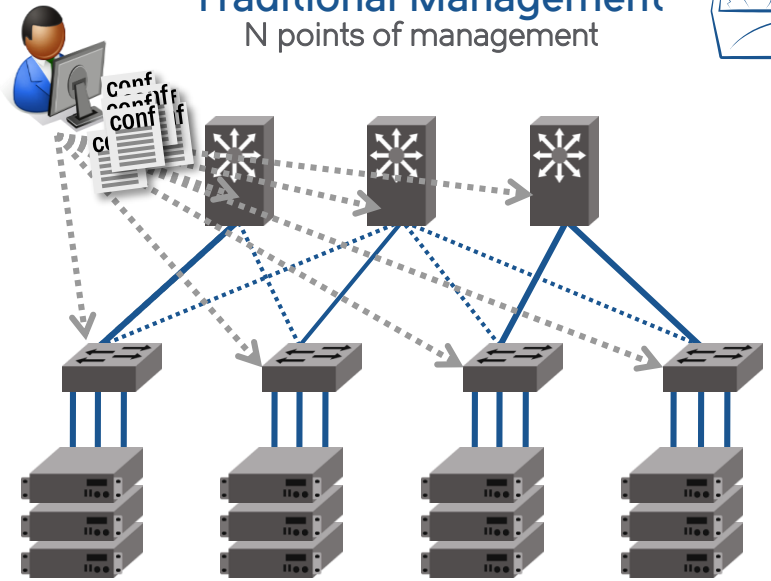

ADMINISTRATIVE COST & COMPLEXITY INCREASES WITH DATA CENTER SCALE

VCS Logical Chassis

Single point of management

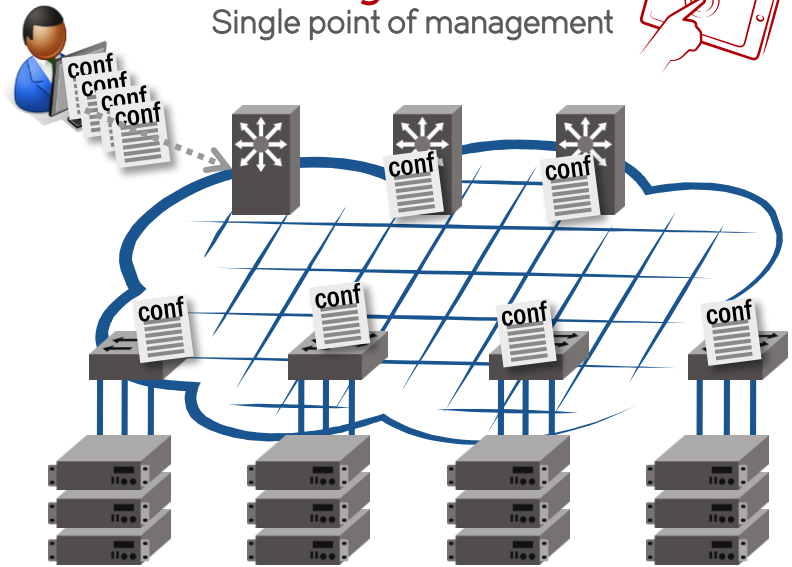

ADMINISTRATIVE COST REMAINS RELATIVELY FLAT WITH DATA CENTER SCALE

Challenge 1 : Operational Complexity

Solution : Automation

- Automatic fabric creation and expansion (two steps, 1 to 2 minutes)
- Automatic trunk creation

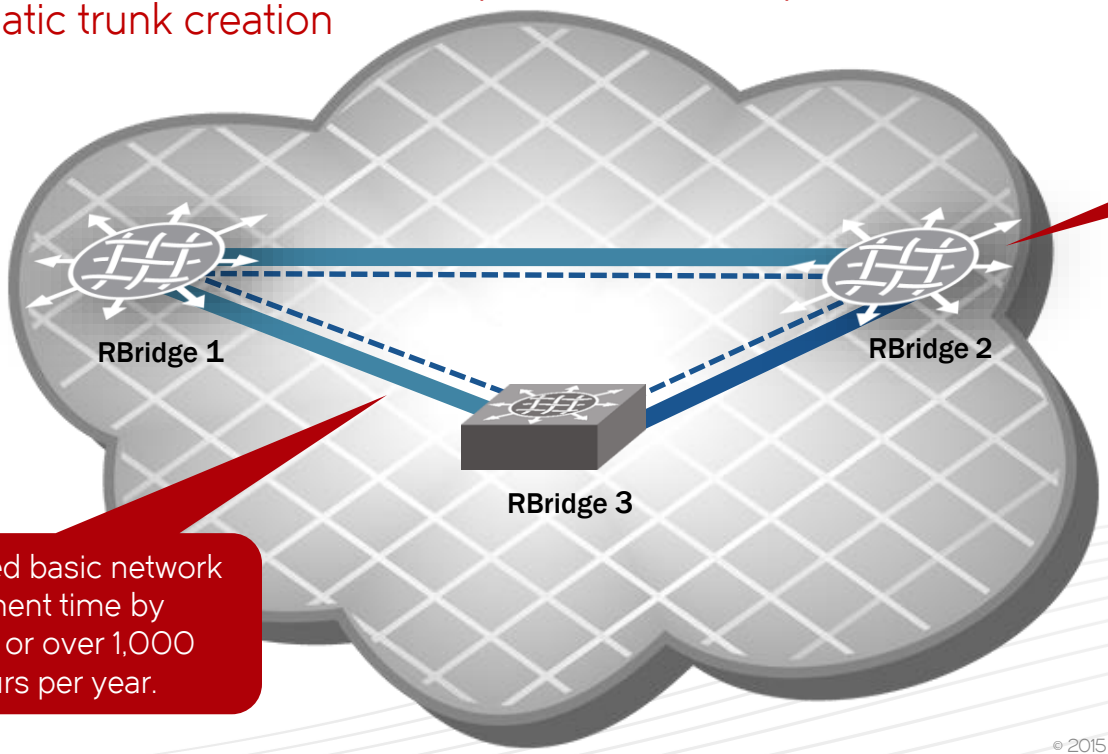

SunPower reduced setup time by over 80 percent compared to its three-tier network.

30 GbE Data Center Bridging (DCB) trunk (3 × 10 GbE)

10 GbE DCB link

GEHA reduced basic network management time by 20 percent, or over 1,000 man-hours per year.

VCS DIFFERENTIATOR :

Automa

Automa

Configuring LAG (for 2 members)	Configuring ISL Trunking (for up to 8 members)
Execute the following commands on one switch: • configure terminal • interface port-channel 1 • switchport • switchport mode trunk • switchport trunk allowed vlan all • qos flowcontrol tx on rx on • mtu 9208 • no shutdown • interface tengigabitethernet 1/0/5 • channel-group 1 mode active type standard • no shutdown • interface tengigabitethernet 1/0/6 • channel-group 1 mode active type standard • no shutdown • exit Repeat same commands on other end switch. Total commands: 30	Absolutely no configuration required. Total commands: 0

(10GbE)

10GbE DCB Link

Challenge 1 : Operational Complexity

Solution : Sharing Port Profiles

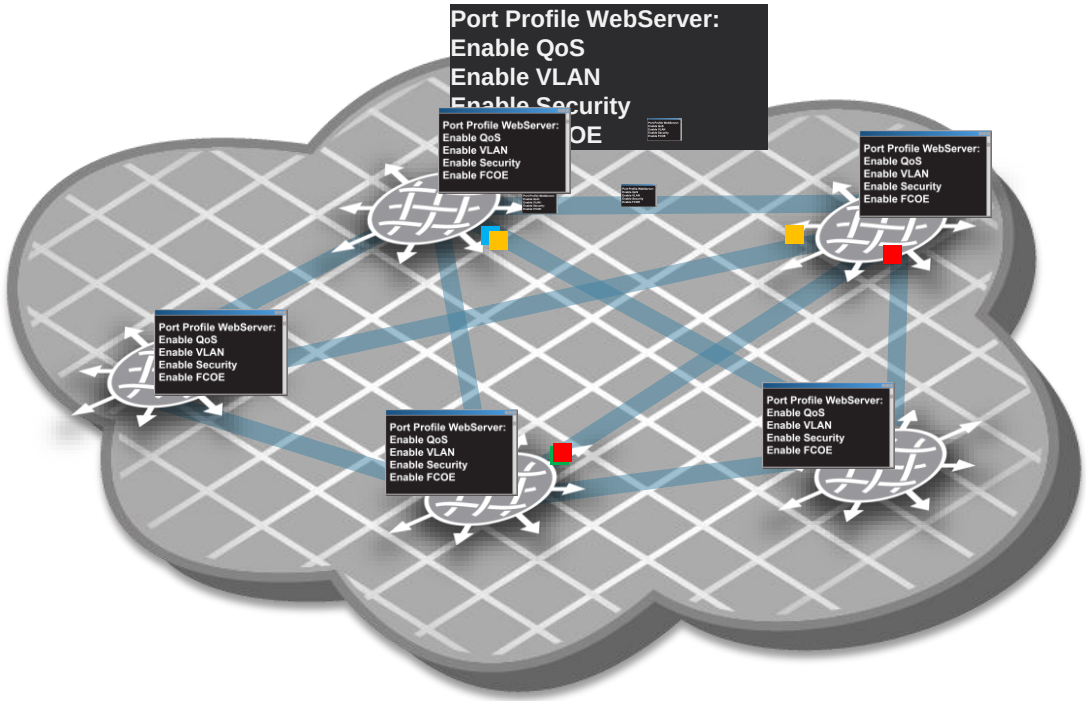

Challenge 1 : Operational Complexity

Solution : Automatic Migration of Port Profiles

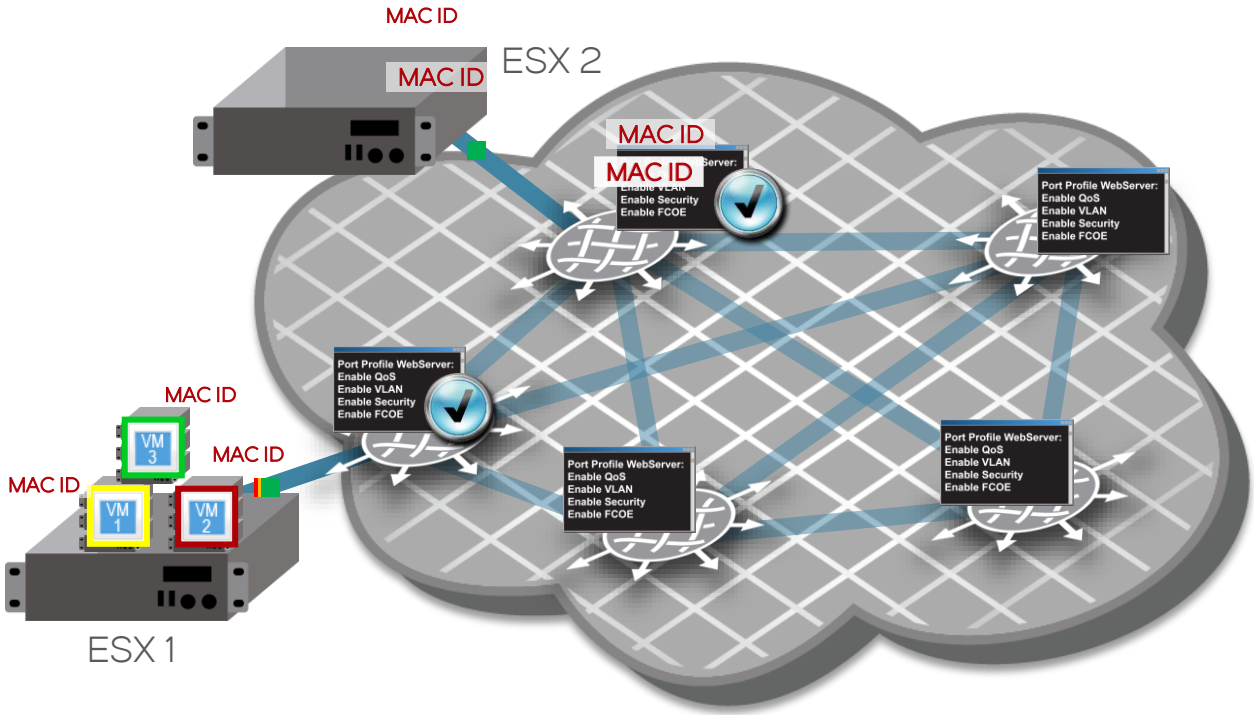

Challenge 1 : Operational Complexity

Solution : VM-aware Network Automation

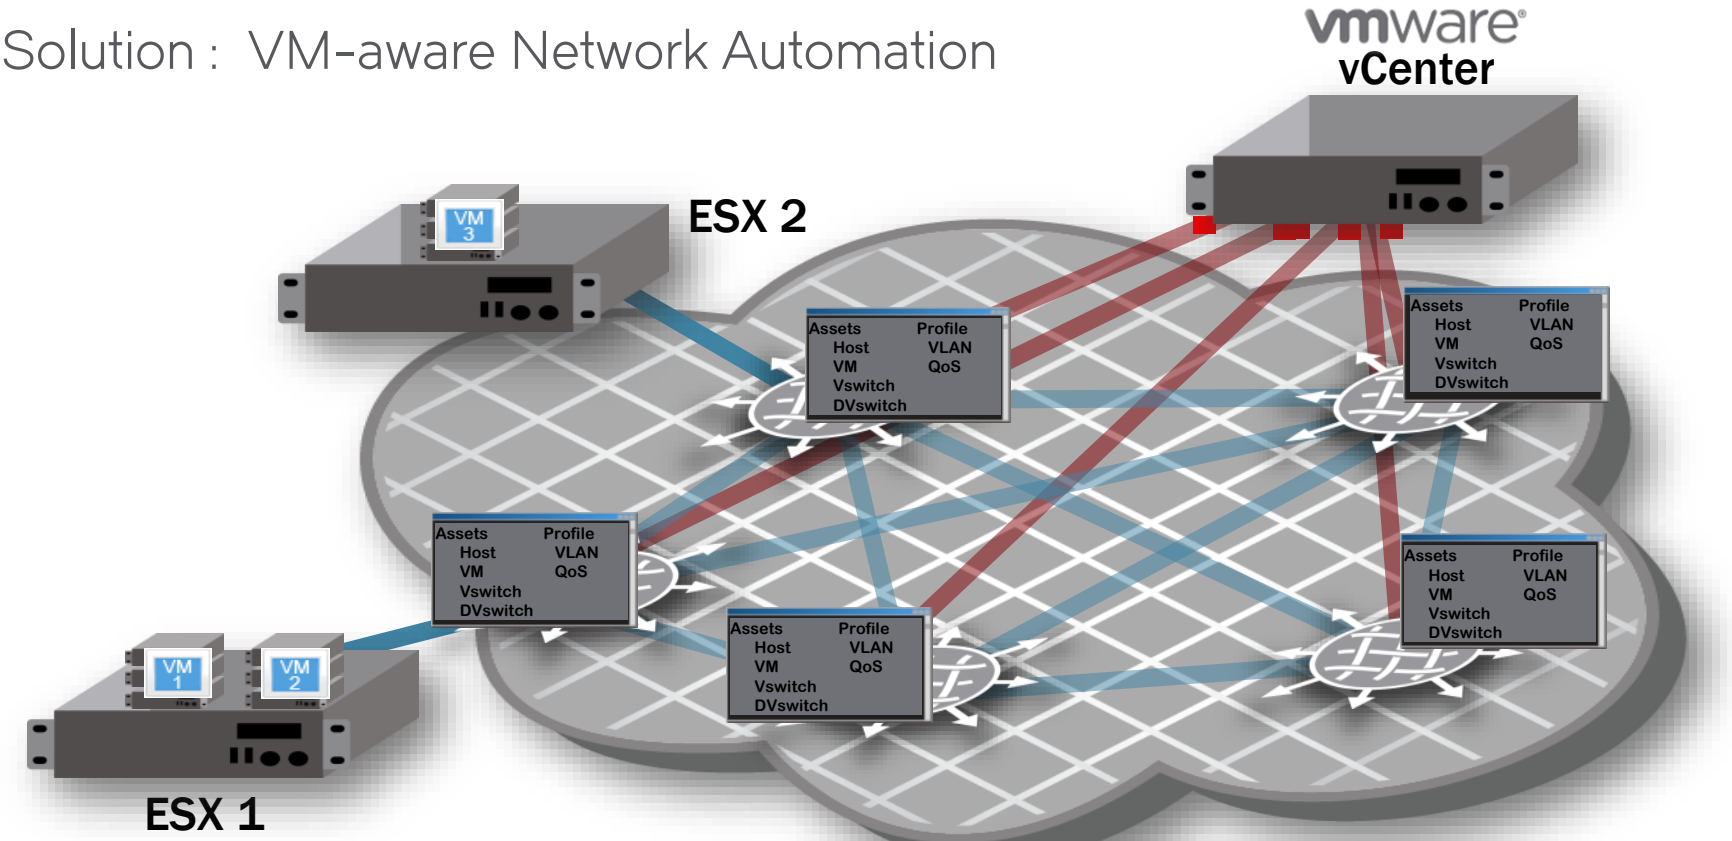

Challenge 1 : Operational Complexity

Solution : Distributed Information to enable VM Mobility

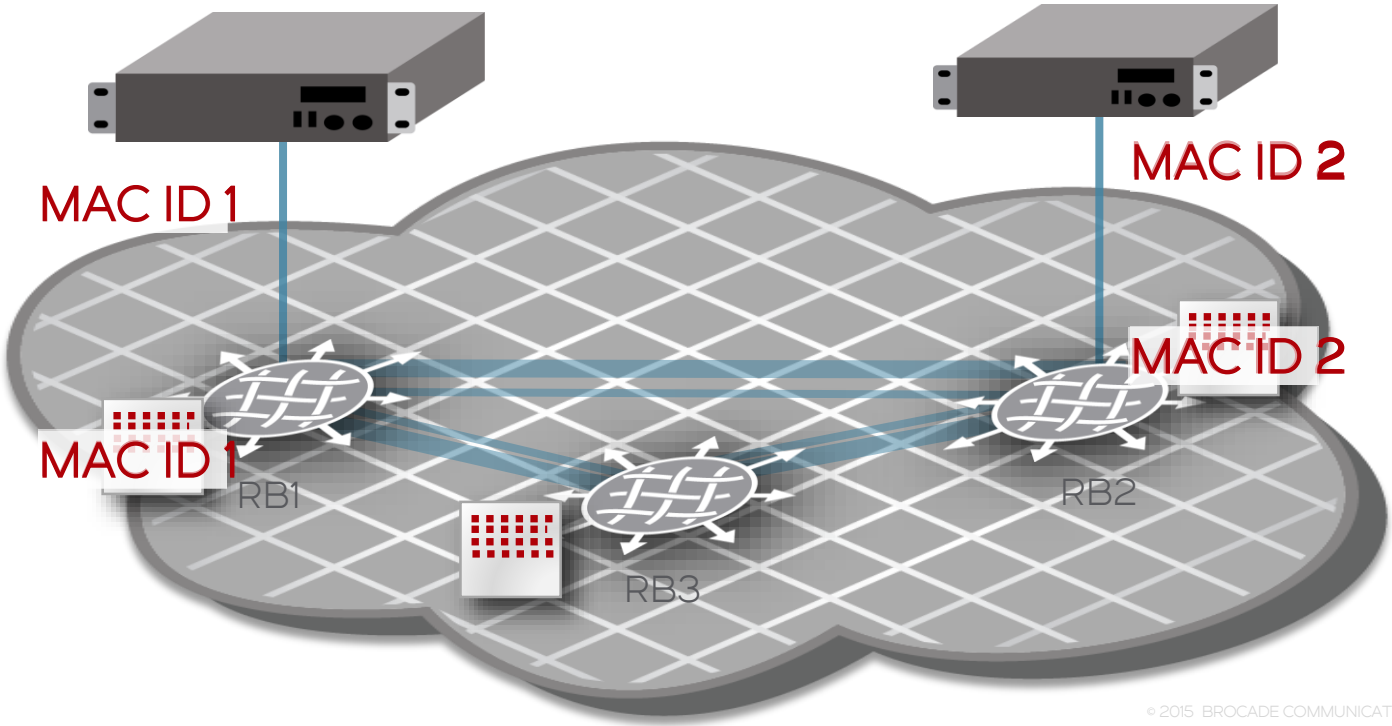

Challenge 2 : Efficiency

Solution : vLAG (LAG to different physical VDX Switches)

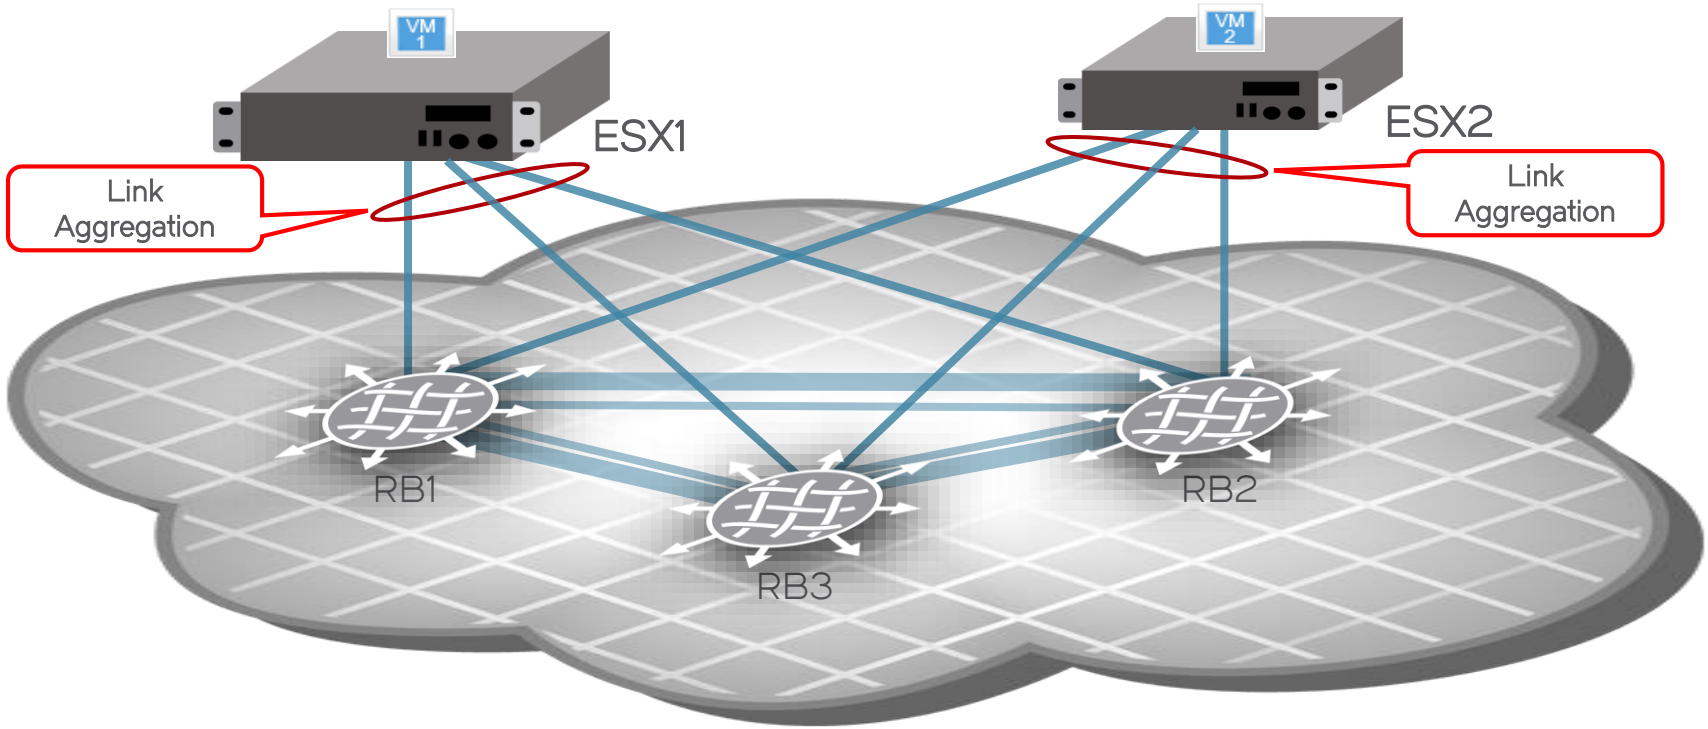

Challenge 2 : Efficiency

Solution : Increase efficiency with Brocade Trunking

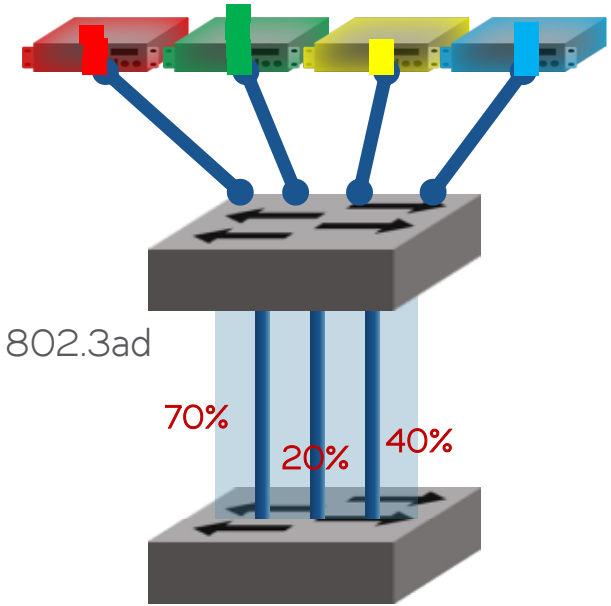

Traditional algorithms cause imbalances

- High link utilization and ease-of-use
  - All 10GE ports are not alike
- Frame-level, hardware-based trunking at Layer 1
  - Near 100% link utilization versus 802.3ad LAG groups ~50-60% link utilization
  - Single flows can be split across all links
  - Frames are distributed across links in the trunk
  - Built into Brocade fabric switching ASIC

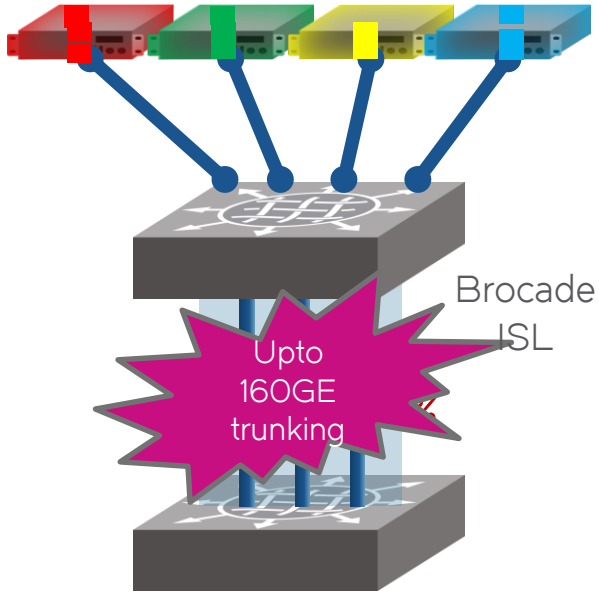

Brocade Trunking increase link efficiency

Challenge 2 : Efficiency

Solution : Brocade VCS Multi-pathing at Multiple Layers

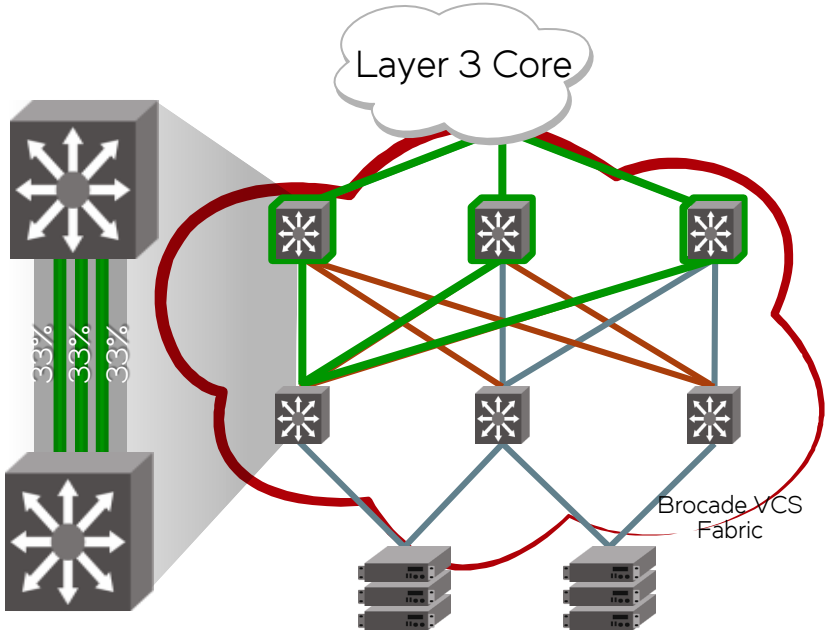

FEATURE	BENEFIT
L1: Trunking with frame striping	Near-perfect load balancing across all links in a trunk group
L2: Equal Cost Multi-Pathing (ECMP)	All links utilized with flow-based load balancing
L3: Fabric load balancing across multiple L3 gateways	Improved scalability and resiliency

Challenge 3 : Multi-tenancy

Solution : Virtual Fabric

CHALLENGE

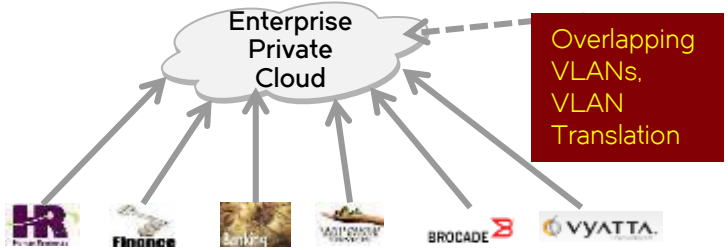

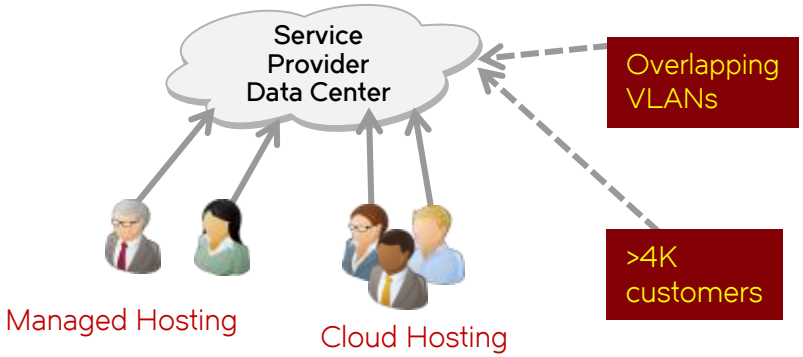

SOLUTION

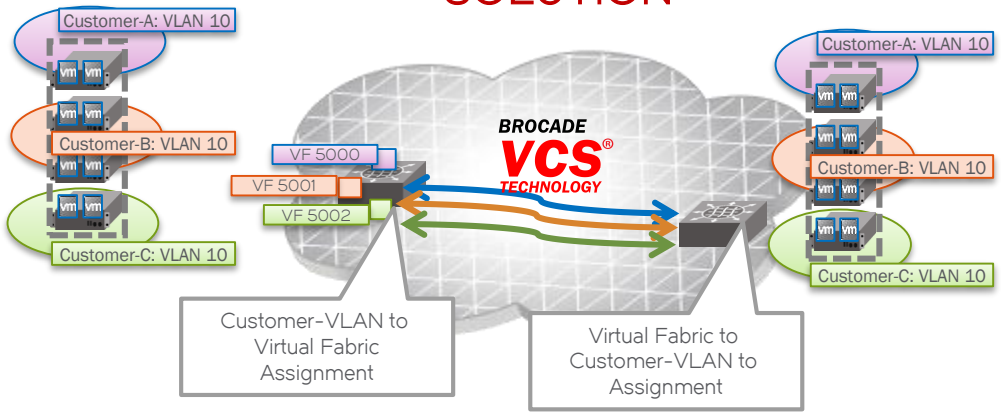

- Virtual Fabric is a native Ethernet Fabric based Multi-Tenancy solution
- It is based on standards based TRILL Fine-grained Labels. RFC # 6325
- Virtual Fabric provides support for overlapping vlans, vlan scale and transparent vlan services.

Challenge 4 : Resiliency

Solution : Multi-layer Fabric HA & ISSU

CHALLENGE

With an Incident length of **90 minutes**, the average cost of a single downtime event was approx. **$505,500**

The Average cost of Data Center downtime is approximately **$5,600 per minute***

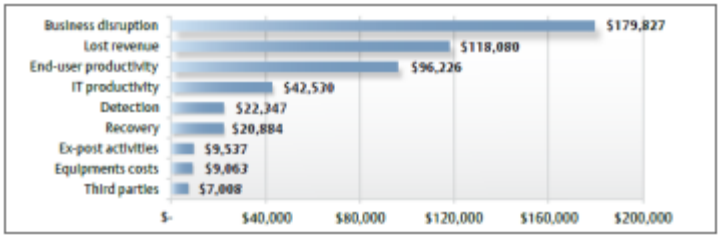

SOLUTION

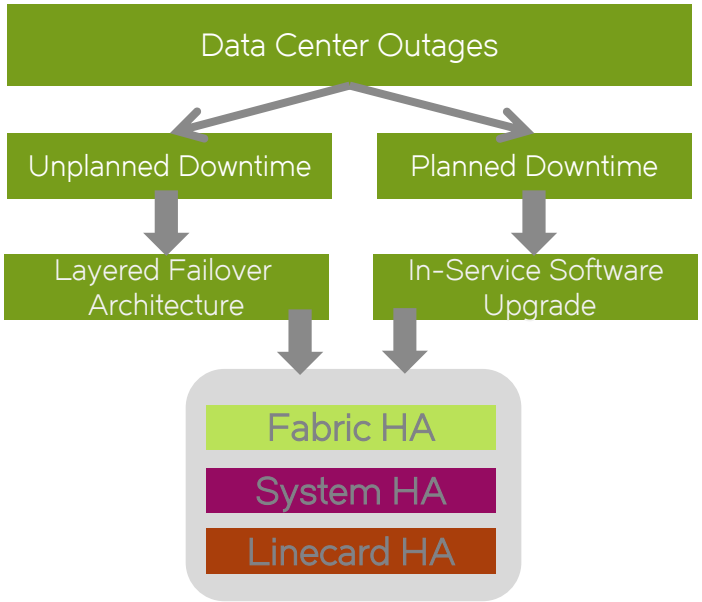

BENEFIT – Minimize risk of Network Downtime by implementing a self healing/intelligent network infrastructure.

Challenge 5: Complexity of Network Virtualization

Not every DC asset understands VXLAN

- Some applications running on physical servers like Oracle DB etc.
- Storage is still physical and non-VxLAN aware
- Existing appliances like firewalls and server load-balancers.
- Access to existing L3 networks via existing routers.

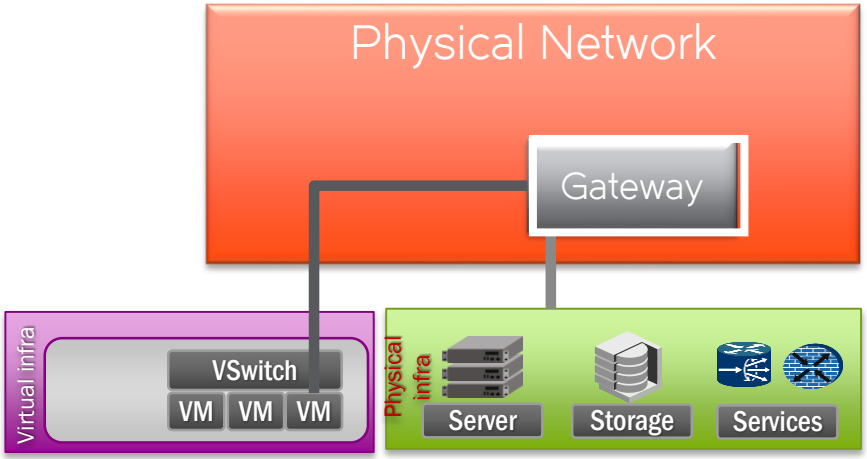

A VXLAN Gateway bridges virtual and physical assets

Network Virtualization with Brocade VCS & VMware NSX

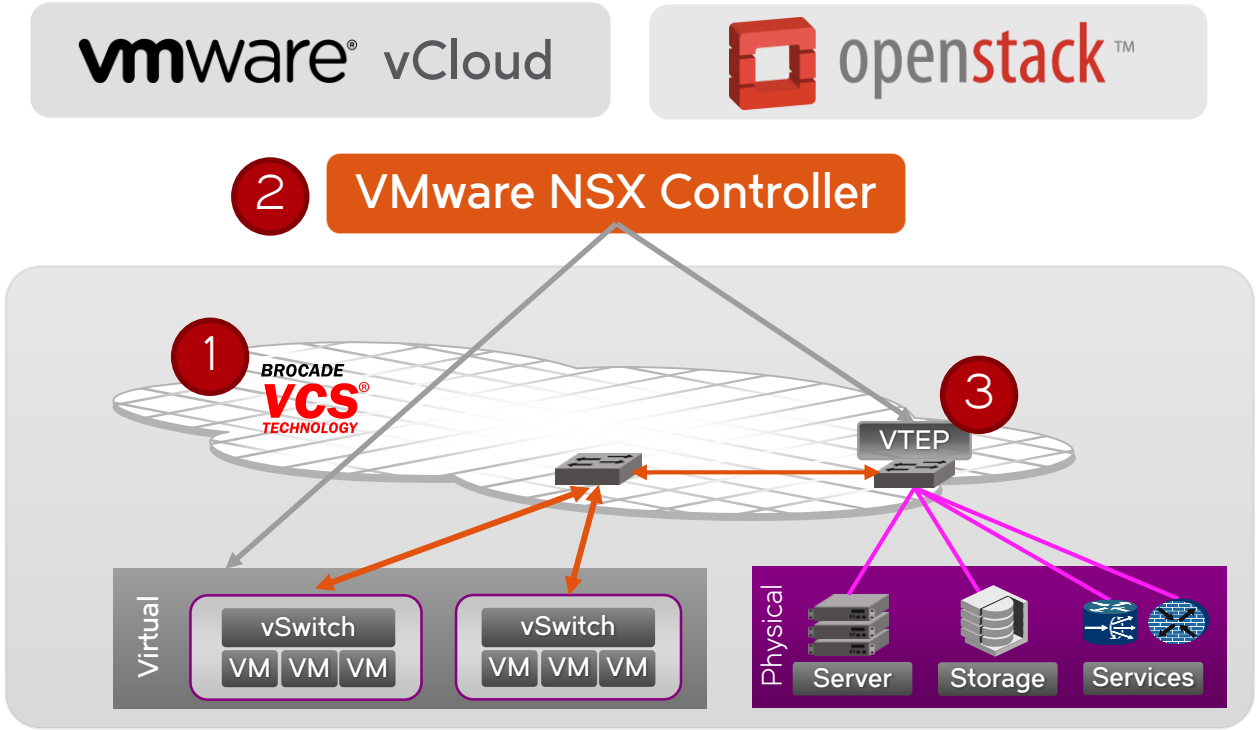

1. **Brocade VCS Fabric:** Automated, resilient physical foundation

2. **VMware NSX:** Network virtualization overlay and controller

3. **Brocade VCS Gateway for VMware NSX:** Bridges virtual and physical infrastructure

VDX 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 Rad Hat and Piston Cloud for commercial versions of OpenStack that include Brocade VCS and FC fabrics

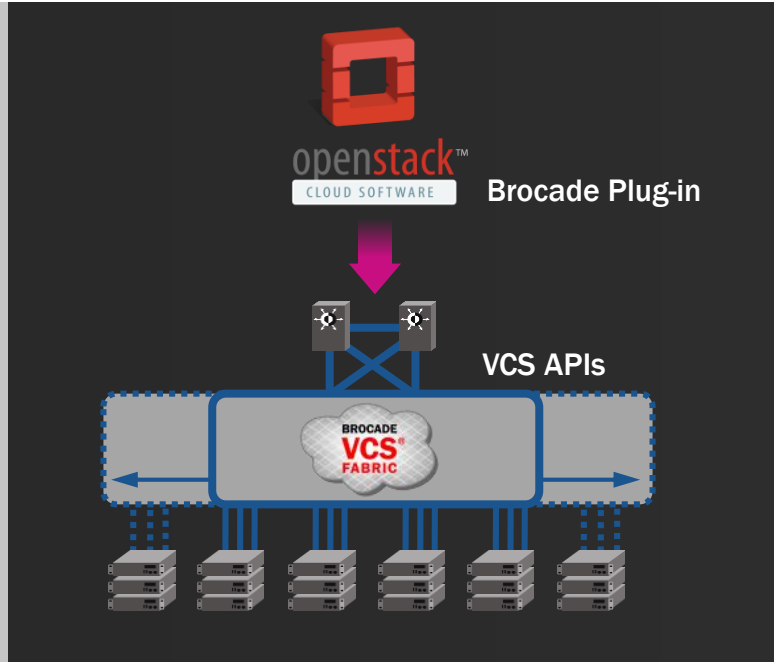

Challenge 6 : Metro VCS

Solution : Stretching VCS fabric over metro distance

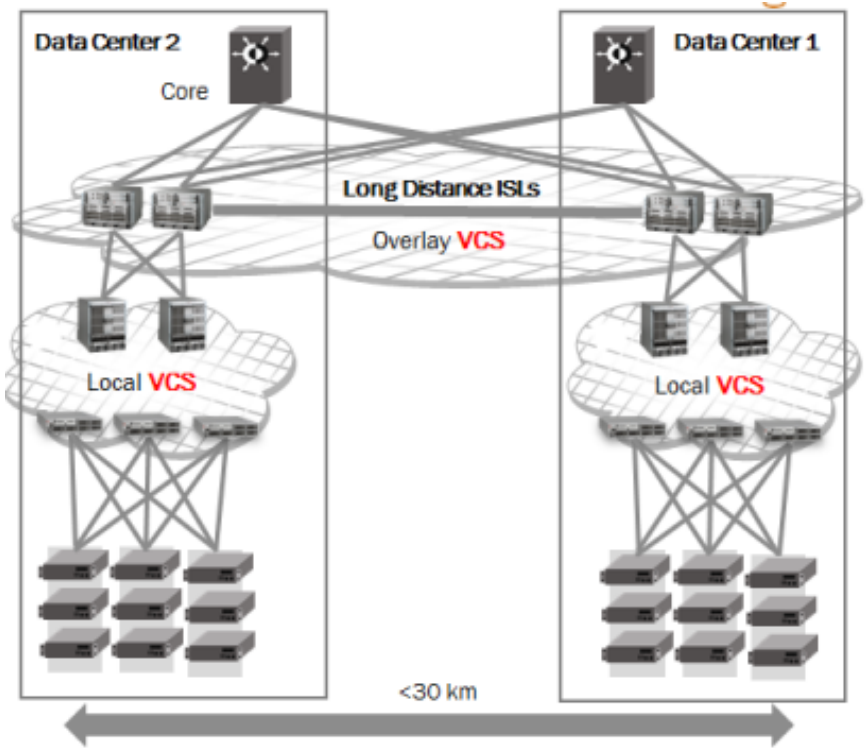

Brocade Multi-Fabric VCS Design
POD Design for Increased Scale - Multi-Fabric VCS deployment with vLAG connectivity between fabrics. Each Data Center can scale independently
Active – Active L3 Gateways - Support up to 4 VRRP-E Gateways in the Overlay VCS cloud.
Distribution of L2 Applications - Shared VLANs over the Overlay VCS
Localized VLANs - Certain VLANs remain native to the respective Datacenter with the Local VCS Fabric
Seamless Transition of Services - Overlay VCS to accommodate Services (IDS, SLB & Firewalls)
Optimize Network Utilization - Contain L2 BUM traffic within Local VCS Clusters

Challenge 7 : IP Storage Connectivity

Solution : Auto NAS Traffic Prioritization

- **Dual NAS Head Architecture** File access to the NAS head, optional block access to the storage nodes.
- **Scale-Up or Out** Scale up the storage by adding storage nodes behind the NAS heads, typically limited to a pair of NAS heads. Or deploy a scale-out architecture such as EMC Isilon.
- **Network Requirements** Ideally highly automated, efficient and simple to deploy. In converged front-end network environments, important to have protection of NAS traffic. Predictable, reliable network transport, no hot spots.

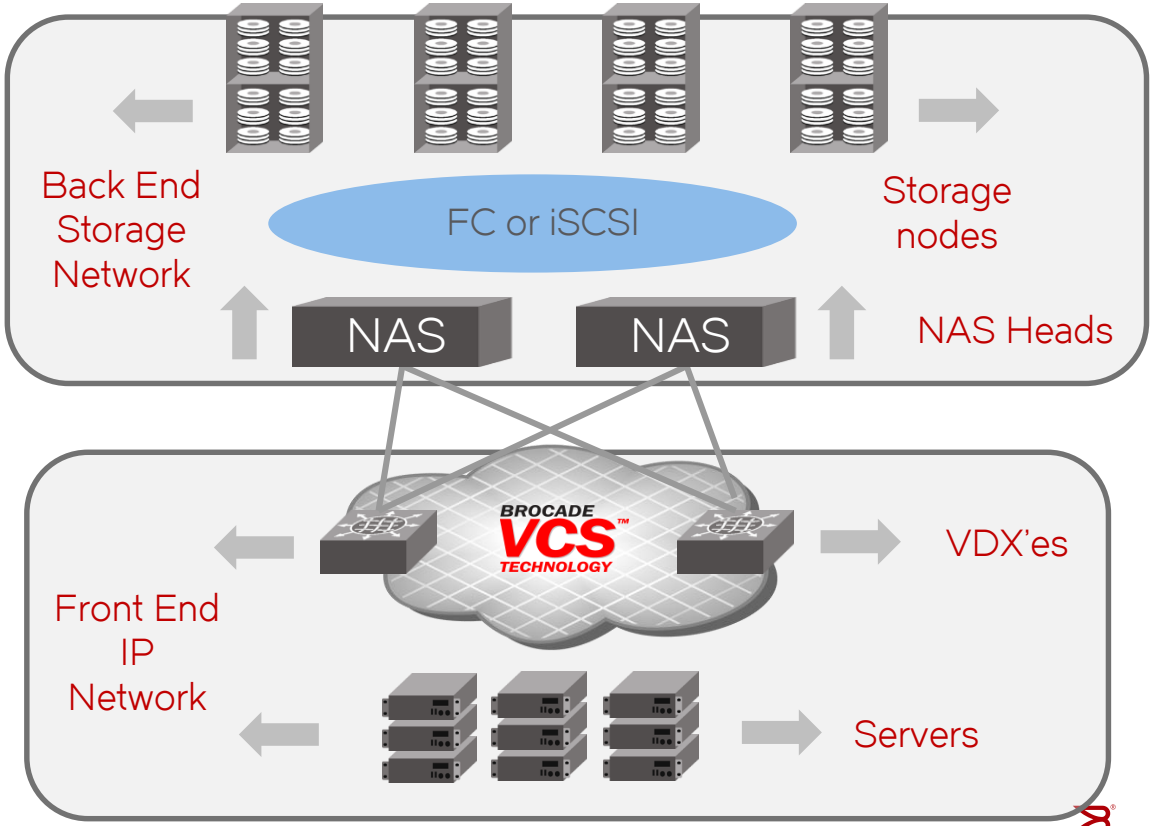

What is Auto NAS

NAS Traffic Prioritization

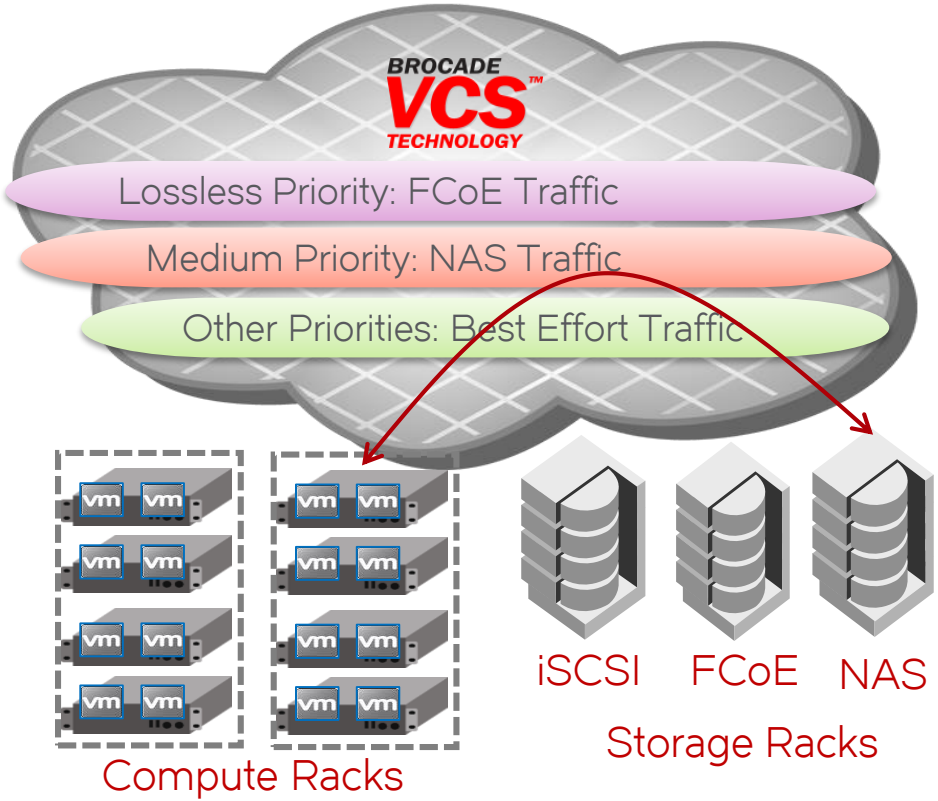

- Fabric is designed for versatile storage technology: NAS, iSCSI, Object, FCoE
- Multipathing, Low Latency, Lossless, Highly Resilient Architecture enables Scale Out Storage
- All the nodes in the fabric will auto prioritize NAS Storage traffic over other traffic types (NFS, SMB/CIFS)
- Ability to monitor IP Storage traffic through ACL Counters

Challenge 8 : Flexibility in deployment

Solution : Supports L2 and L3 Fabrics, SDN and classic models also

DIFFERENTIATOR

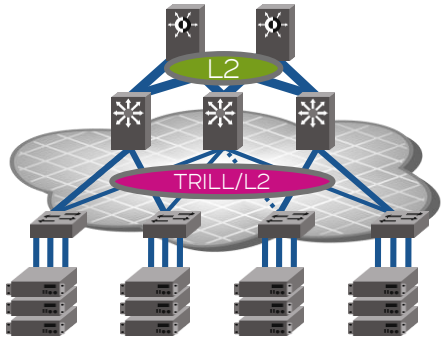

TRILL + L2

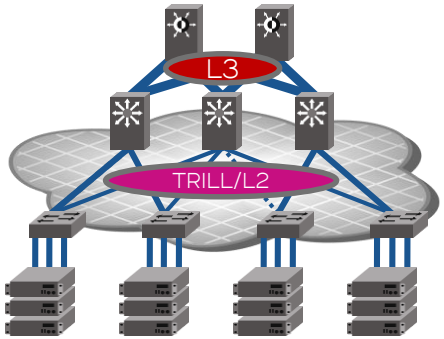

TRILL + L3

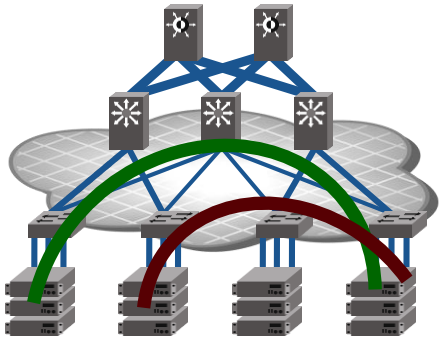

TRILL + Overlays

CLASSIC

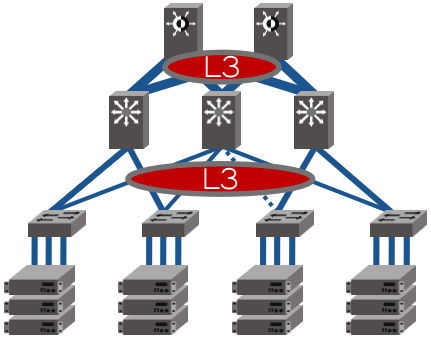

Classic L3

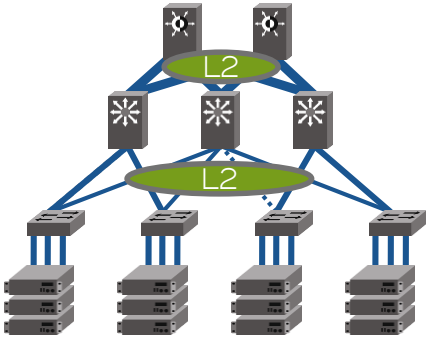

Classic L2

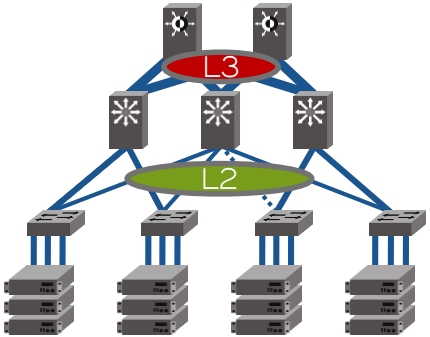

L2 + L3

Challenge 8 : Flexibility in deployment

Solution : Classic to Fabric Transition

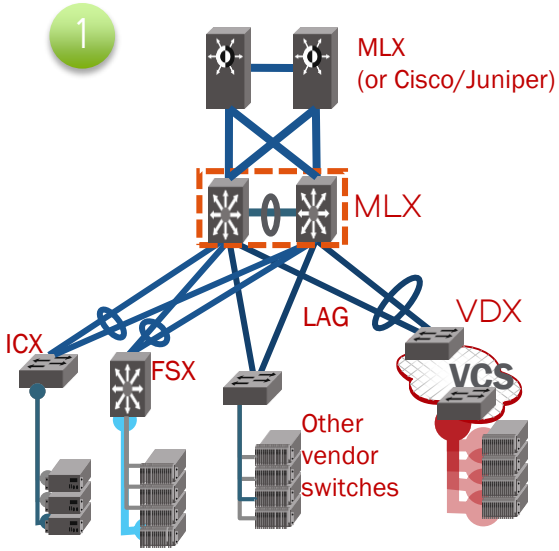

- Intro VCS in a POD
- Flat layer 2
- Co-exist with legacy access switches

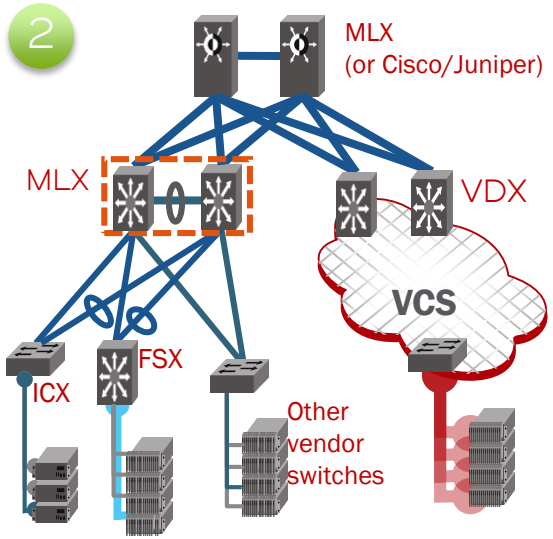

- Scale Out VCS
- Co-exist with legacy access/aggregation

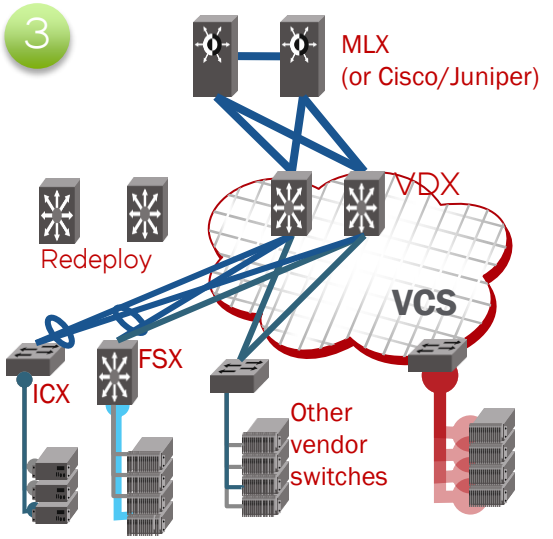

- Fully deployed fabric
- Core-Edge Solution
- Legacy network support

Brocade VDX Fixed Switch Family

VCS fabric-enabled switches

Brocade VDX 6740T-1G Switch

- 1RU form factor & Single ASIC design/ISSU
- 48 × 1000BASE-T ports
- Low latency— 3µs (All packet sizes)
- Industry's 1st 1GbE to 10GbE S/W upgradable switch
- SDN Ready (Open Flow 1.3 support)
- VXLAN & VTEP support
- VCS Virtual Fabric Support
- Ports on Demand (POD)

Brocade VDX 6740/6740T Switches

- 1RU form factor & Single ASIC design/ISSU
- 48 × 1/10GbE SFP+ OR 48 × 1/10GBASE-T and 4 × 40GbE QSFP ports
- 32 Flexports (FC/Ethernet)
- SDN-ready (Open Flow 1.3 support)
- VXLAN & VTEP support
- VCS Virtual Fabric Support
- Low latency—Fiber (850ns) / Copper (3us)
- Ports on Demand (POD)

Brocade VDX 6940-36Q/ /6940-144S Switch

- 1/2RU form factor & Single ASIC design/ISSU
- 36 × 40GbE QSFP ports
- 96 x 10GbE SFPP + 12 x 40GbE QSFP/4 x 100GbE QSFP28
- 32 Flexports (FC/Ethernet)
- SDN-ready (Open Flow 1.3 support)
- VXLAN & VTEP support
- VCS Virtual Fabric Support
- Low latency—650ns
- Ports on Demand (POD)

Brocade VDX 8770 Modular Switch

- 4 and 8RU form factors with ISSU
- Ultra-high availability
- 384 × 1/10GbE SFP+ ports
- 384 × 10GBASE-T ports
- 216 × 40GbE QSFP ports (576 × 10GbE w/ breakout)
- 48 x 100GbE CFP2 ports with Ports on Demand (POD)
- SDN-ready (Open Flow 1.3 support)
- VCS Virtual Fabric Support
- Low latency—3.5µs

Summary of VCS Fabric Value Propositions

AUTOMATED

- Zero touch provisioning
- Zero-touch VM discovery, configuration, and mobility
- Self-forming trunks
- Manage many switches as single logical device

50% lower opex

EFFICIENT

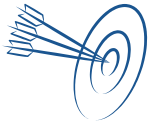

- All links fully active, none on standby
- Multi-pathing at all layers of the network: L1/L2/L3
- Most efficient platform for IP storage

2x greater network utilization

AGILE

- Network virtualization with VCS Virtual Fabrics or VMware NSX
- Scale-out non-disruptively
- Orchestration thru OpenStack

Quicker to deploy

BROCADE

Thank you!