BROCADE

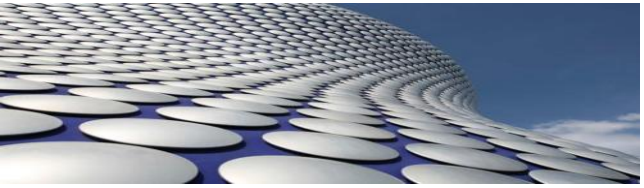

BROCADE VIRTUAL CLUSTER SWITCHING (VCS)

Tomasz Pokora

Solutions Consultant Eastern Europe

February 2011

Brocade One

Classic Data Centre

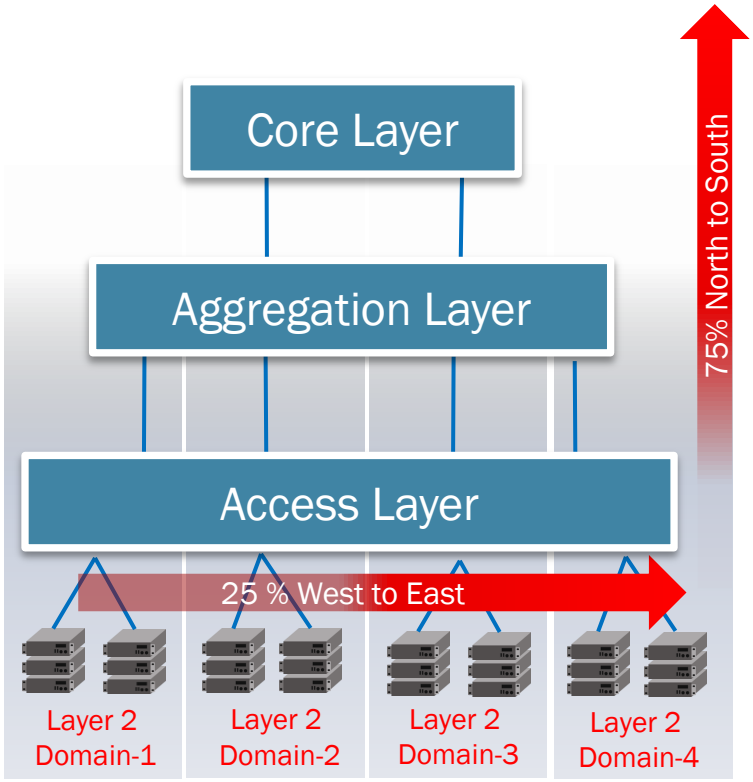

- Generation One Data Centre
  - Designed for North to South Traffic
  - Client to Server traffic model
  - Designed for transport, not the application
- Generic Enterprise Solution
  - Enterprise technologies -stacking
  - Enterprise topologies- STP, MSTP
  - Enterprise limitations – STP, stacking
  - Minimize Layer 2 fault domains
  - Increased Management footprint
  - Multi-layered, multi-protocol architectures for scalability

Brocade One

The Next Generation Data Centre

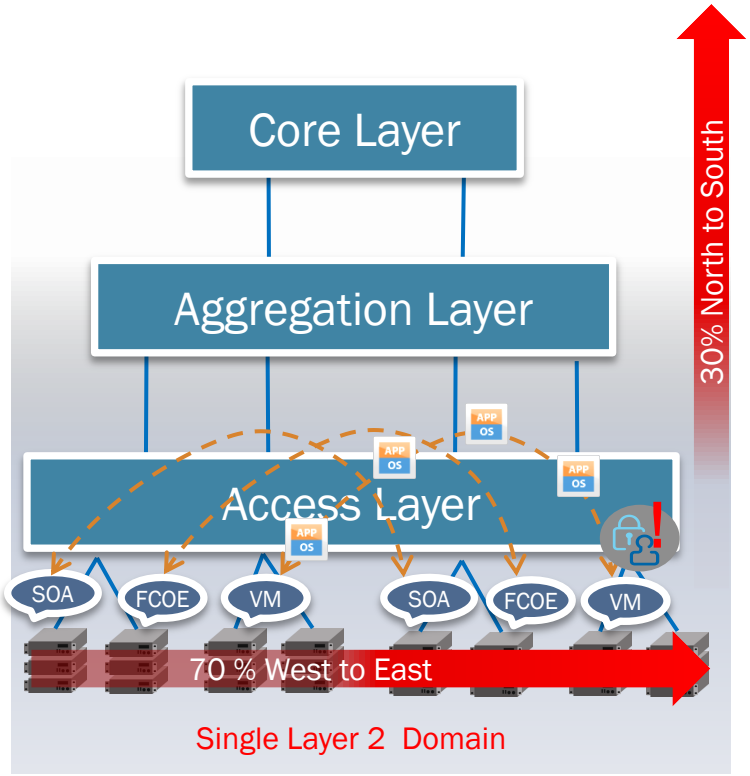

- Increased West to East traffic
  - Next Generation Apps (SOA, SAS. Web 2.0)
  - Server Virtualisation (VM)– Server to Server
  - Convergence (FCOE) – Server to Storage
- Drive for applications awareness
  - Applications the business enabler
  - DC designed around the application
  - Network needs to be aware of the apps
- The New DC needs to be flat
  - Single scalable Layer 2 Domain

Brocade One

Ethernet Fabric

The Next Generation Data Centre is a Fabric Architecture

- Fabric is a single logically flat Layer 2 network
- Wire-once, plug-in-play management
- End to end application awareness within the fabric
- Active-active, non-stop forwarding, zero downtime
- Delivering optimal any-to-any forwarding across the fabric

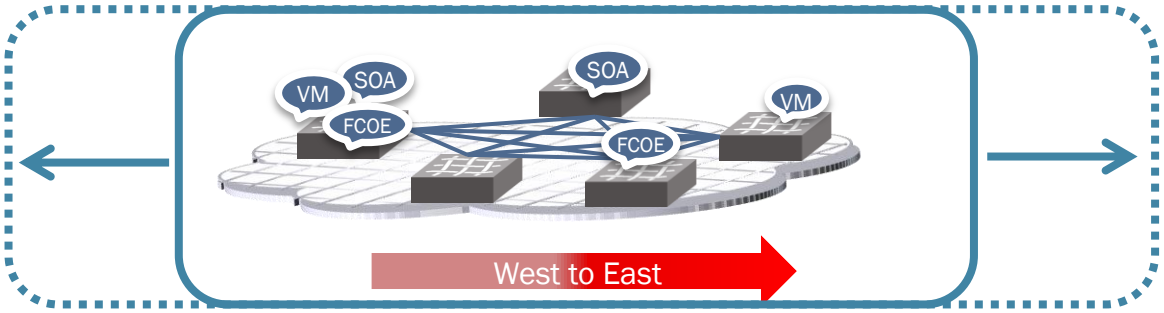

The Fabric architecture purpose built for the challenges of East-to-West traffic

Brocade One

Virtual Cluster Switching

- Brocade created the fabric architecture
  - Market leader in building SAN Fabric
  - SAN Fabric are scalable layer 2 network
  - Robust 10 year proven technology in the Data centre
  - Present in 90 of Global 100 Data Centre
- Virtual Cluster Switching
  - Brocade's Vision for the Next generation Data Centre
  - Ethernet Fabric built on Brocade's SAN fabric

VCS Evolution not Revolution

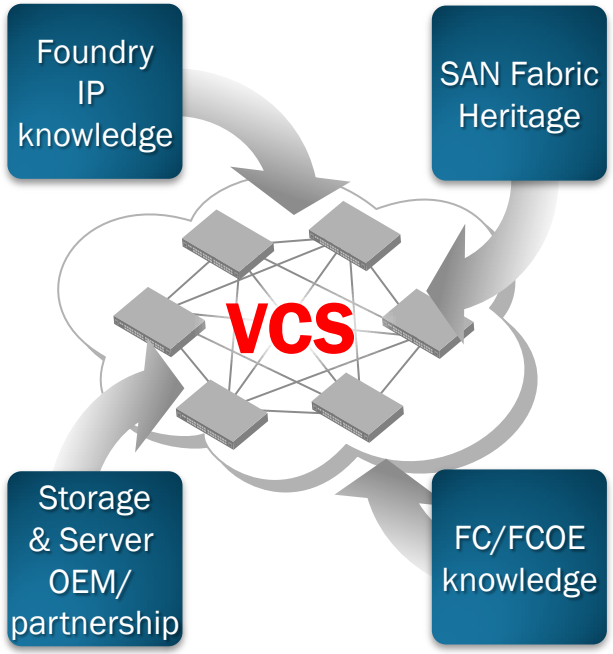

Virtual Cluster Switching

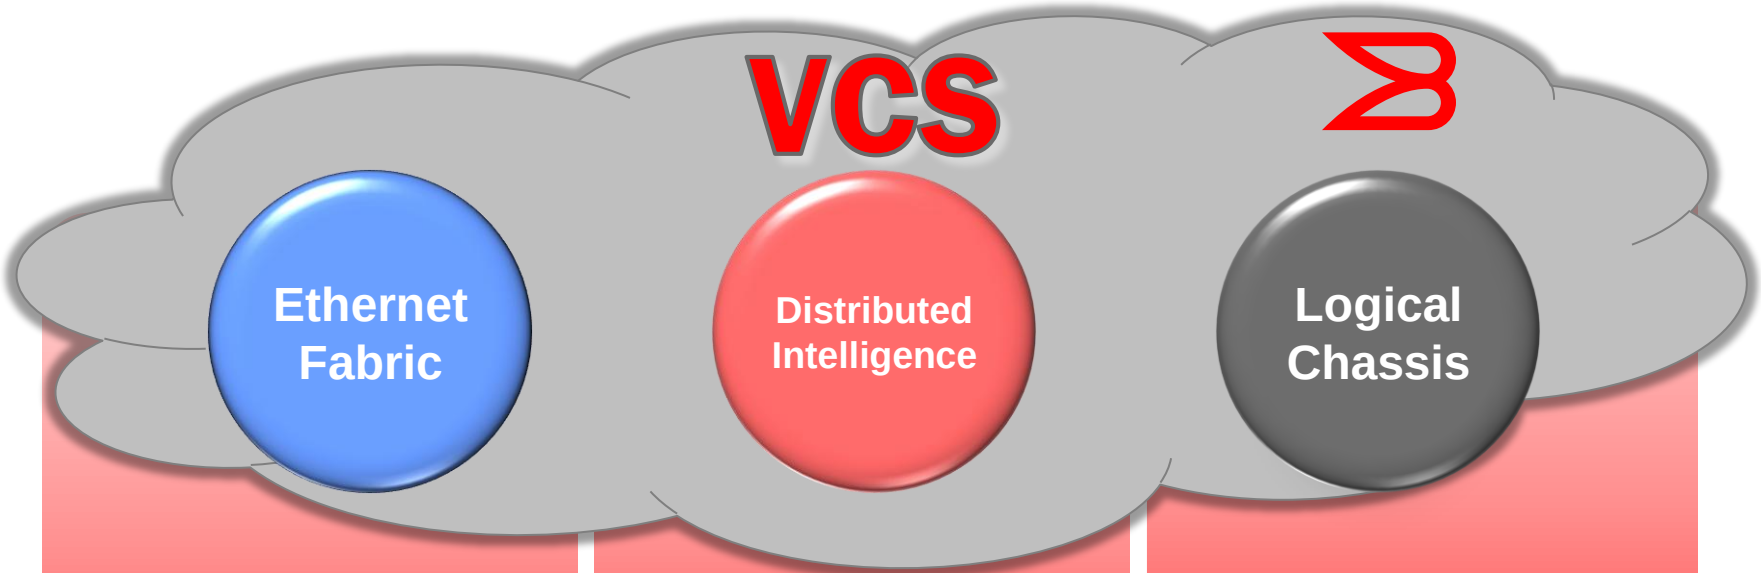

- No Spanning Tree Protocol
- Multi-path, deterministic
- Auto-healing, non-disruptive
- Lossless, low latency
- Convergence-ready

- Self-forming
- Arbitrary topology
- Fabric is aware of all members, devices, VMs
- Masterless control, no reconfiguration
- VAL interaction

- Logically flattens and collapses network layers
- Scale edge and manage as if single switch
- Auto-configuration
- Centralized or distributed mgmt; end-to-end

Dynamic Services — Connectivity over Distance, Native Fibre Channel, Security Services, Layer 4-7, etc.

BROCADE

ETHERNET FABRICS

Scaling Virtual Server Environments

Technical Challenges Today

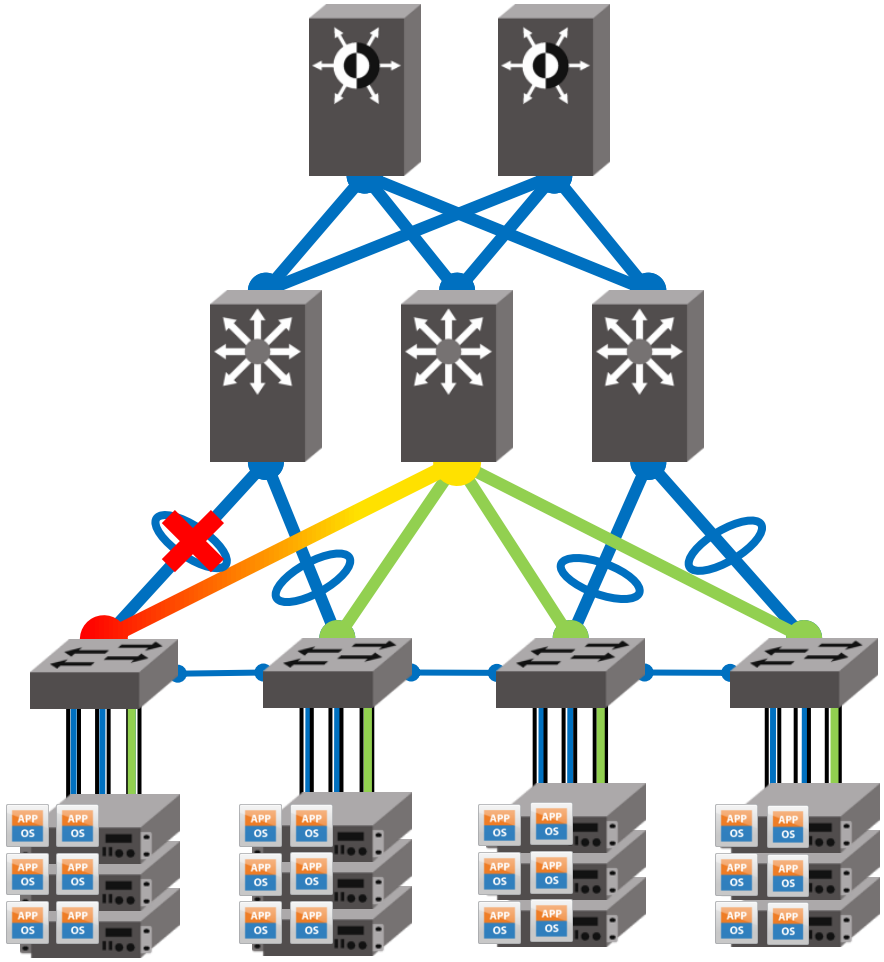

- Layer 2: only 1 active path
  - STP disables other paths
  - Not “virtualization optimized”
- Add Virtual Machines
  - Add additional GbE connections
  - Move to 10 GbE for simplicity and more performance
  - Uplinks are stressed; need more connections in LAG
- Increase utilization using MSTP (spanning tree per VLAN)
  - Increases complexity
  - Creates multiple single-path networks; limits sphere of mobility
- Link failure
  - STP reconvergence – network is down
  - Broadcast storms stress network
- Layer 3 as an alternative
  - Greater complexity; higher cost
  - VM mobility limited to rack

Adding a new switch to an Ethernet Fabric

Automatic Fabric Creation and Expansion

Automatic Trunk Creation

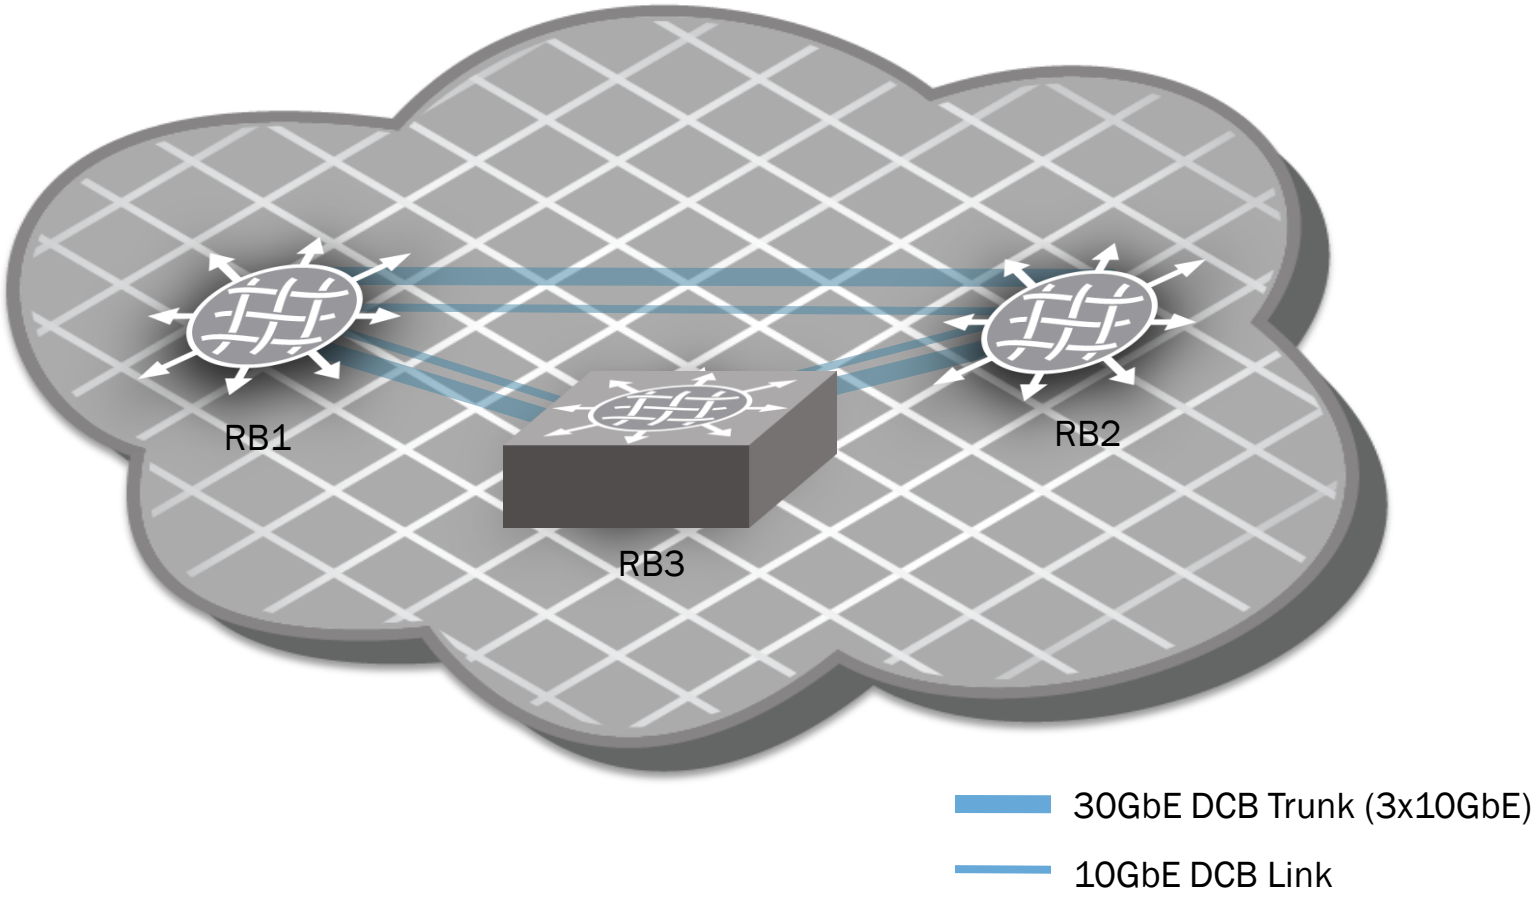

CLI Callouts

New VDX Switch added to cluster

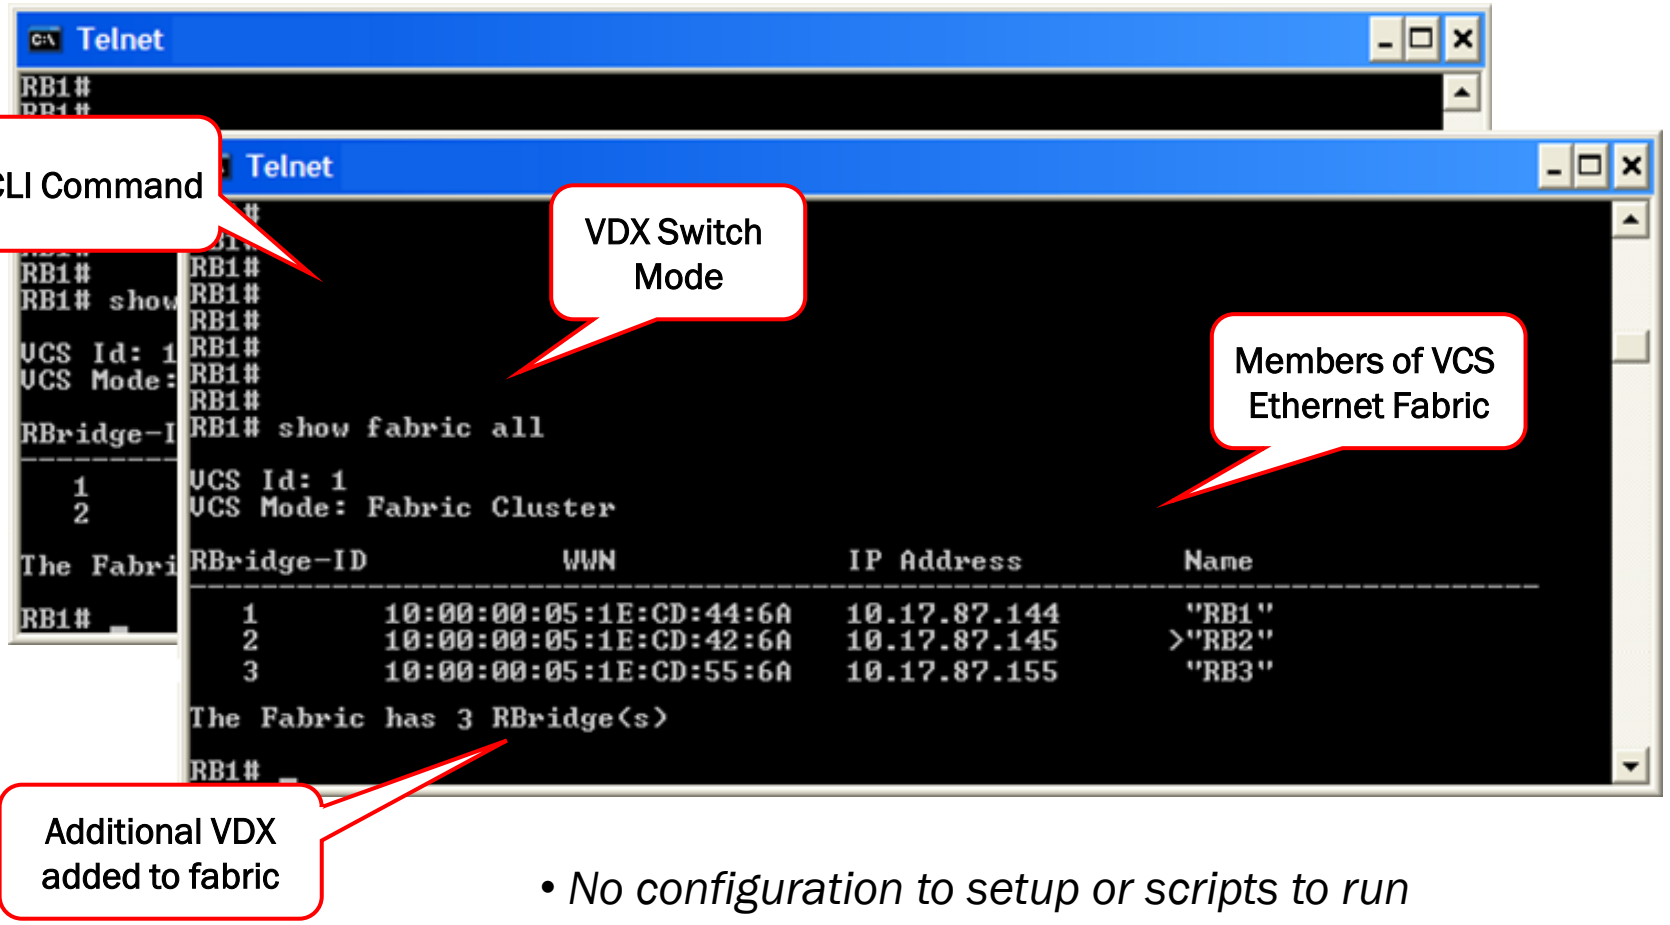

- *No configuration to setup or scripts to run*
- *Existing traffic is uninterrupted*

Ethernet Fabrics

Distributed MAC address tables

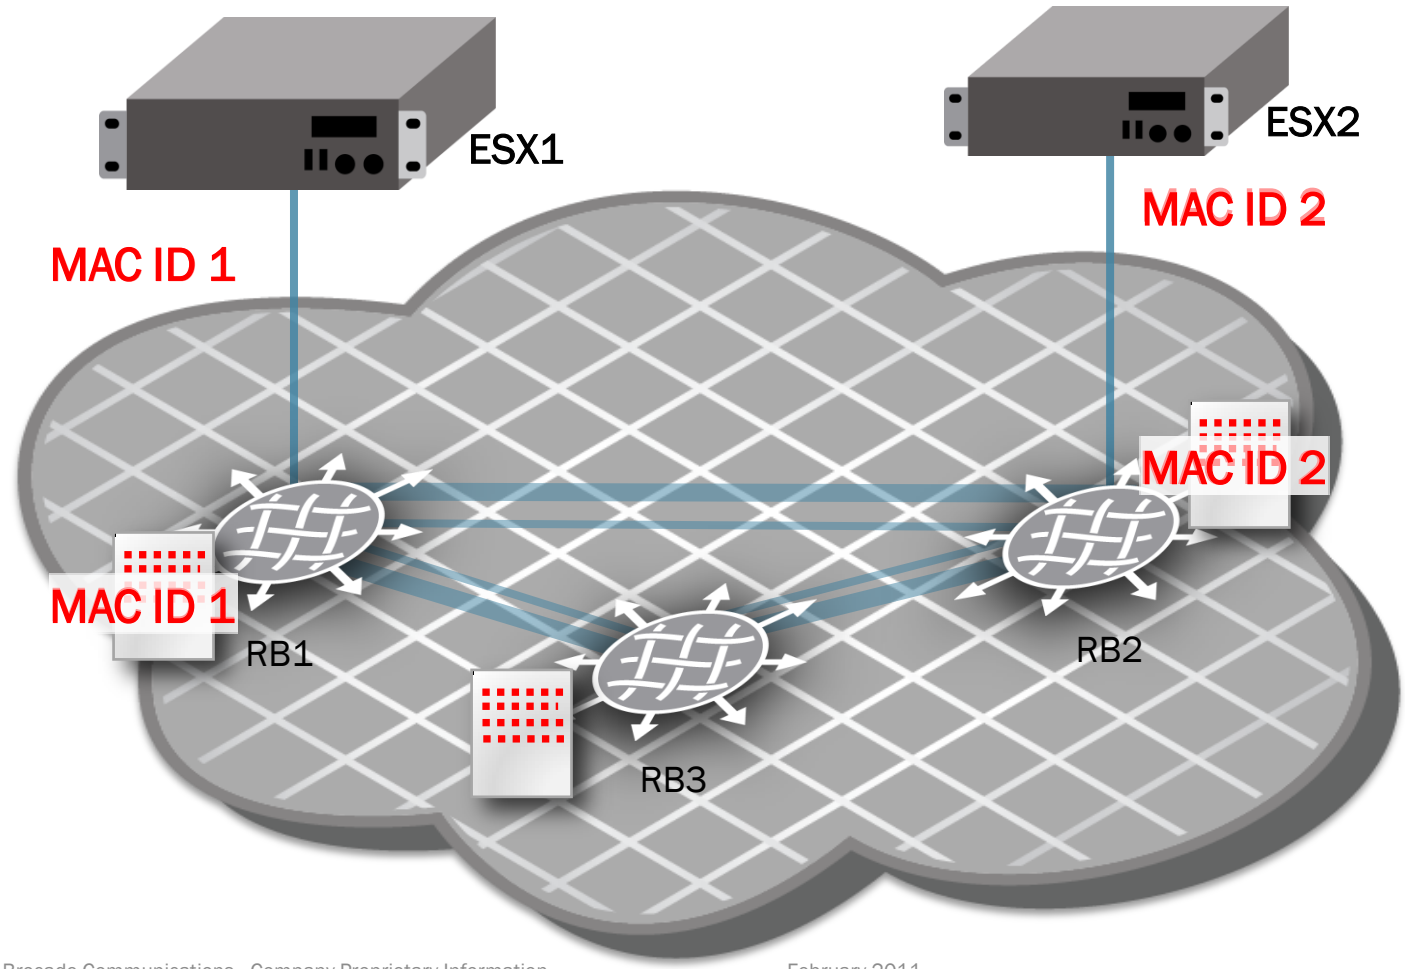

CLI Callouts

MAC Learning

CLI Command

Learned MAC addresses

Members of VCS Ethernet Fabric

See all devices on all switches

Switch and Port MAC Learned

Telnet 10.17.87.144

```
VCS Mode: Fabric Cluster

RBridge-ID          WWN               IP Address        Name
-----------------------------------------------------------------------
   1        10:00:00:05:1E:CD:44:6A   10.17.87.144       "RB1"
   2        10:00:00:05:1E:CD:42:6A   10.17.87.145      >"RB2"
   3        10:00:00:05:1E:CD:55:6A   10.17.87.155       "RB3"

The Fabric has 3 RBridge(s)

RB1# show mac-address-table
VlanId    Mac-address         Type      State      Ports
1         0005.1ed9.f024      Dynamic   Active     Te 2/0/16
1         0005.3326.574b      Dynamic   Active     Te 2/0/6
1         0005.3326.57e2      Dynamic   Active     Te 1/0/3
1         0010.9400.9910      Dynamic   Active     Te 2/0/11
1         0010.9400.9999      Dynamic   Active     Te 1/0/4
1         0015.1515.0ccc      Dynamic   Active     Te 2/0/16
1         0050.5678.cd2e      Dynamic   Active     Te 1/0/6
1         0050.5679.9c60      Dynamic   Active     Te 2/0/4
1         0050.56bf.0000      Dynamic   Active     Te 2/0/4
1         0050.56bf.0005      Dynamic   Active     Te 1/0/6
1         0050.56bf.0007      Dynamic   Active     Te 1/0/6
1002      0efc.0001.2900         FPMA   Active     Te 1/0/17
Total MAC addresses     :  12
RB1#
```

Telnet 10.17.87.145

```
VCS Mode: Fabric Cluster

RBridge-ID          WWN               IP Address        Name
-----------------------------------------------------------------------
   1        10:00:00:05:1E:CD:44:6A   10.17.87.144       "RB1"
   2        10:00:00:05:1E:CD:42:6A   10.17.87.145      >"RB2"
   3        10:00:00:05:1E:CD:55:6A   10.17.87.155       "RB3"

The Fabric has 3 RBridge(s)

RB2# show mac-address-table
VlanId    Mac-address         Type      State      Ports
1         0005.1ed9.f024      Dynamic   Active     Te 2/0/16
1         0005.3326.574b      Dynamic   Active     Te 2/0/6
1         0005.3326.57e2      Dynamic   Active     Te 1/0/3
1         0010.9400.9910      Dynamic   Active     Te 2/0/11
1         0015.1515.0aaa      Dynamic   Active     Te 2/0/16
1         0015.1515.0ccc      Dynamic   Active     Te 2/0/16
1         0050.5678.cd2e      Dynamic   Active     Te 1/0/6
1         0050.5679.9c60      Dynamic   Active     Te 2/0/4
1         0050.56bf.0000      Dynamic   Active     Te 2/0/4
1         0050.56bf.0001      Dynamic   Active     Te 2/0/4
1         0050.56bf.0005      Dynamic   Active     Te 1/0/6
1         0050.56bf.0007      Dynamic   Active     Te 1/0/6
1002      0efc.0002.1e00         FPMA   Active     Te 2/0/6
Total MAC addresses     :  13
RB2#
```

- *New servers are seen on all switches*

Ethernet Fabrics

Host Based LAG – LAG to different physical VDX Switches

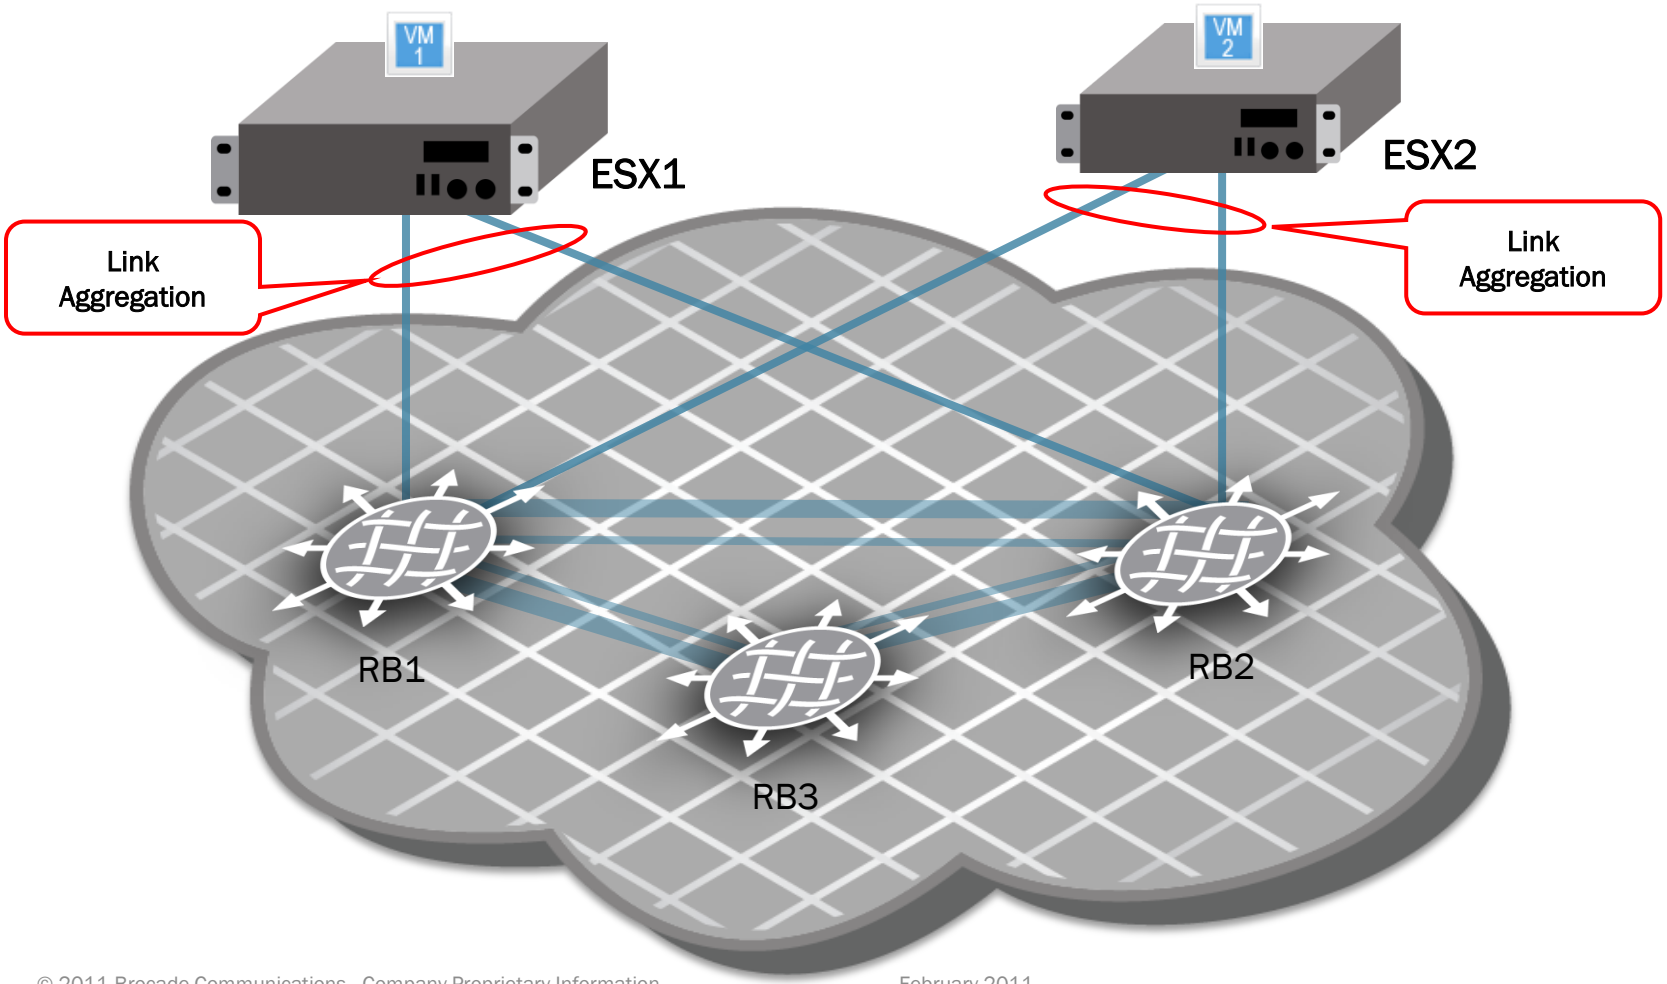

CLI Callouts

Fabric Topology

CLI Command

ISL Cost

ISL Bandwidth

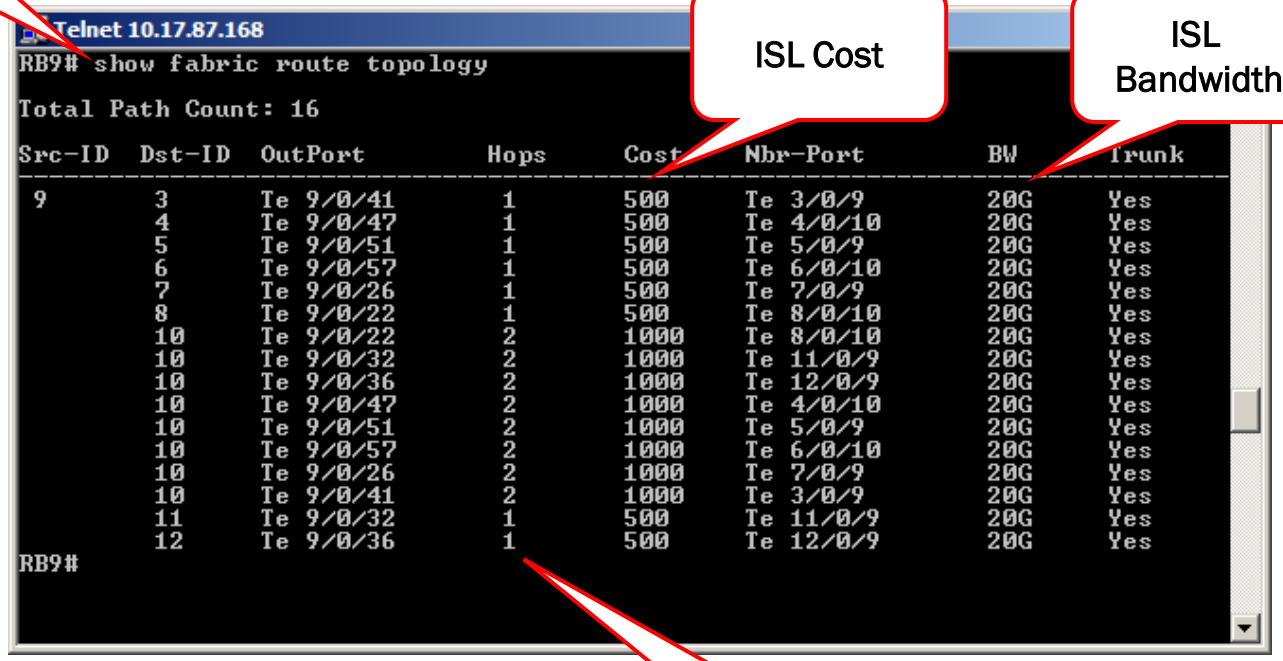

Hop Count to VCS Member Switch

Traditional Link Aggregation

Switch-to-Switch LAG – Hashing algorithms cause imbalances

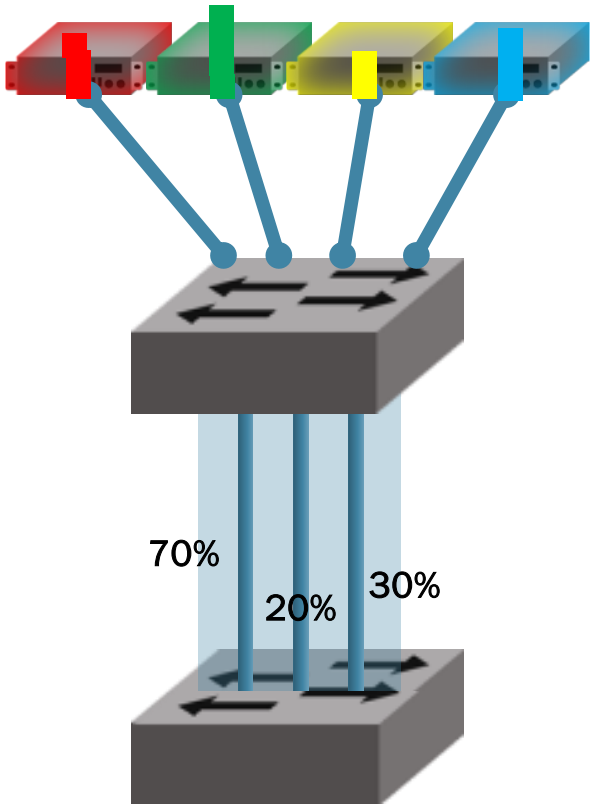

Virtual Link Aggregation (Trunking)

Balanced Connections with no hashing algorithms

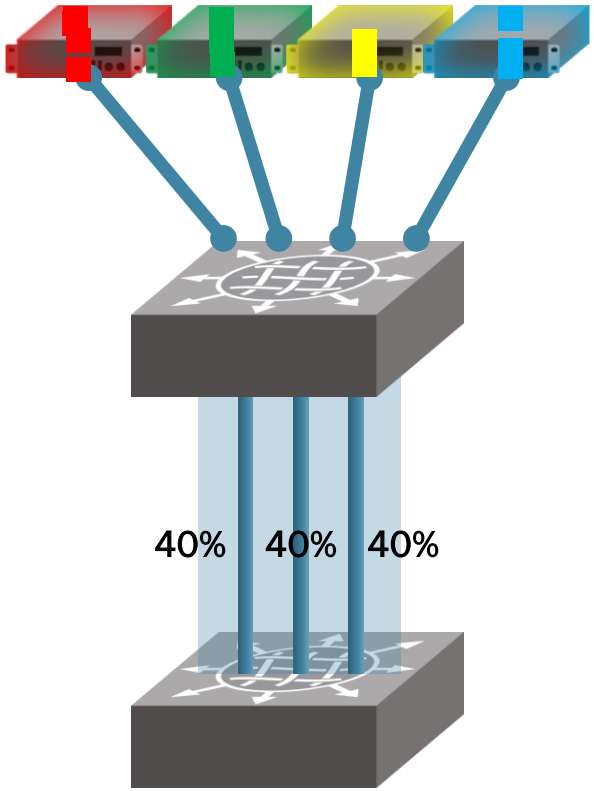

Ethernet Fabrics

Equal Cost Multi-Pathing

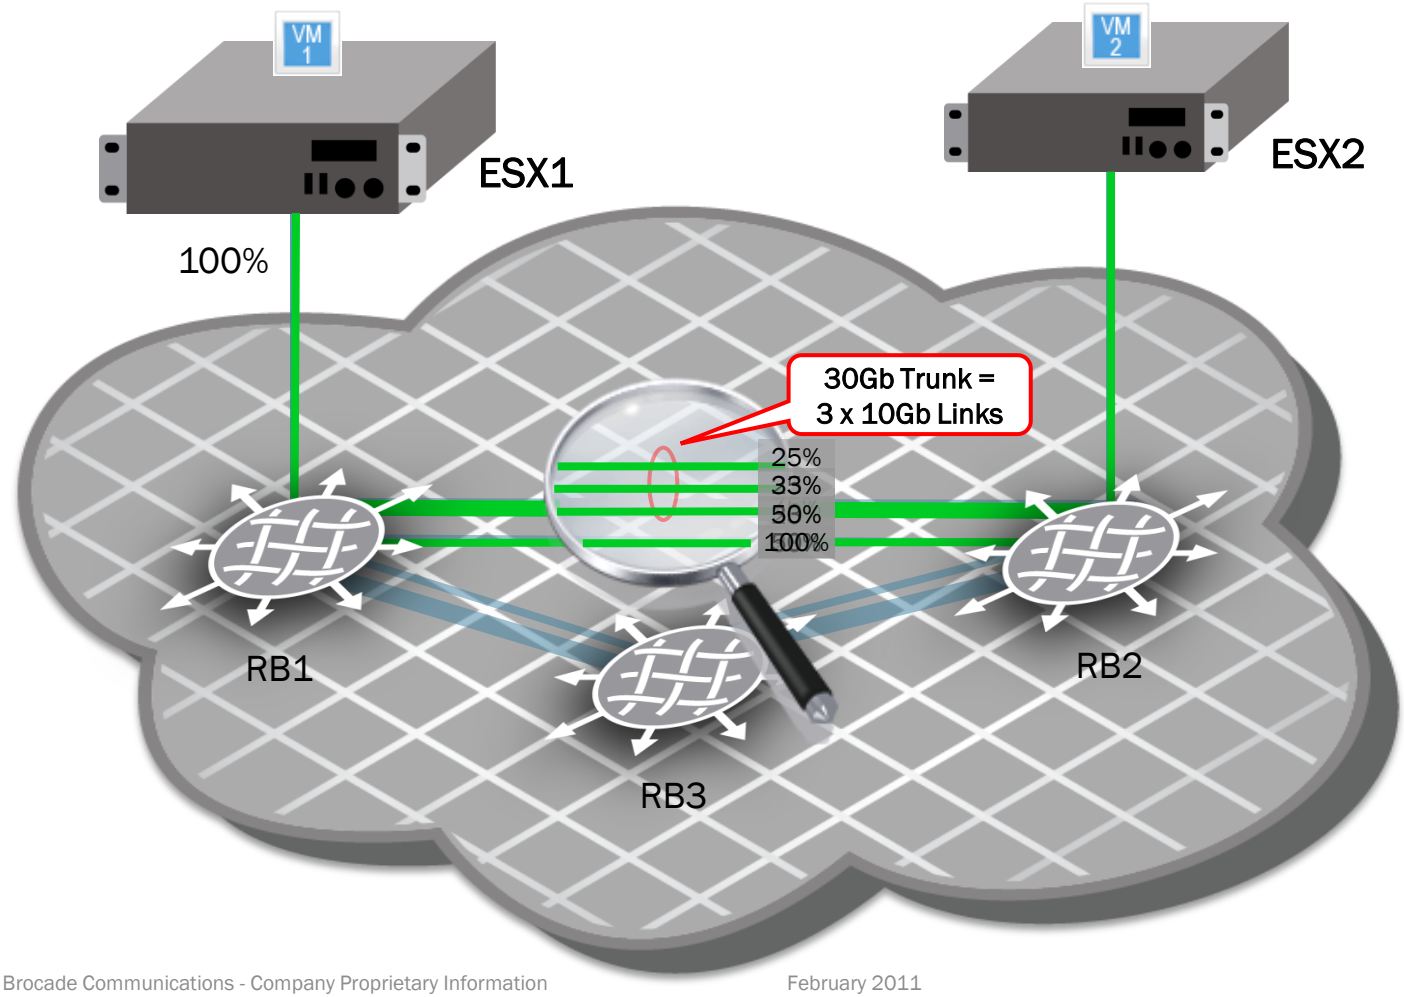

CLI Callouts

ISL Bandwidth

CLI Command

Ports Participating In ISL

ISL Bandwidth

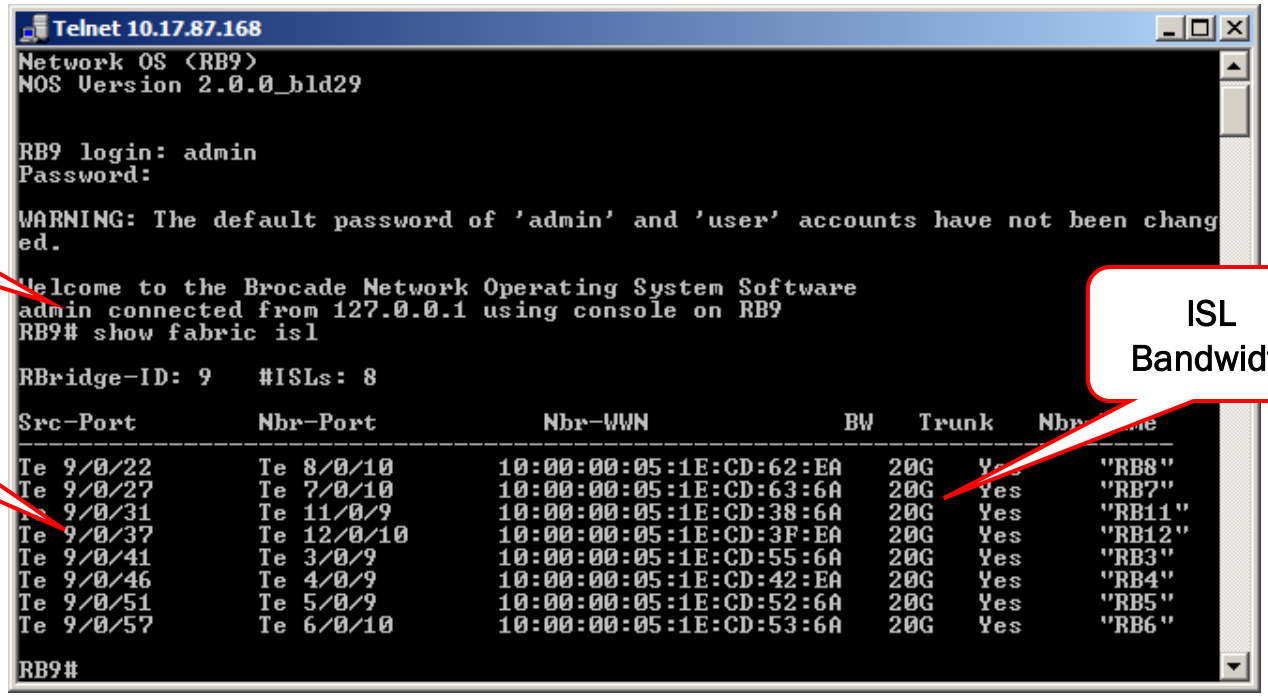

BROCADE

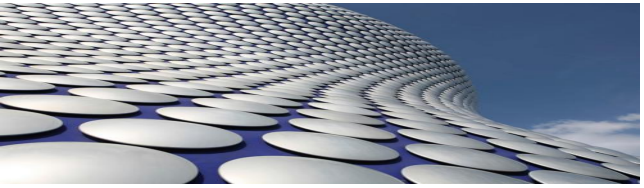

VIRTUAL MOBILITY

Automatic Migration Port Profiles

Technical Challenges of Migrating Virtual Machines

Network Complexity

Security and Policy Enforcement

VM MAC Association

Support for multiple hypervisors

Virtual Machine Mobility

Challenges Today

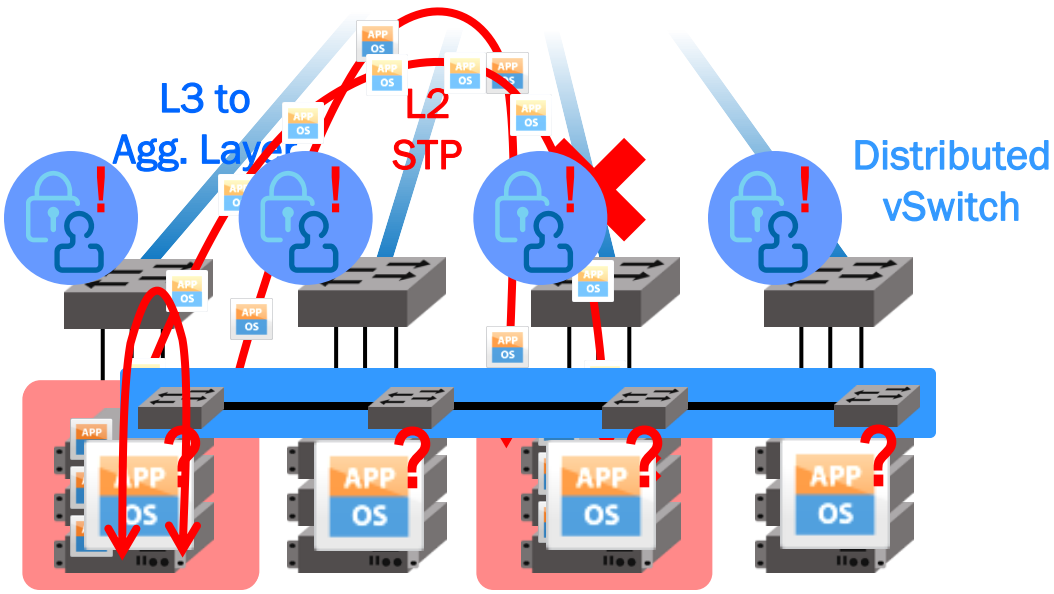

- Limited sphere of mobility
  - STP limits flexibility to a minimized, defined tree of switches
  - L3 limits mobility to a single rack
- VM migration can break network/application access
  - Port setting information must be identical at destination
- Map services (VLANs, QoS, security, etc.) to all physical ports
  - Eases mobility, but undermines network and security best practices
- Distributed Virtual Switch
  - Addresses configuration needs
  - Consumes server resources and still restricted by physical limits
- Limited insight into where VMs are running
  - VMs exist anywhere in the cluster

What Does It Mean To Virtualize The Network?

- Automatically distribute and bind policies at a per-VM level
- Allow seamless VM migration across the network
- Enforce VM level policies in a consistent fashion across the network
- Be Hypervisor agnostic to meet project requirements

Sharing Port Profiles

Automatic Sharing to simplify management

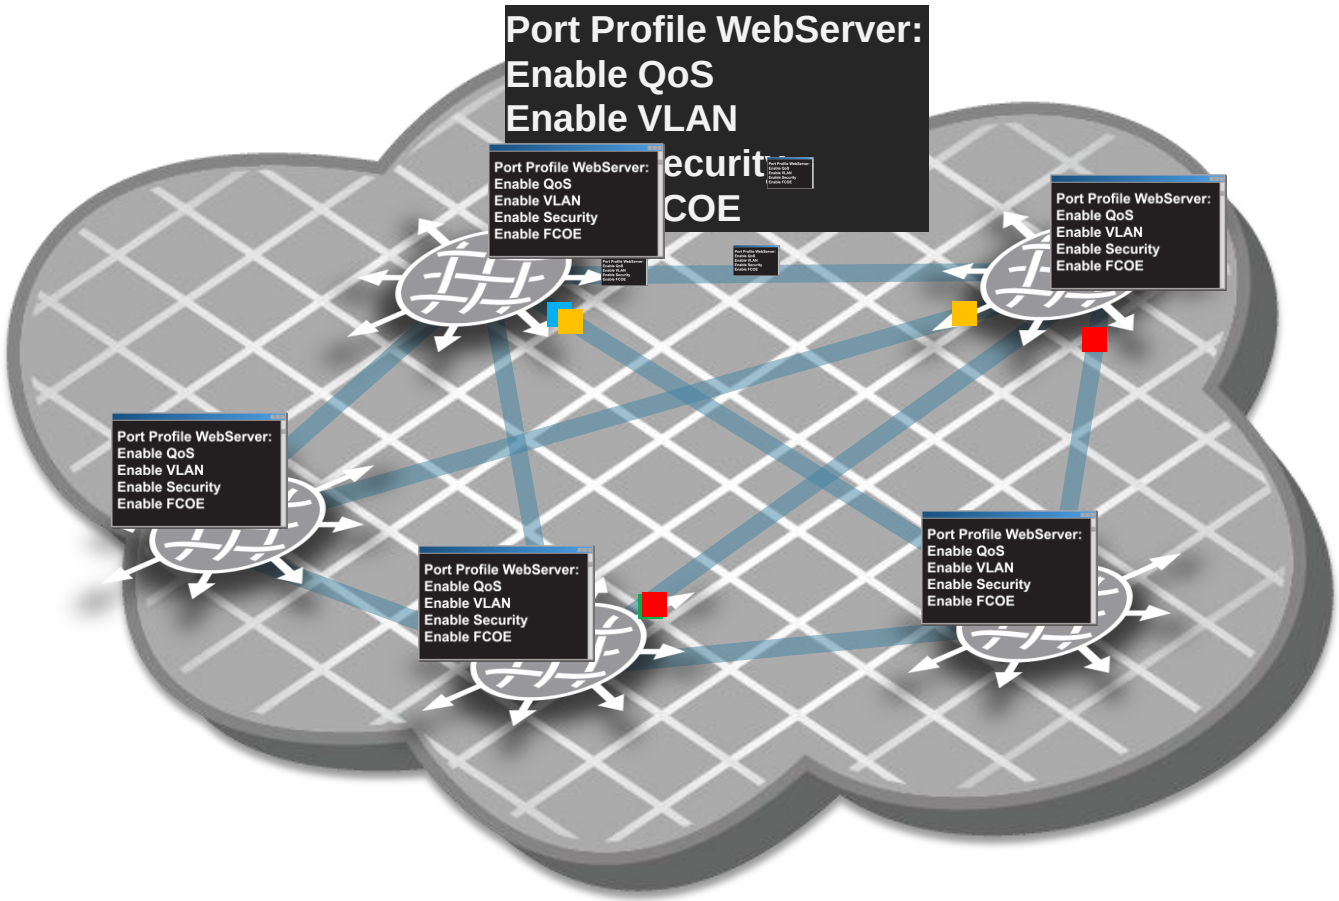

Simplified Virtual Machine Migration

Automatic Migration of Port Profiles

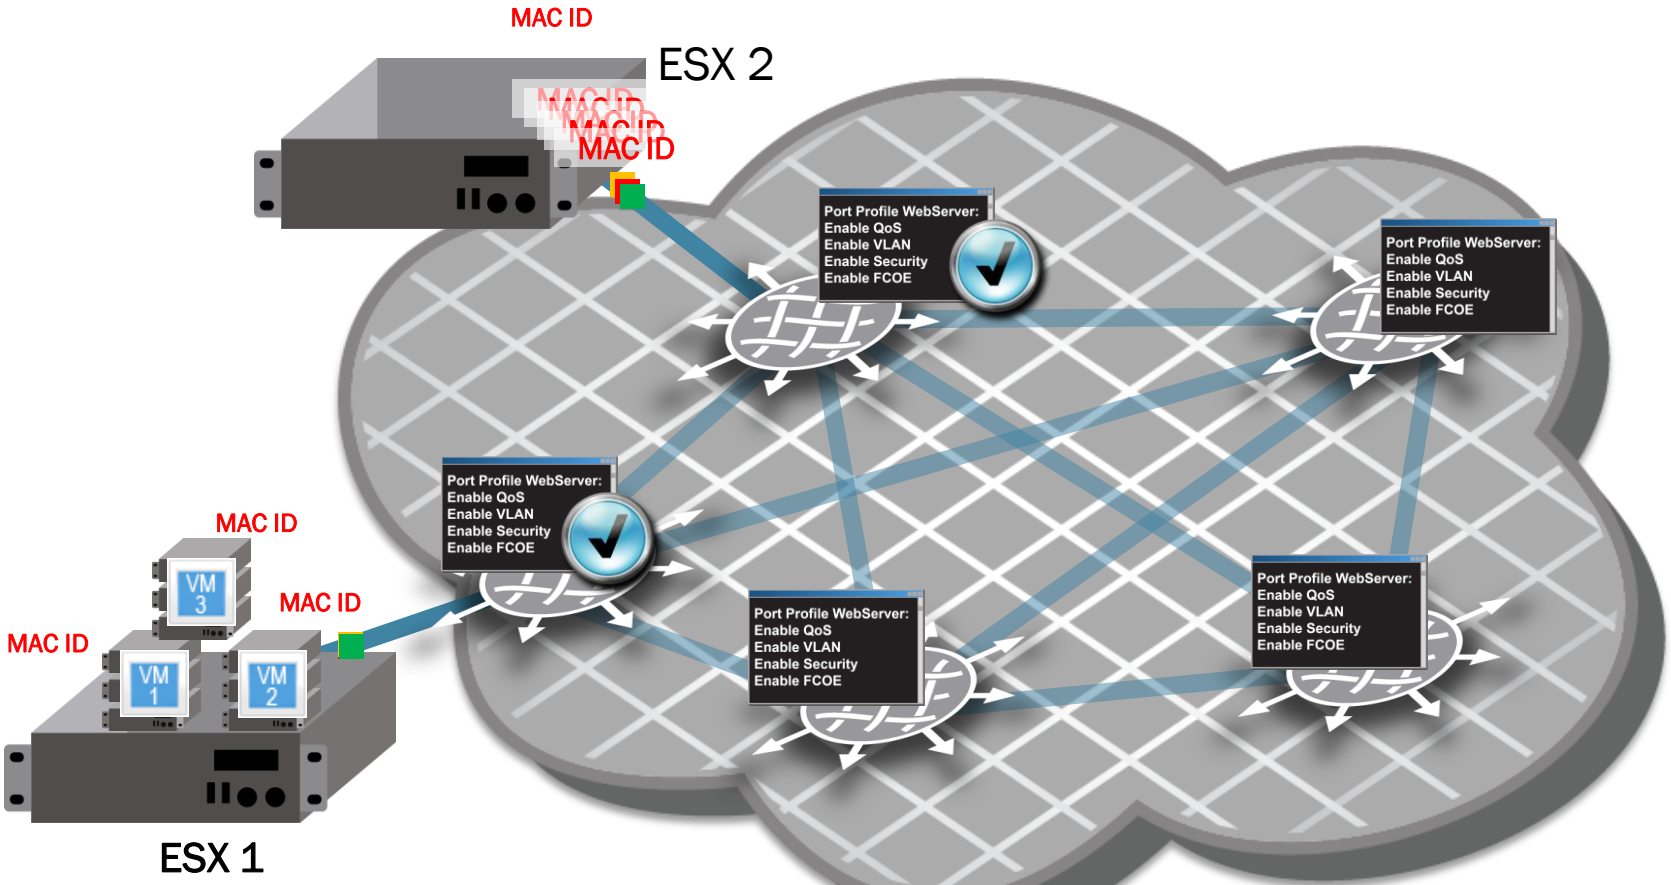

Hypervisor Agnostic

Tied only to VM Mac Addresses

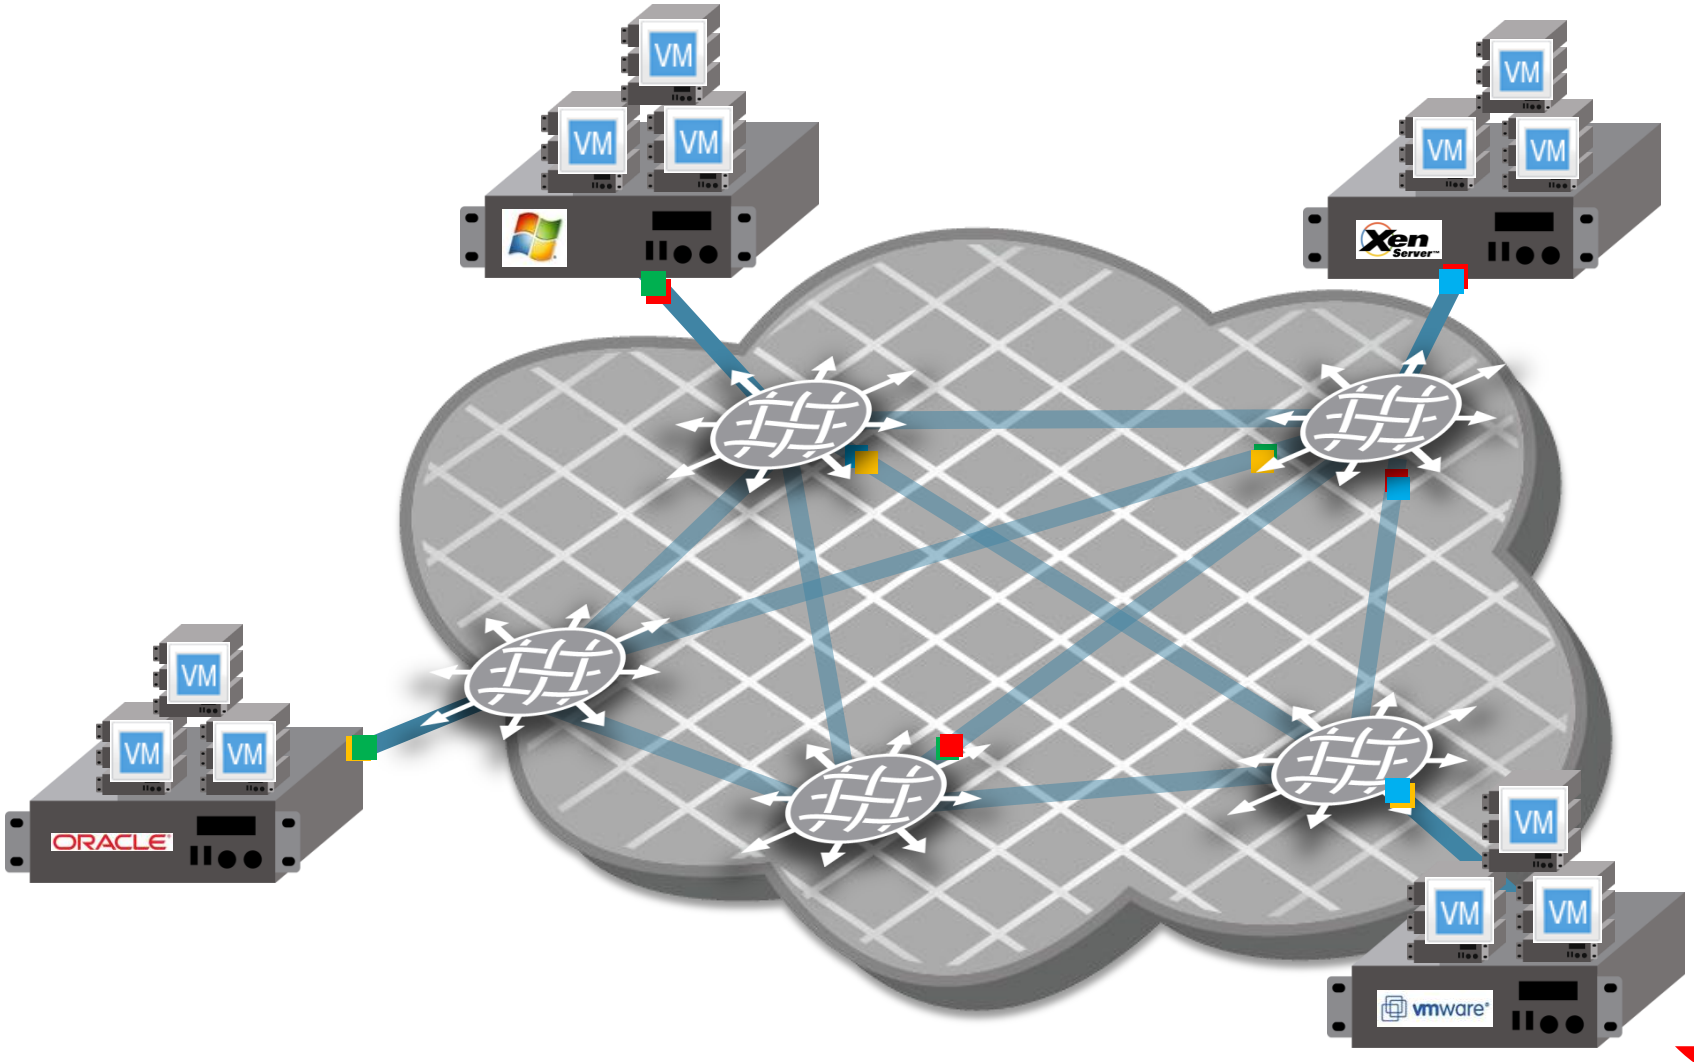

BROCADE

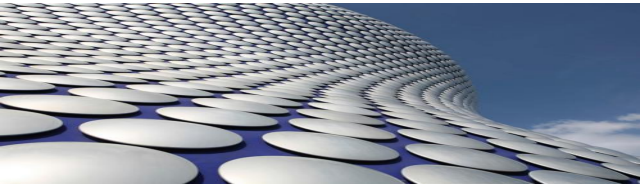

ELIMINATING BOTTLENECKS & INCREASING RELIABILITY

Multi-Chassis Trunking (MCT) & vLAG

The Network Challenge

- Better Performance
- Better Scalability
- Fewer Disruptions
- Fewer Path Failures
- Fewer Hardware Malfunctions

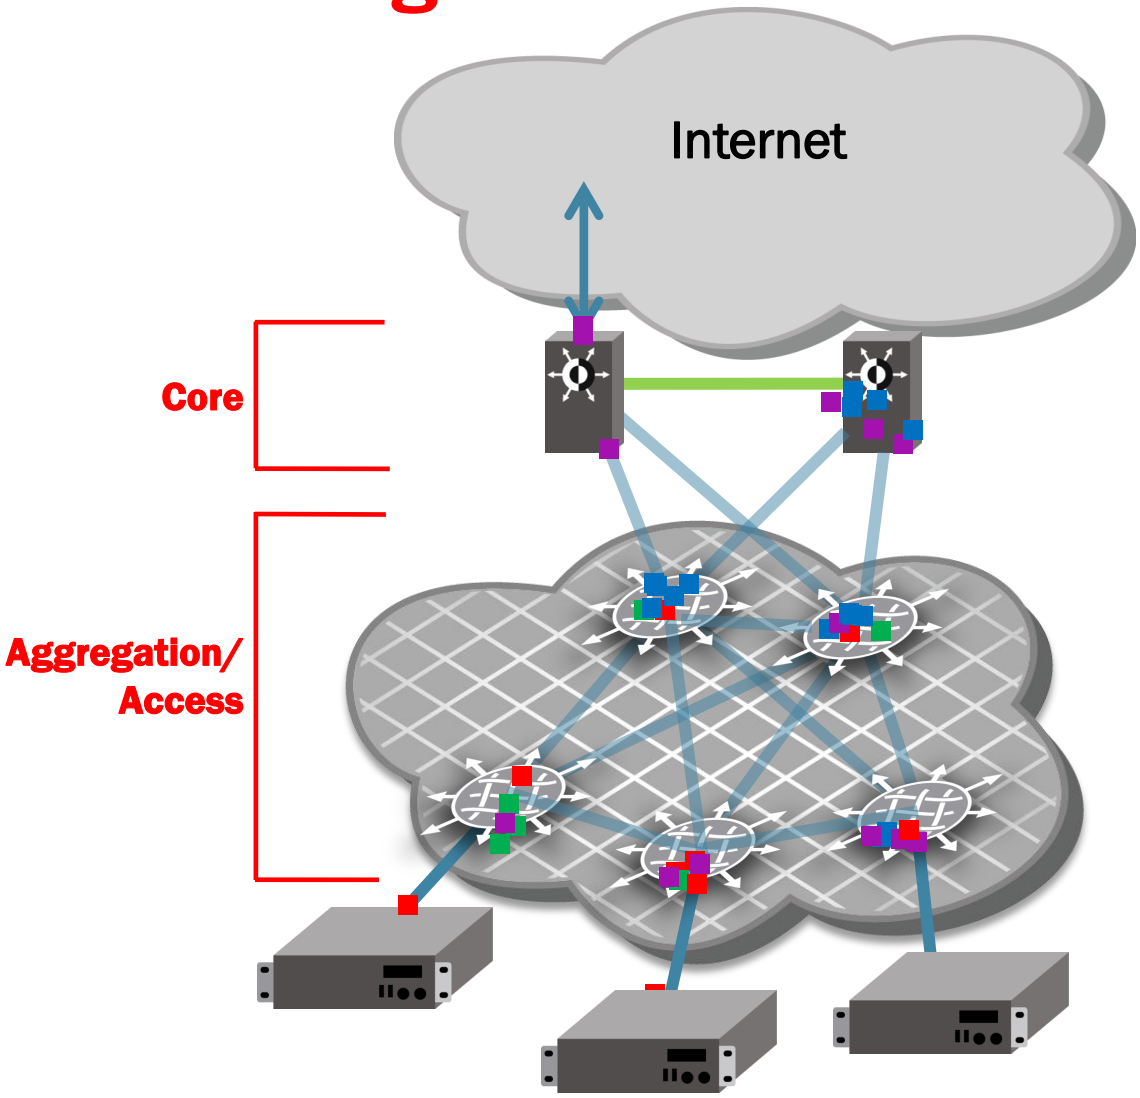

Multi-Chassis Trunking & vLAG

- High-Speed Links
- Active-active load balancing
- Faster recovery
- Interoperable with existing networks

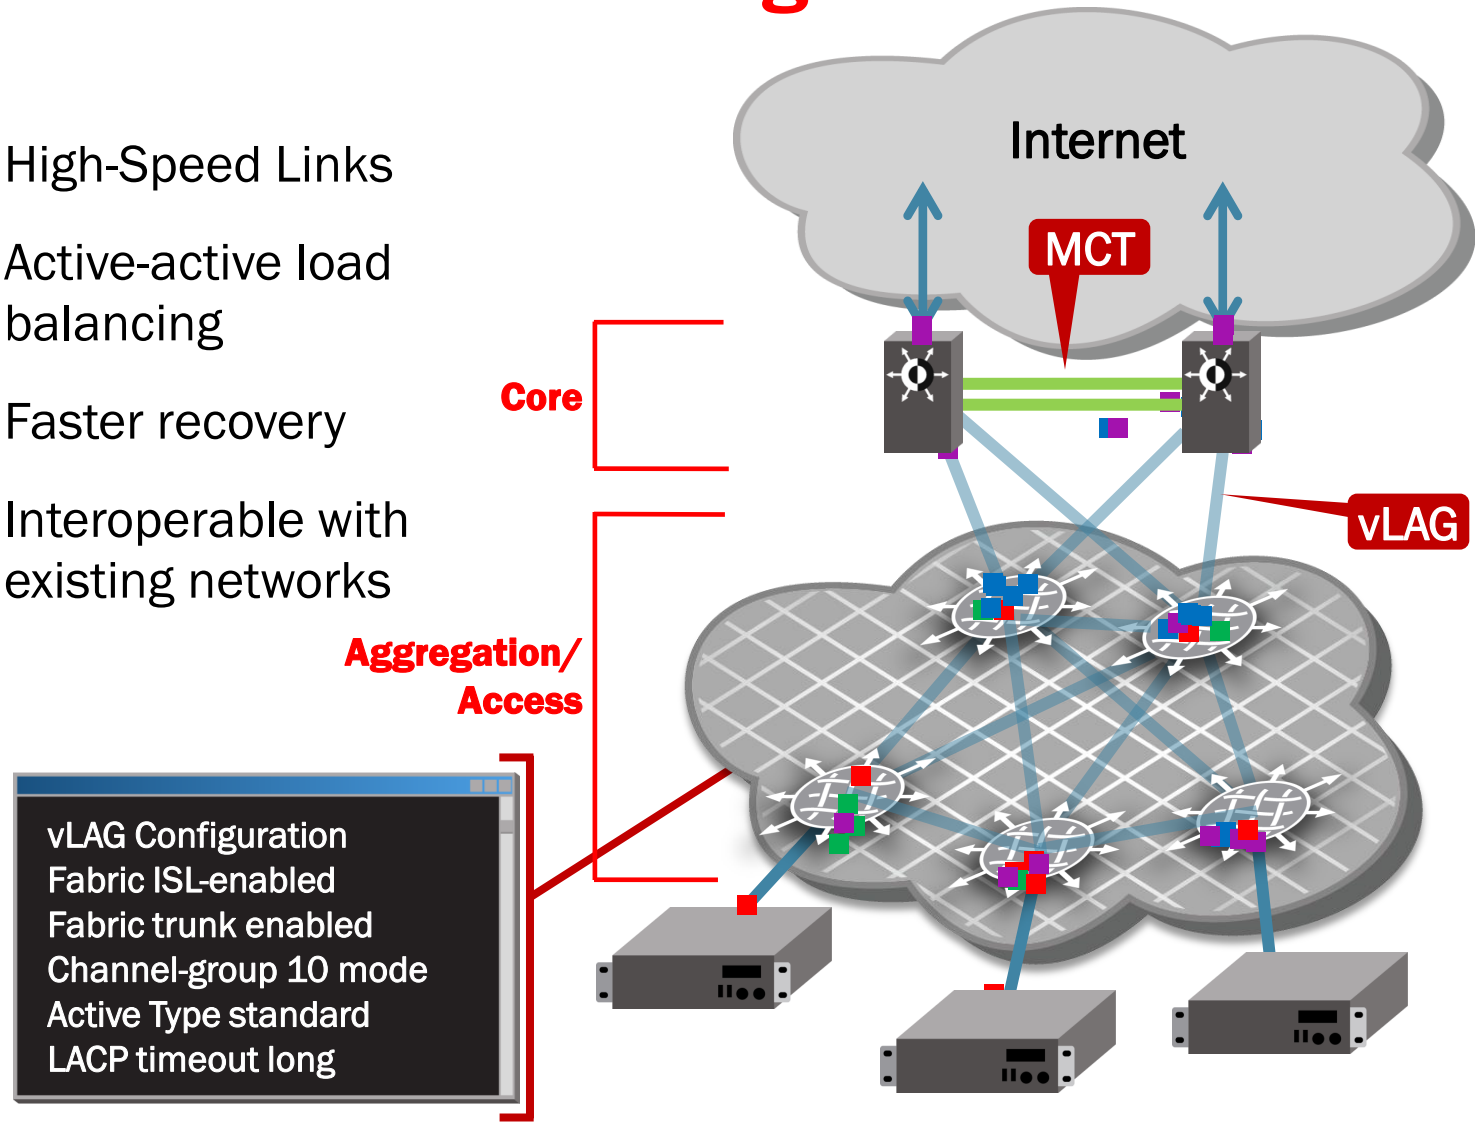

Simplified Logical View

- Handle traffic spikes
- Greater Scalability
- Reduced downtime
- Single logical access/agg layer to manage
- Single core configuration

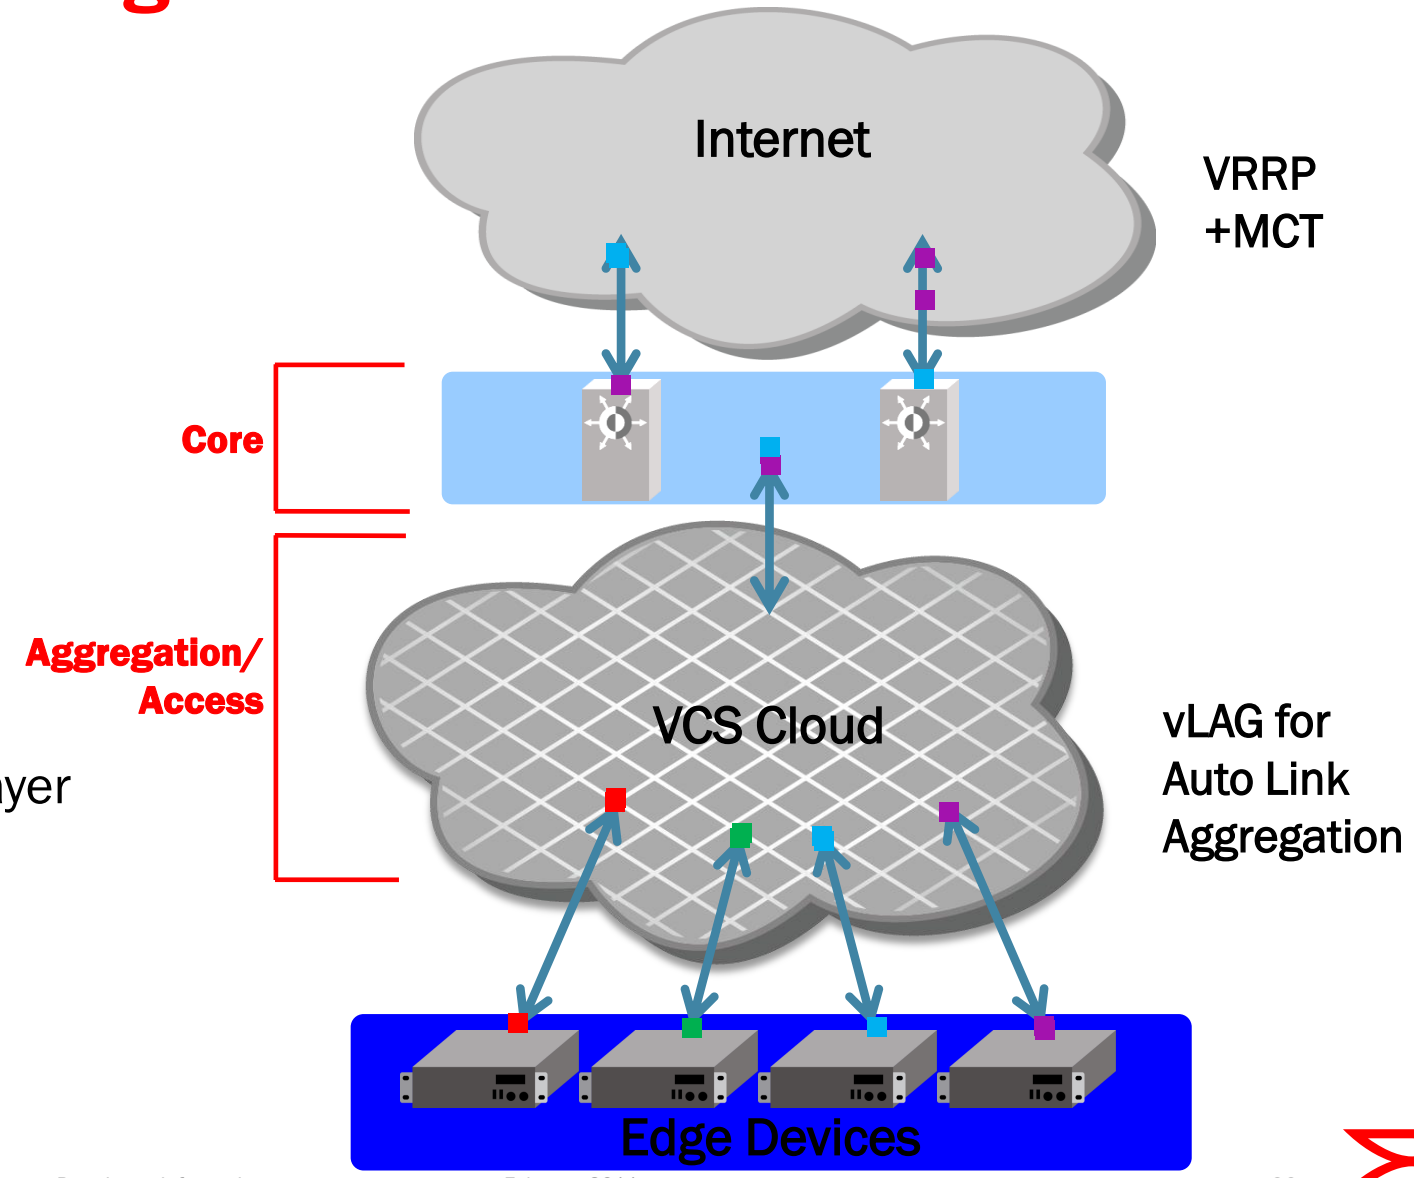

BROCADE

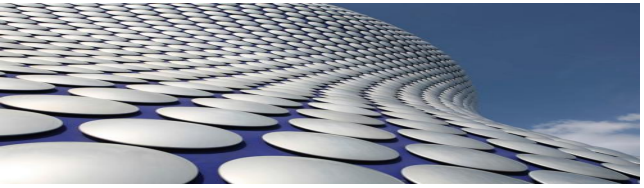

BRINGING FCoE INTO THE ETHERNET FABRIC

You don’t know anyone like this..

do you?

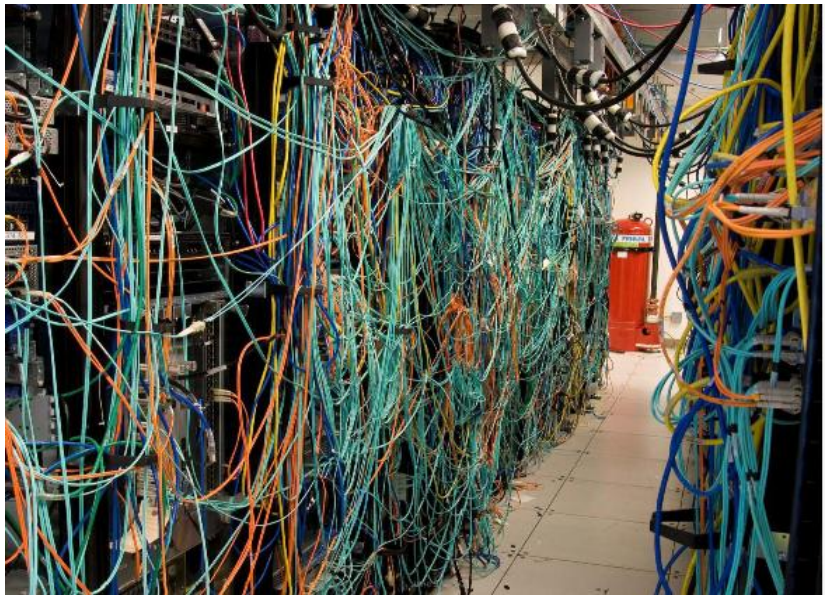

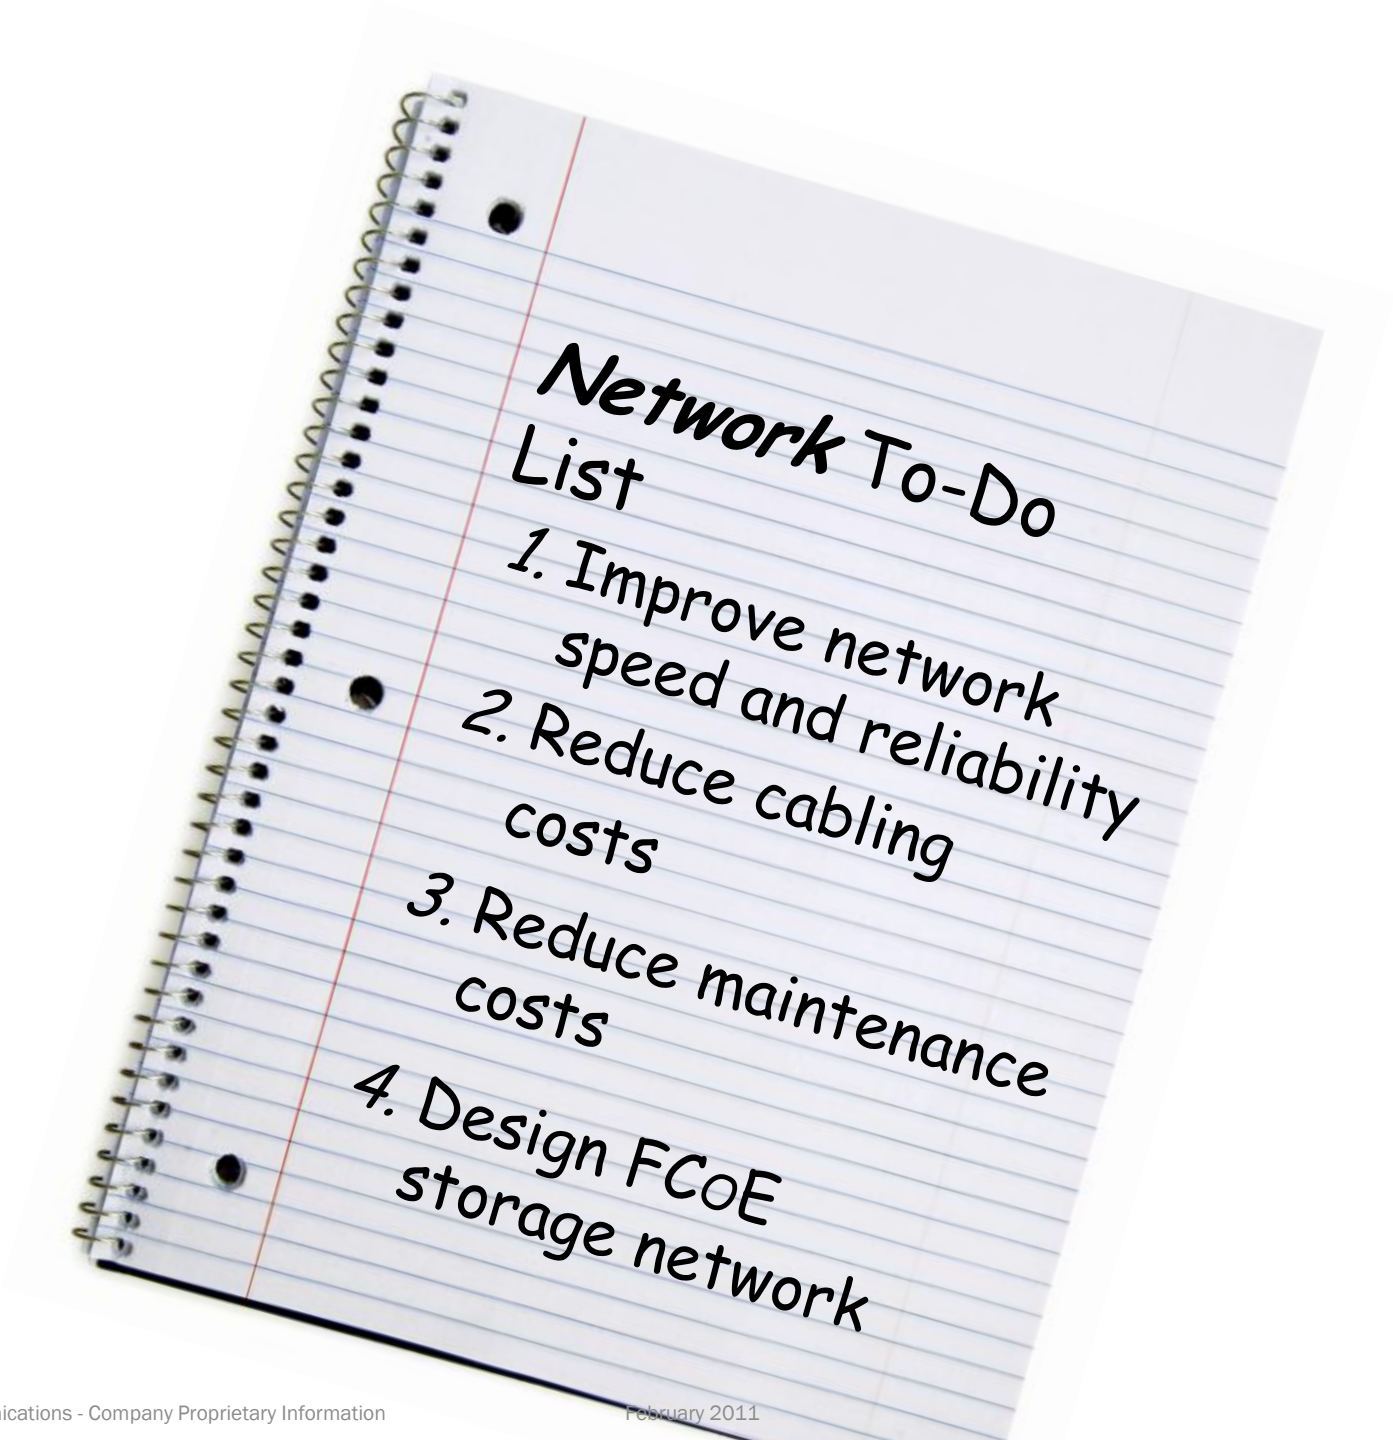

Challenges of FCoE & Convergence

Even if not today, be prepared and infrastructure ready!

- FCoE has been limited to single hop configurations on all DCB products…until Virtual Cluster Switching
  - Ethernet fabrics provide meet the requirements for multi-hop storage connectivity including low latency, lossless and ECMP
  - The network MUST be storage aware, and Brocade knows storage
  - Secure roles & responsibilities for LAN & SAN administrators
- Improves network performance and reliability for iSCSI, NAS and backup traffic as well
- Upcoming bridging for existing SANs & FC storage
  - Preserves storage investment but allows converged server connectivity to simplify the access layer

What really changes with Convergence?

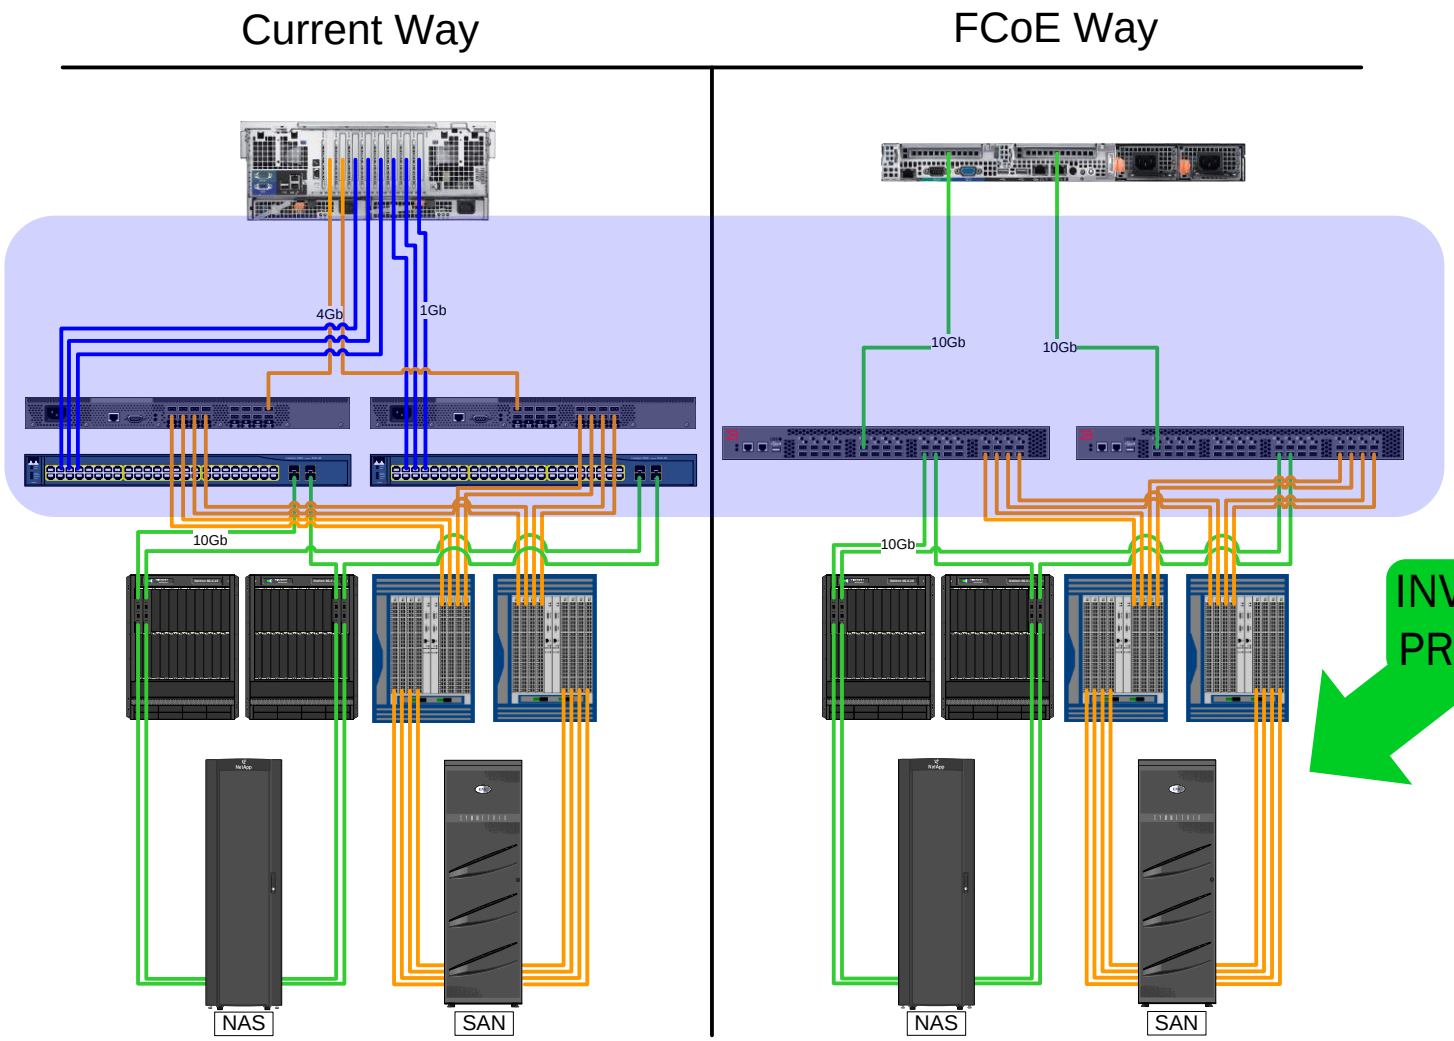

Example Environment
6 IP Connections and 2 FC per Server

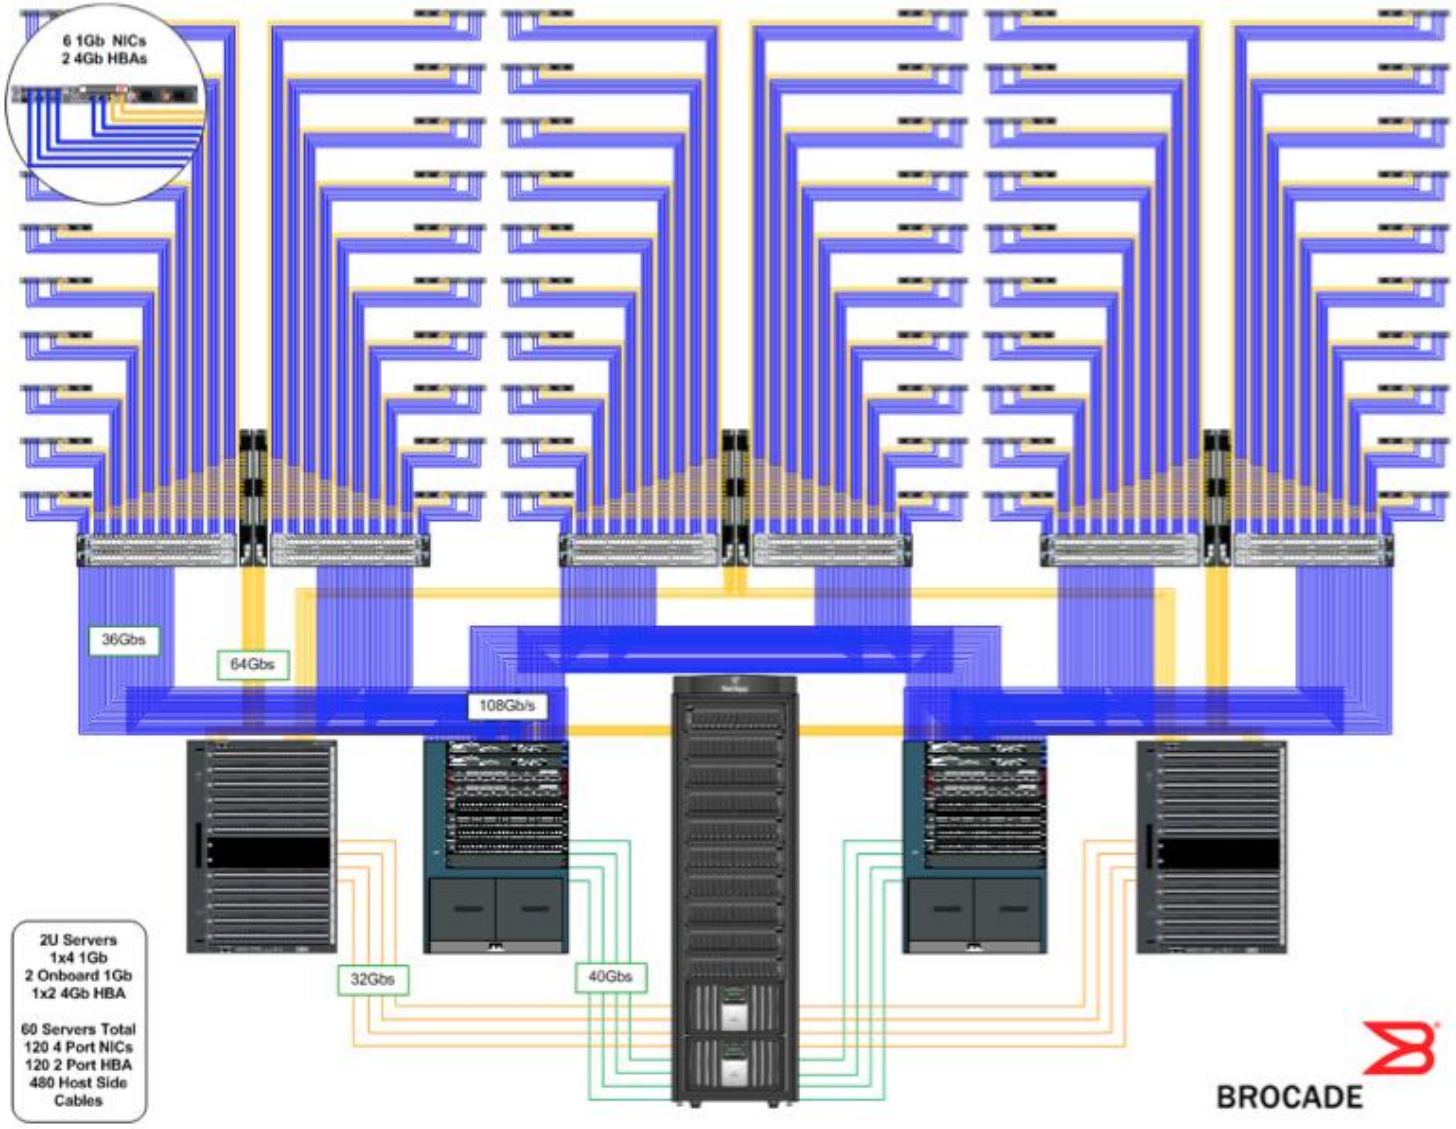

Same Example Environment
360 Host Side Cables Removed

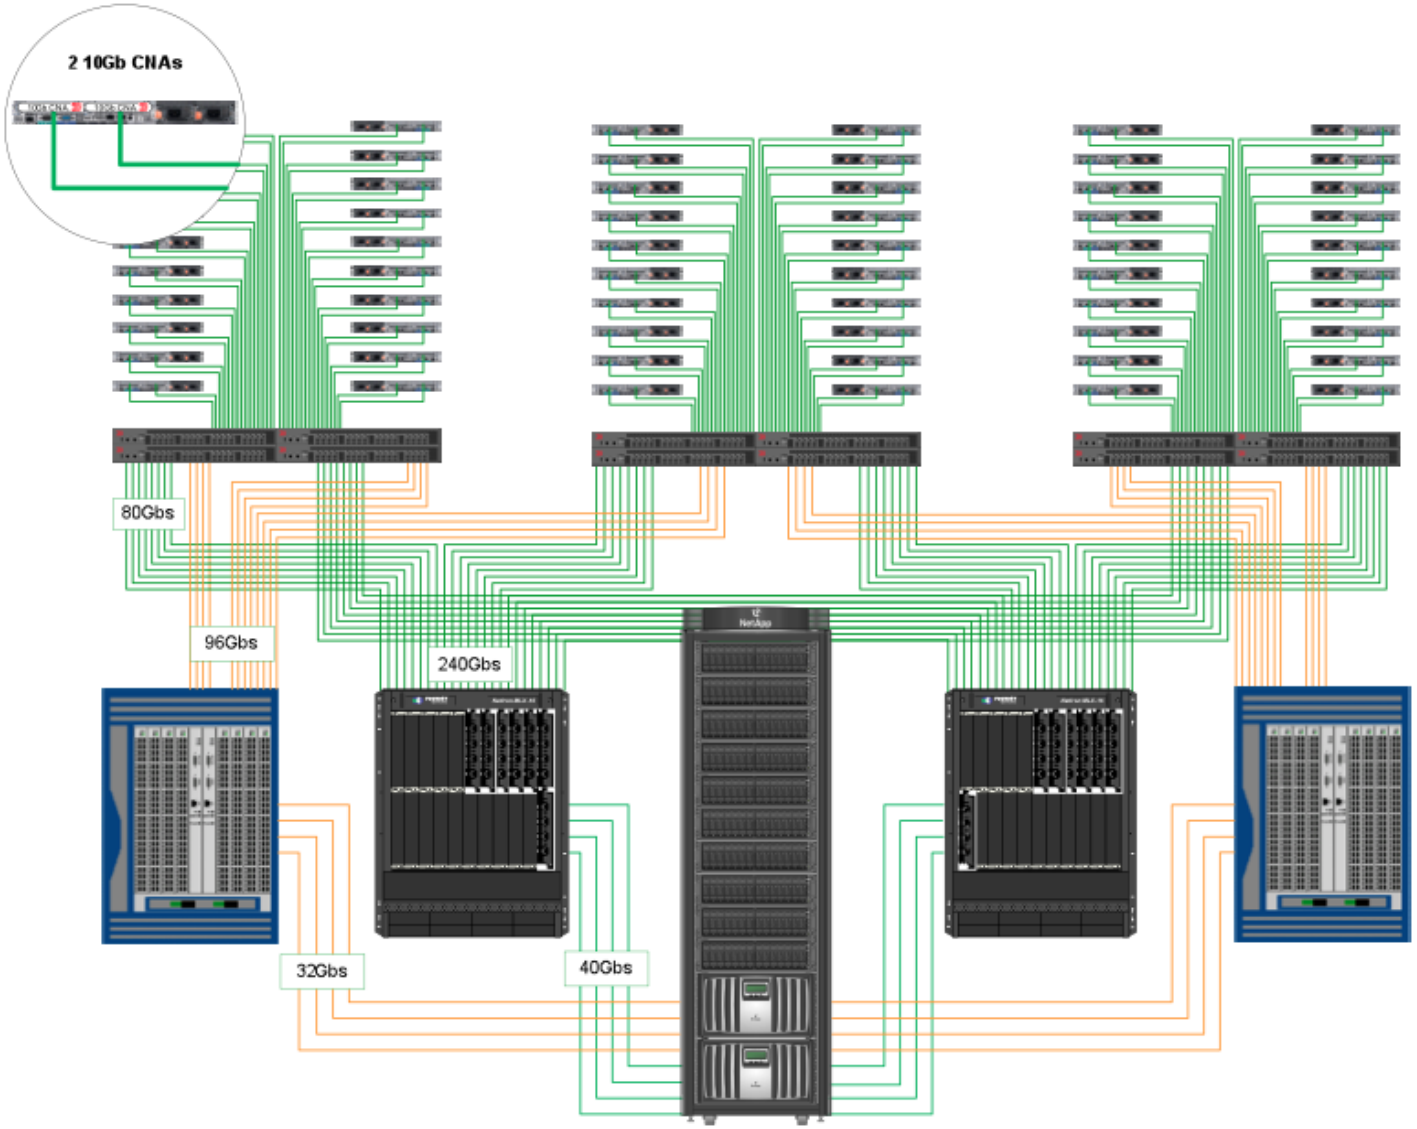

Multi-hop FCoE

Low Latency, Lossless, Resilient, Storage Aware Network

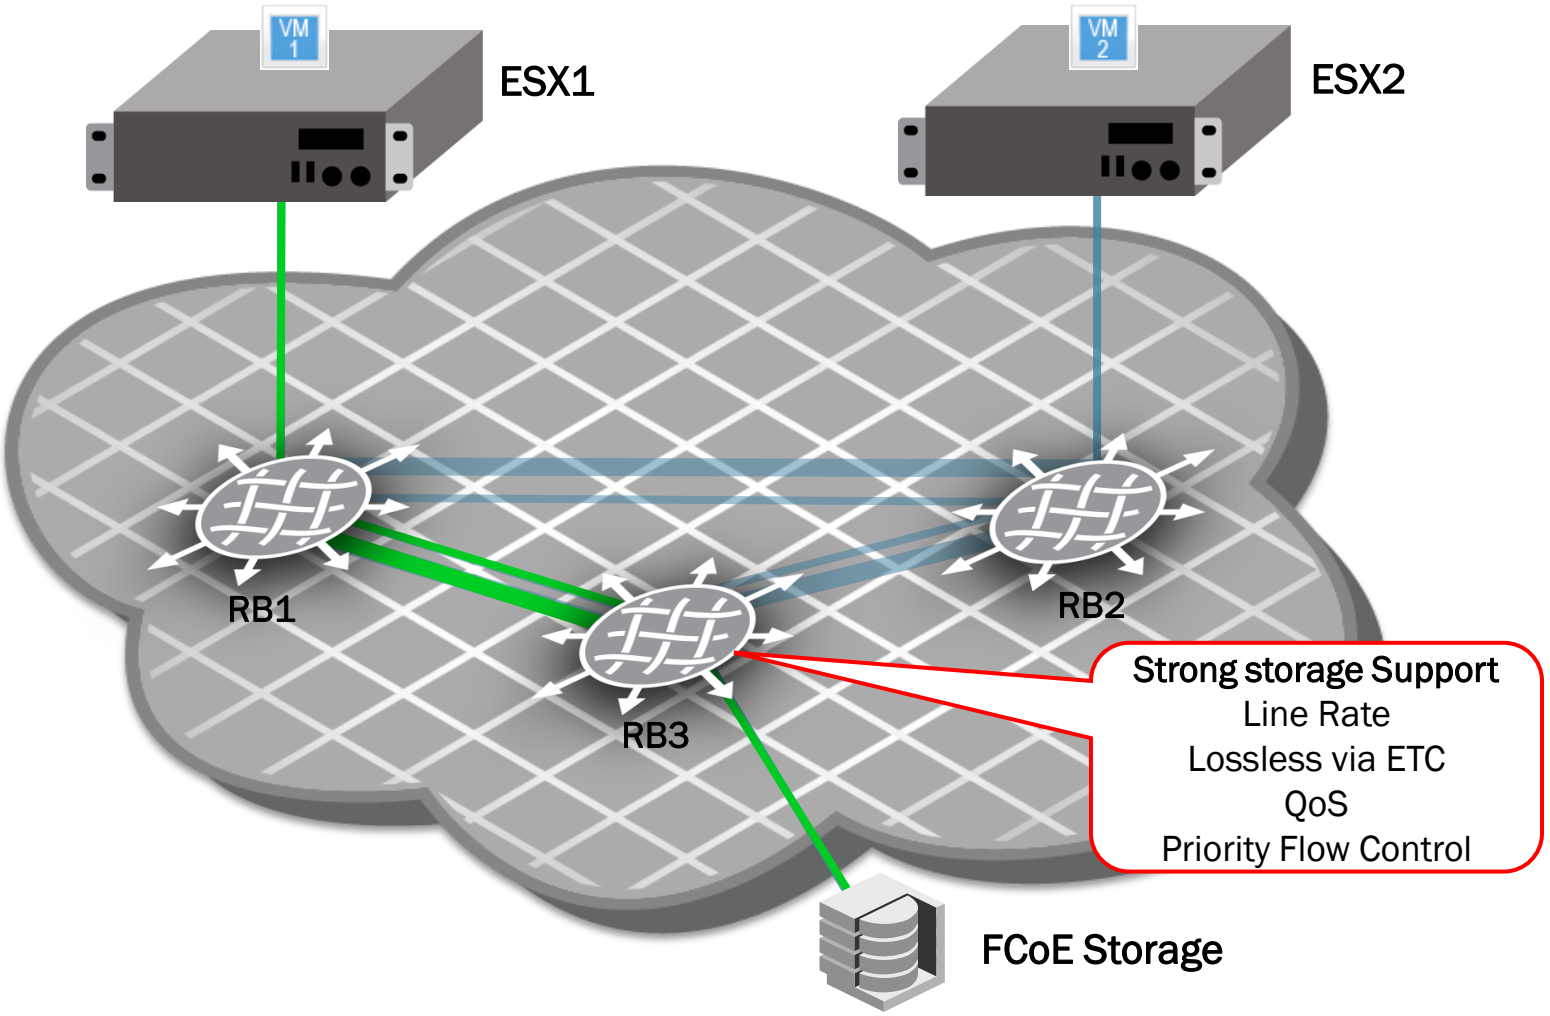

Easy Configuration & Setup

Included default CEE Map is the recommended starting place

```
RB1# show cee map default

CEE Map 'default'
   Precedence: 1
   Remap Fabric-Priority to Priority 0
   Remap Lossless-Priority to Priority 0
   Priority Group Table
     1:  Weight 40, PFC Enabled, BW% 40
     2:  Weight 60, PFC Disabled, BW% 60
     15.0: PFC Disabled
     15.1: PFC Disabled
     15.2: PFC Disabled
     15.3: PFC Disabled
     15.4: PFC Disabled
     15.5: PFC Disabled
     15.6: PFC Disabled
     15.7: PFC Disabled
   Priority Table
      CoS:    0    1    2    3    4    5    6    7
     ----------------------------------------
     PGID:    2    2    2    1    2    2    2    2
   Enabled on the following interfaces:
        Te 1/0/5

RB1#
```

Easy Configuration

Configuration on Switch is pushed to CNA

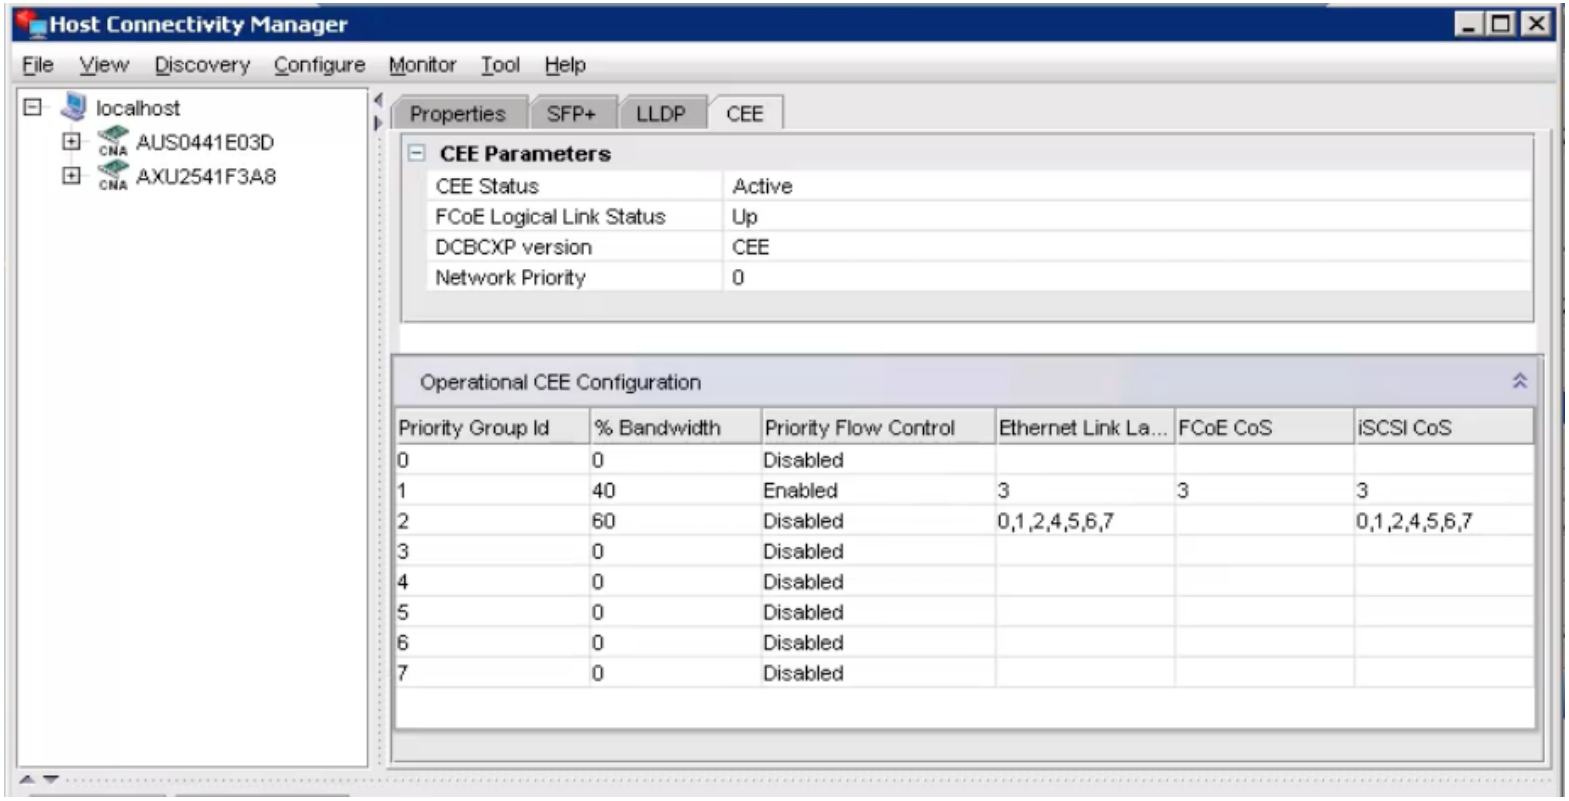

Performance is Outstanding

FCoE over 10Gb has very low overhead

9.7Gb/s

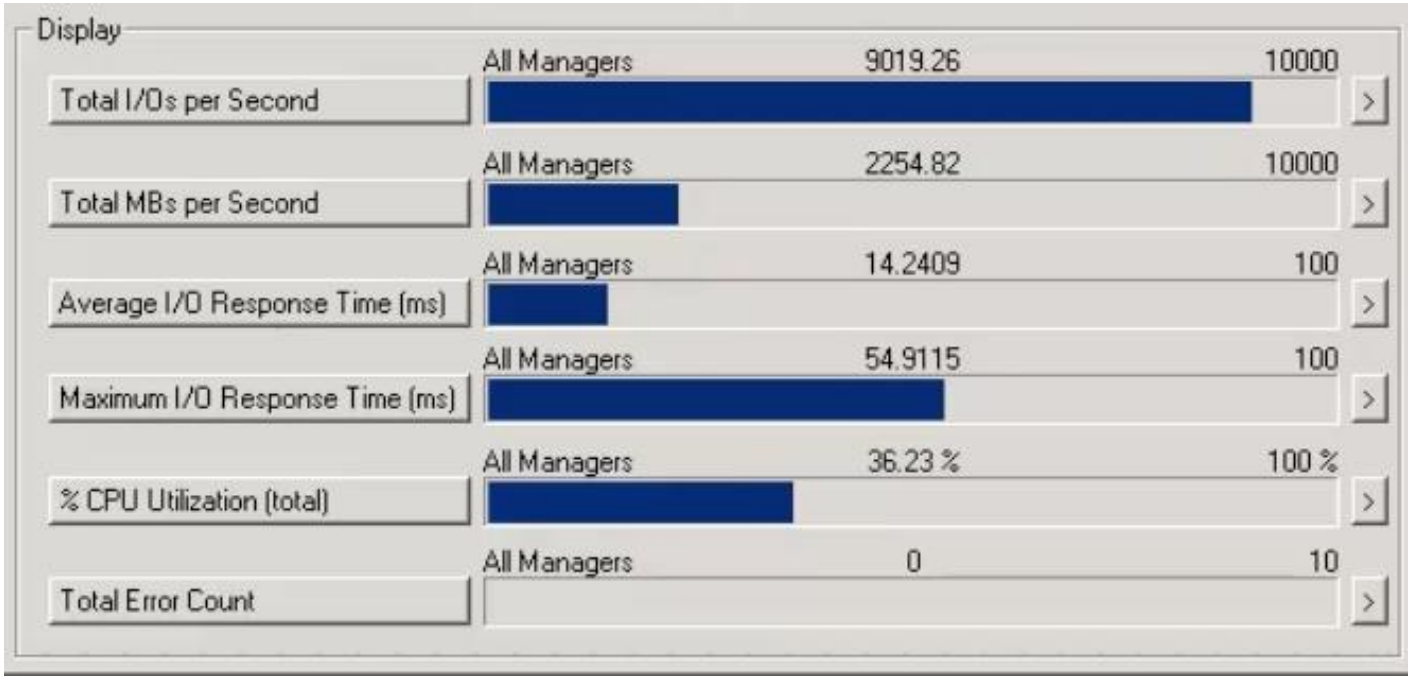


```
RB3(conf-if-te-3/0/12)# do show interface tengigabitethernet 3/0/12 | include rate
Queueing strategy: fifo
    Input 9567.779070 Mbits/sec, 562874 packets/sec, 95.68% of line-rate
    Output 9734.734532 Mbits/sec, 574924 packets/sec, 97.35% of line-rate
RB3(conf-if-te-3/0/12)#
```

ECMP provides very even traffic balance

4 Links (1 single link, 3 in a trunk) with balanced traffic

```
RB2# show interface tengigabitethernet 2/0/17 | include rate
Queueing strategy: fifo
    Input 2427.565040 Mbits/sec, 151336 packets/sec, 24.28% of line-rate
    Output 2519.642712 Mbits/sec, 146842 packets/sec, 25.20% of line-rate
RB2# show interface tengigabitethernet 2/0/1 | include rate
Queueing strategy: fifo
    Input 2464.649672 Mbits/sec, 151497 packets/sec, 24.65% of line-rate
    Output 2463.273888 Mbits/sec, 143580 packets/sec, 24.63% of line-rate
RB2# show interface tengigabitethernet 2/0/2 | include rate
Queueing strategy: fifo
    Input 2441.272776 Mbits/sec, 148736 packets/sec, 24.41% of line-rate
    Output 2476.909300 Mbits/sec, 144242 packets/sec, 24.77% of line-rate
RB2# show interface tengigabitethernet 2/0/3 | include rate
Queueing strategy: fifo
    Input 2438.666696 Mbits/sec, 152191 packets/sec, 24.39% of line-rate
    Output 2470.860516 Mbits/sec, 144166 packets/sec, 24.71% of line-rate
RB2#
```

ECMP Redistribution

After removing one link, traffic immediately rebalanced

```
RB2(conf-if-te-2/0/1)# do show interface tengigabitethernet 2/0/2 | in ra
Queueing strategy: fifo
Transmit Statistics:
    Input 3265.368904 Mbits/sec, 201013 packets/sec, 32.65% of line-rate
    Output 3287.184864 Mbits/sec, 191695 packets/sec, 32.87% of line-rate
RB2(conf-if-te-2/0/1)# do show interface tengigabitethernet 2/0/3 | in ra
Queueing strategy: fifo
Transmit Statistics:
    Input 3311.434088 Mbits/sec, 200685 packets/sec, 33.11% of line-rate
    Output 3267.580168 Mbits/sec, 190552 packets/sec, 32.68% of line-rate
RB2(conf-if-te-2/0/1)# do show interface tengigabitethernet 2/0/17 | in ra
Queueing strategy: fifo
Transmit Statistics:
    Input 3280.195824 Mbits/sec, 199546 packets/sec, 32.80% of line-rate
    Output 3281.832392 Mbits/sec, 191864 packets/sec, 32.82% of line-rate
RB2(conf-if-te-2/0/1)#
```

Brocade 1010 and 1020 CNAs

DCB-enabled Converged Network Adapters (CNAs)

BR-1010-0010 (optical)
BR-1010-1010 (copper)

BR-1020-0010 (optical)
BR-1020-1010 (copper)

Key features:

- **10 Gbps** Data Center Bridging (DCB) for TCP/IP and storage (iSCSI/FCoE)
- **Single-** or **dual-port** models
- **Copper** (active twinax) and **optical** options
- Up to **500,000** IOPS per port
- **Common** FC and FCoE **driver**
- x8-lane **PCI Express 2.0**
- **Low-profile design;** LC-style pluggable SFP+
- Single chip, low power **(8.5 watts)**

VDX 6720 Ethernet Fabric Switches

- **Built for the virtual data center**
  - Uses Brocade fabric switching ASICs
  - First switches to run new Brocade Network Operating System
  - Virtual Cluster Switching (VCS) fabric technology
  - Automatic Migration of Port Profiles (AMPP)
- **Best-in-class performance and density**
  - 24- and 60-port models
  - Non-blocking, cut-through architecture, wire-speed
  - **600 ns port-to-port latency; 1.8 us across port groups**
- **Environmental flexibility**
  - 10 GbE and 1 GbE supported on every port
  - Twinax, direct-attached optical, and SFP optical connectivity options
  - Less than 17" switch depth and reversible front-to-back airflow
- **Enables network convergence**
  - Complete FCoE support, multi-hop
  - iSCSI DCB support
- **Highly resilient and efficient design**
  - Hot code load and activation
  - Remote "lights out" management
  - Simplistic design, optimal power efficiency

VDX 6720 ETHERNET FABRIC

Brocade VDX Product Family

Delivering virtual cluster switching

BROCADE VCS

A new family of Ethernet Fabric switches

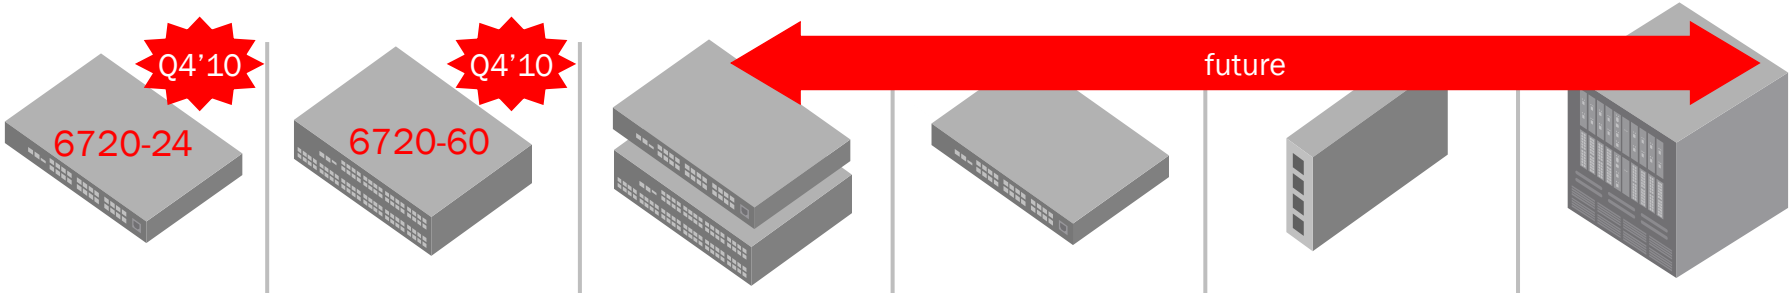

24-port 1/10 GbE 600 ns latency Fastest Ethernet switch available	60-port 1/10 GbE High density Wire-speed	24- and 60-port 1/10 GbE Fibre Channel ports for connectivity to SAN	48-port 1 GbE High-density 1 GbE VCS connectivity	VCS in blade server chassis 1/10 GbE Blade switches are part of the logical chassis	Wire-speed chassis with VCS Allows Ethernet fabrics to scale further

Ethernet Fabrics

Easing into a new A New Network Architecture

Classic Hierarchical Ethernet Architecture

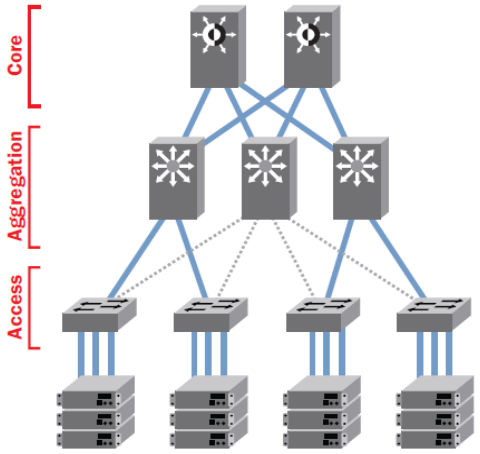

Servers with 10 Gbps Connections

Hybrid Ethernet Fabric Architecture

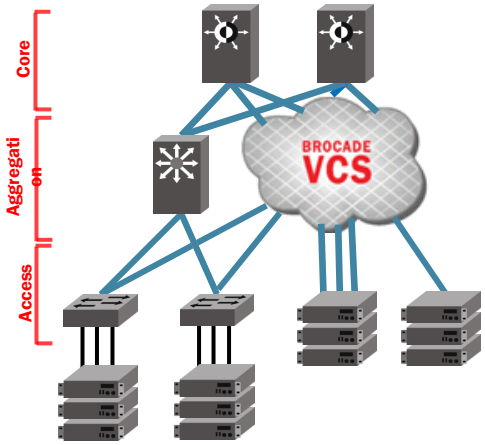

Servers with 1 and 10Gbps Connections

Ethernet Fabric Architecture

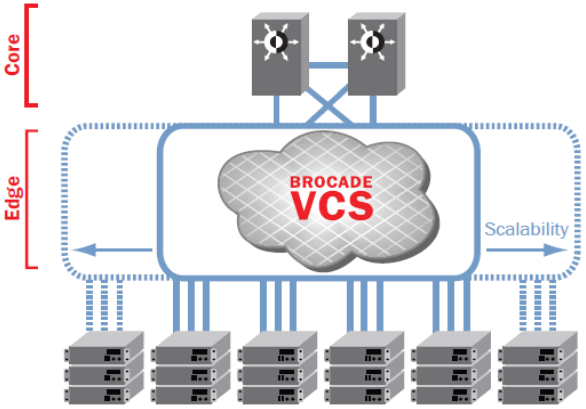

Servers with 10 Gbps Connections

Migrating to an Ethernet Fabric is Easy

Seamless integration into existing environments

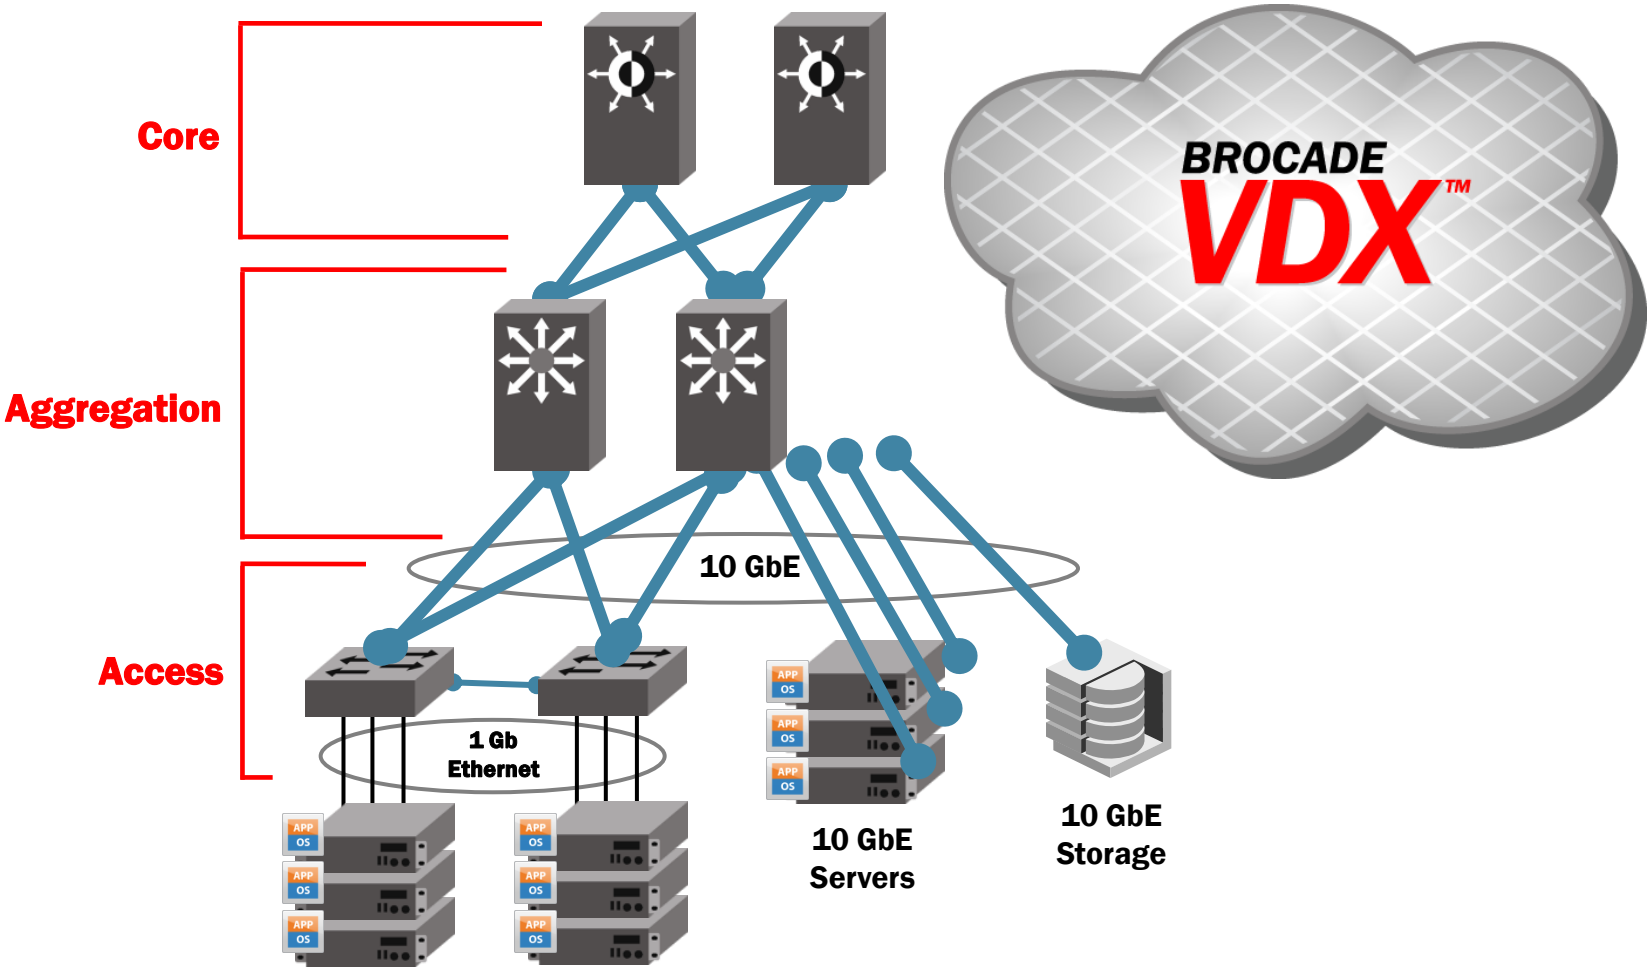

Future VCS Enabled Datacenter

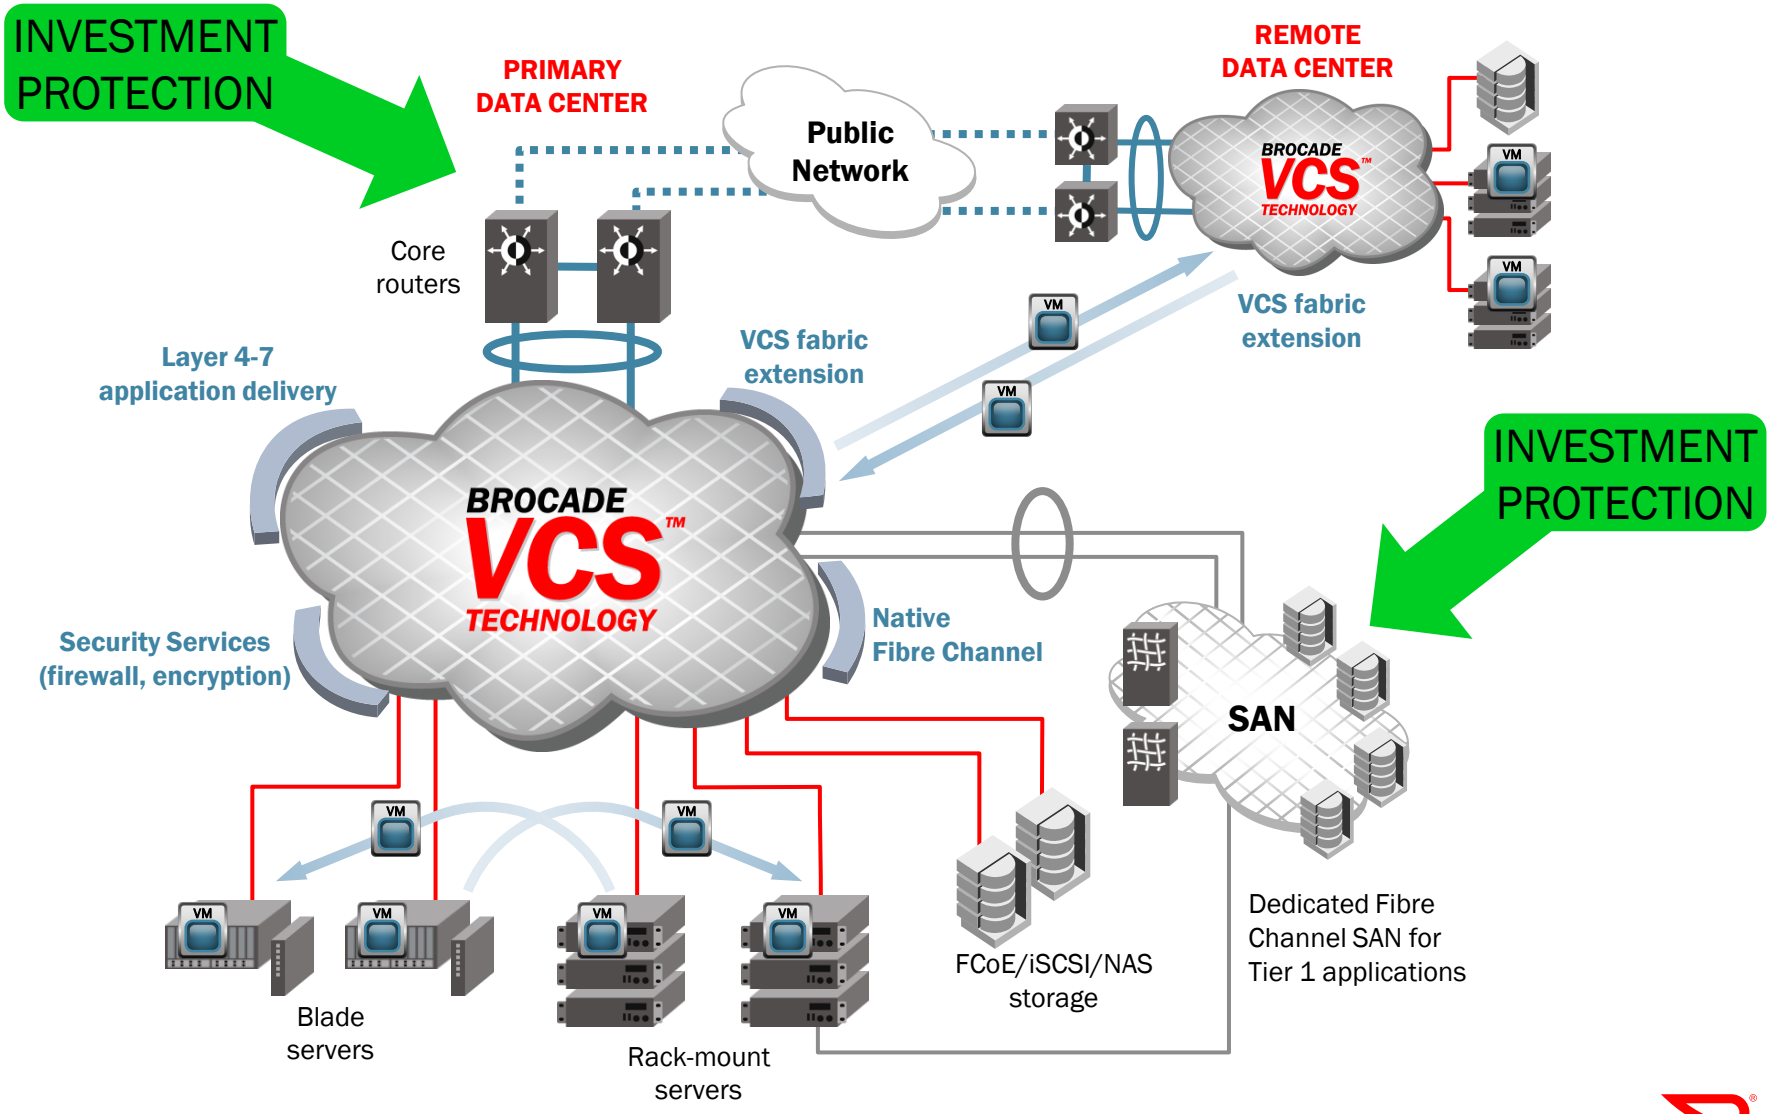

Virtual Cluster Switching (VCS)

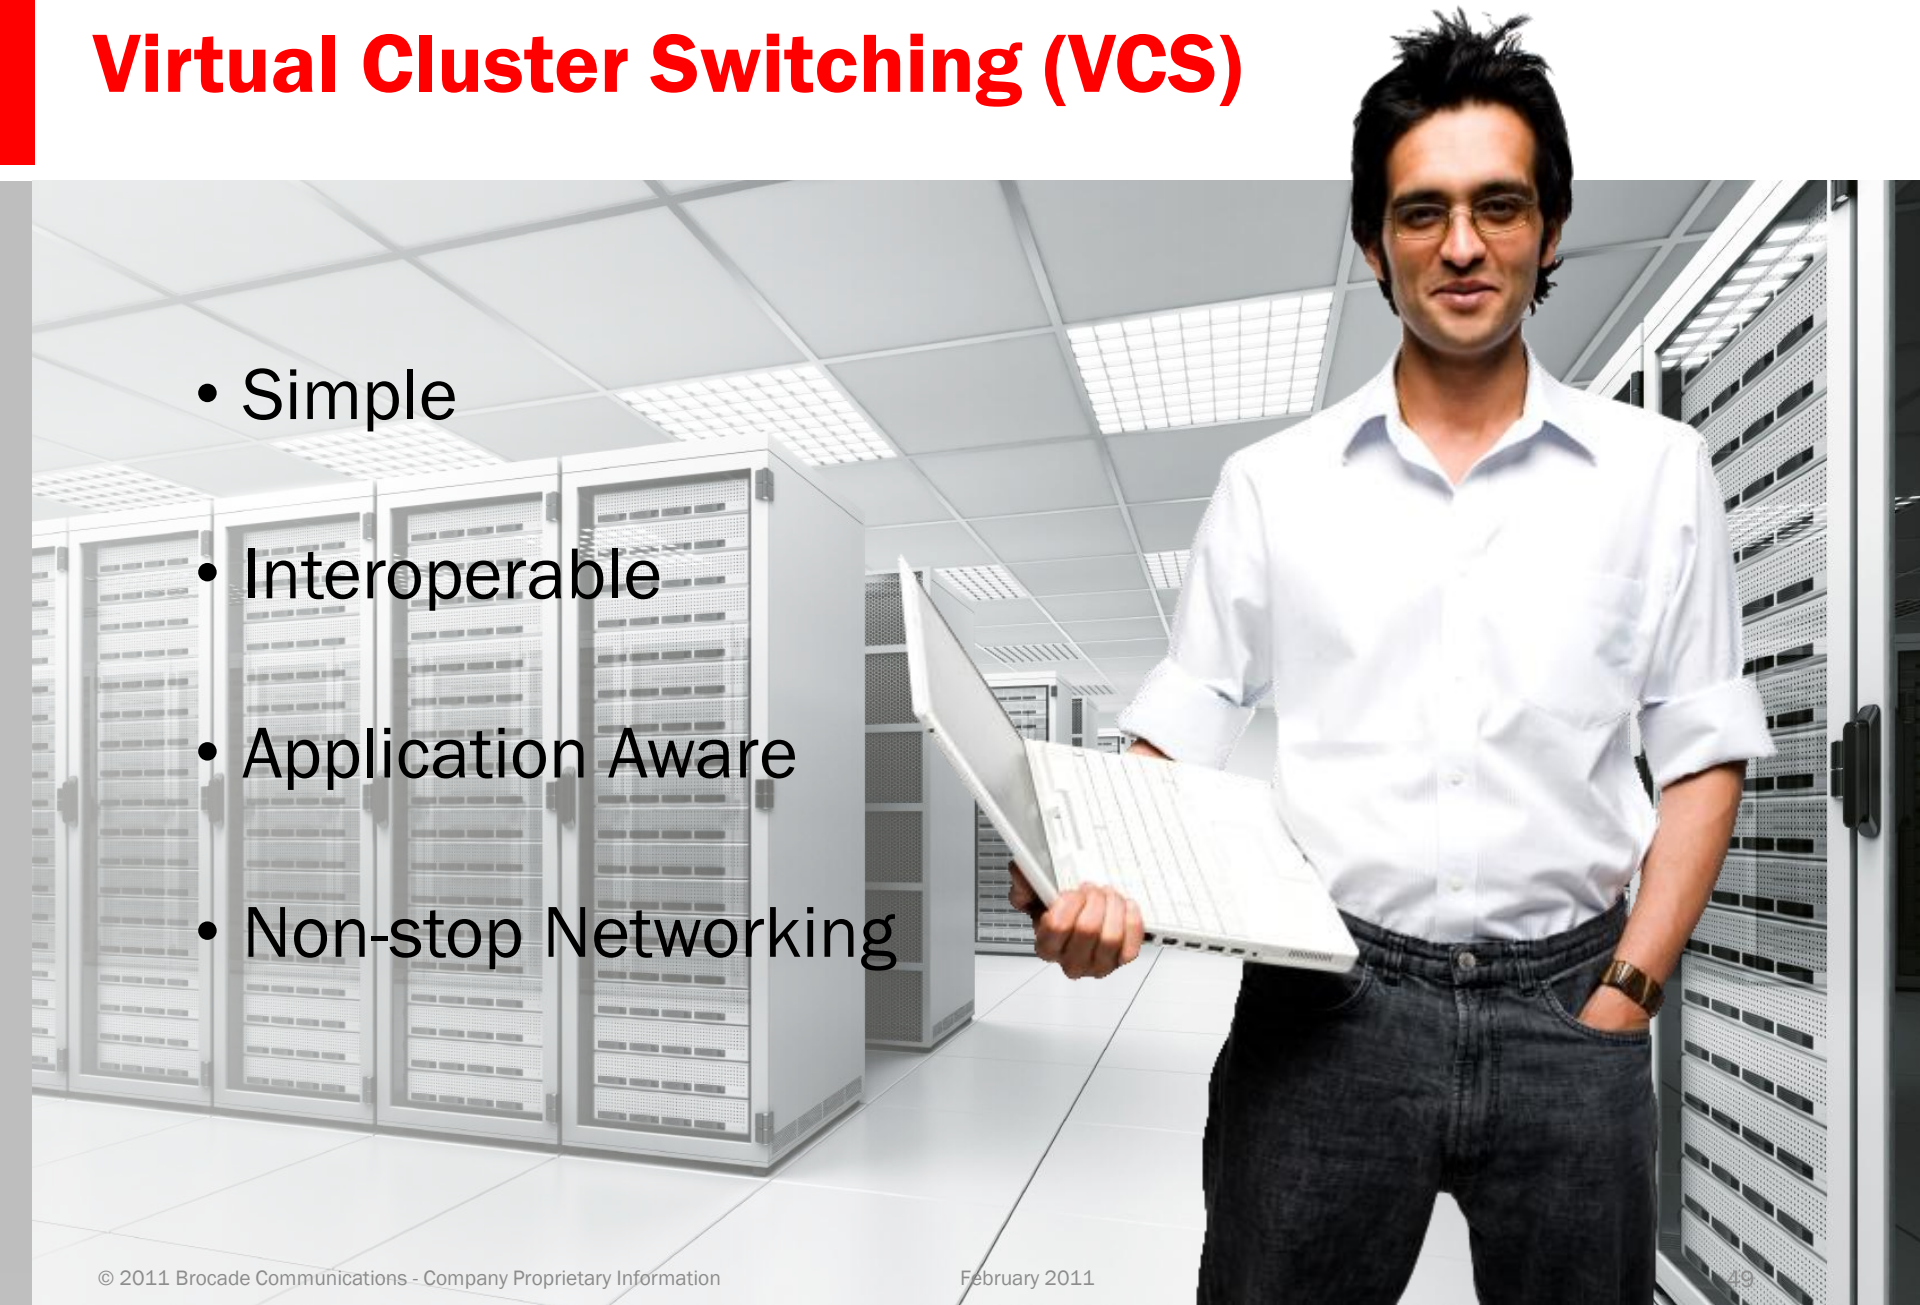

- Simple
- Interoperable
- Application Aware
- Non-stop Networking

BROCADE

Thank You