BROCADE

VIRTUAL CLUSTER SWITCHING

SWITCHES AS A CLOUD FOR THE VIRTUAL DATA CENTER

Emil Kacperek
Systems Engineer

Brocade Communication Systems

Mar, 2011

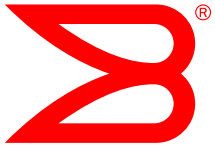

BROCADE

Acquired Foundry 2008

- Data center networking experts
- Storage networking pioneer and leader
- 70% SAN market share

- Price/performance leader in IP networks
- Powering 90% of Internet Exchange Points
- 15,000+ customers worldwide

Powering Mission Critical Networks

> 60%	Of Internet Traffic Traverses Brocade Infrastructure
1st	Multi-terabit Core Router in a Single Chassis
#1	IPv4 Transit Network (Level3) IPv6 Transit Network (Hurricane Electric)
4	Top 4 Carrier Ethernet Switching and Routing Vendors
6	Largest Internet Exchange Points in the World
7	Of 10 Largest Content Delivery Networks
10,000+	IP/MPLS Core and Edge Routers Deployed World-wide

Brocade One for Service Providers

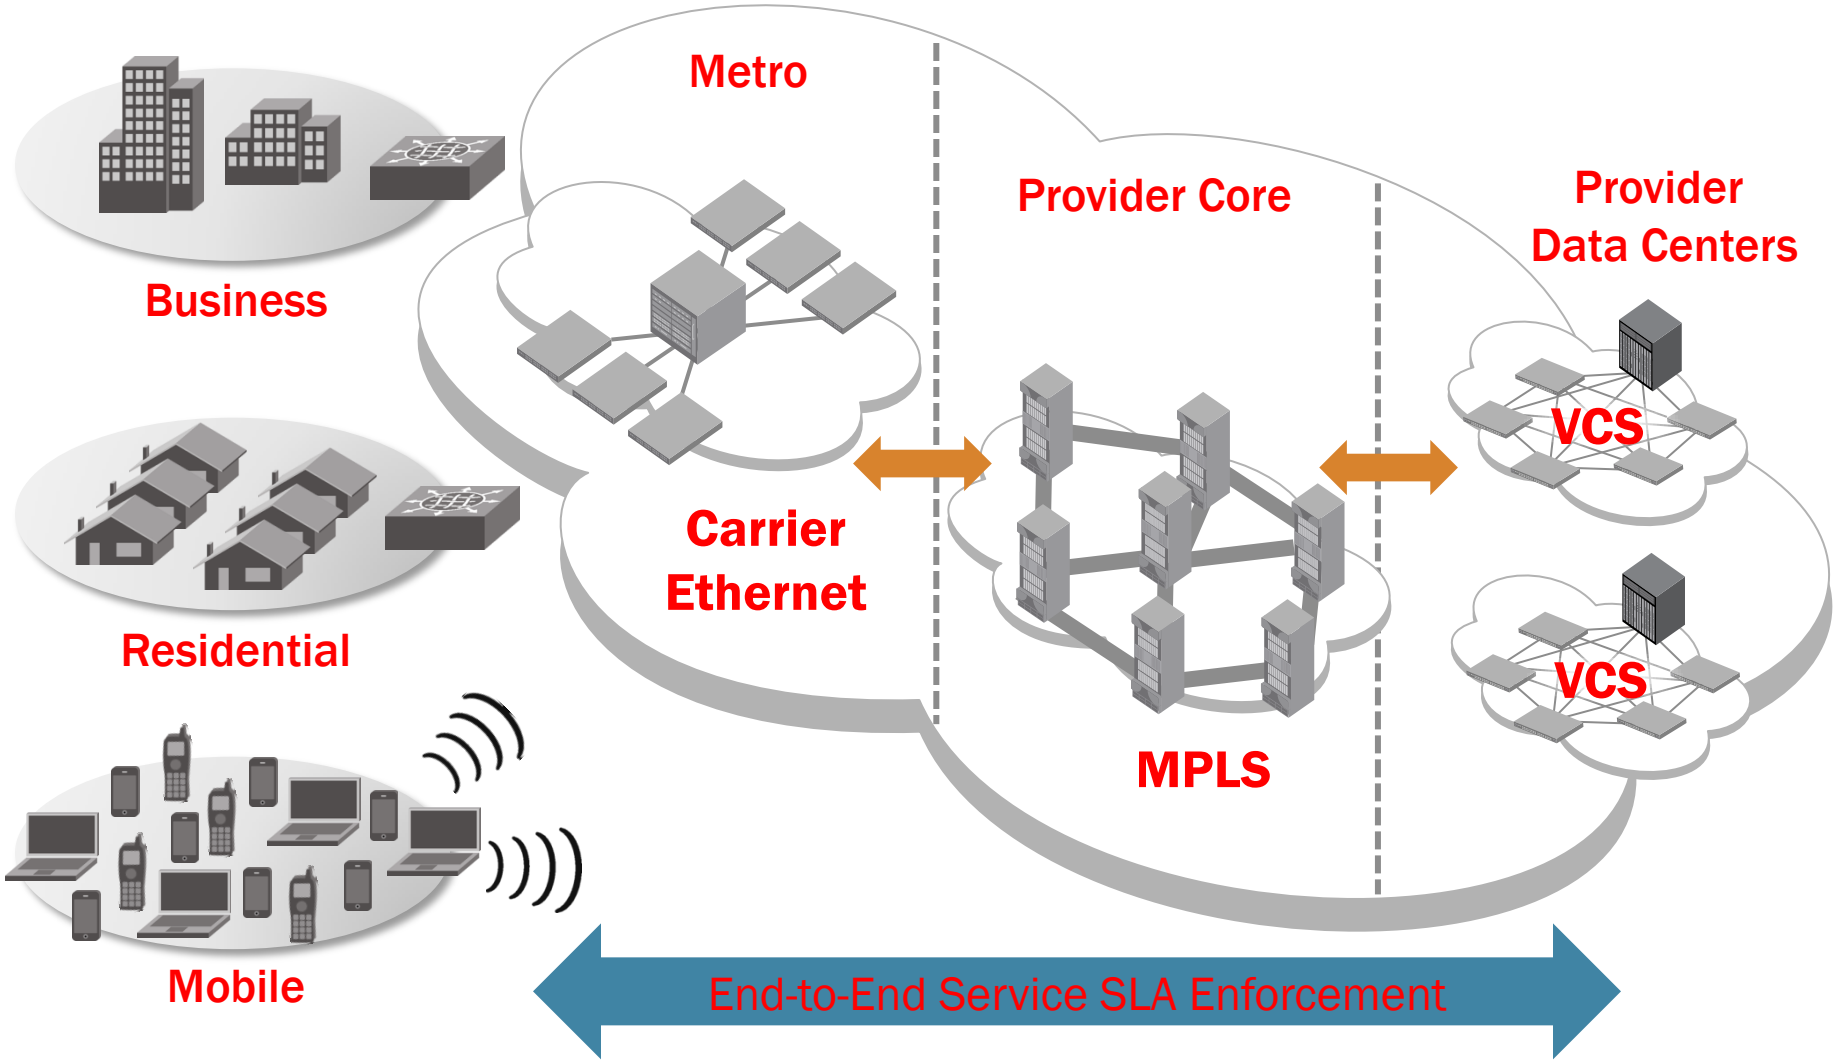

Brocade High-Performance Routers

Delivers performance from service provider edge to core

Access

Metro Aggregation and Provider Edge

Brocade
MLX Routers

NetIron CER
2024/48/C/F

Provider Core

Brocade
MLXe Routers

New

Brocade
XMR series

Brocade MLXe Routers

Product Highlights

A Unified Platform That Scales From Data Center Core to Service Provider Core

- Industry-leading switch fabric capacity 15.36 Tbps
- Half-slot capacity 480 Gbps
- Industry-leading 10G wire-speed density of 256 ports
- Industry-leading 100G wire-speed density of 32 ports
- Investment protection - Supports all the existing MLX modules
- Future-ready for timing distribution (Sync-E, 1588)
- Ability to use XMR modules- XMR level Scalability
- Rear exhaust for all MLXe models
- NEBS Level 3 Certified (pending)

Brocade NetIron CER Series

Compact 1–rack unit IP/MPLS Edge Router with 10 GbE upgradability

- **Scalable routing and VPN services**
  - Up to 512K IPv4 routes, 128 BGP peers
- **Common software across Brocade CER, CES, MLX, XMR**
- **Full MPLS routing**
  - VPLS, VLL, IP over MPLS, Layer 3 VPNs
  - Advanced traffic engineering
  - Fast reroute with sub-50 ms resiliency
- **Best-in-class flexibility**
  - 24- and 48-port copper and fiber models
- **Wire-speed performance for Layer 2/IPv4/IPv6/MPLS/VRF**
- **High availability**
  - Redundant AC/DC power supplies
  - N+1 redundant, replaceable cooling system

Multi-Service Edge Router

Enterprise Border Router

Brocade One for Service Providers

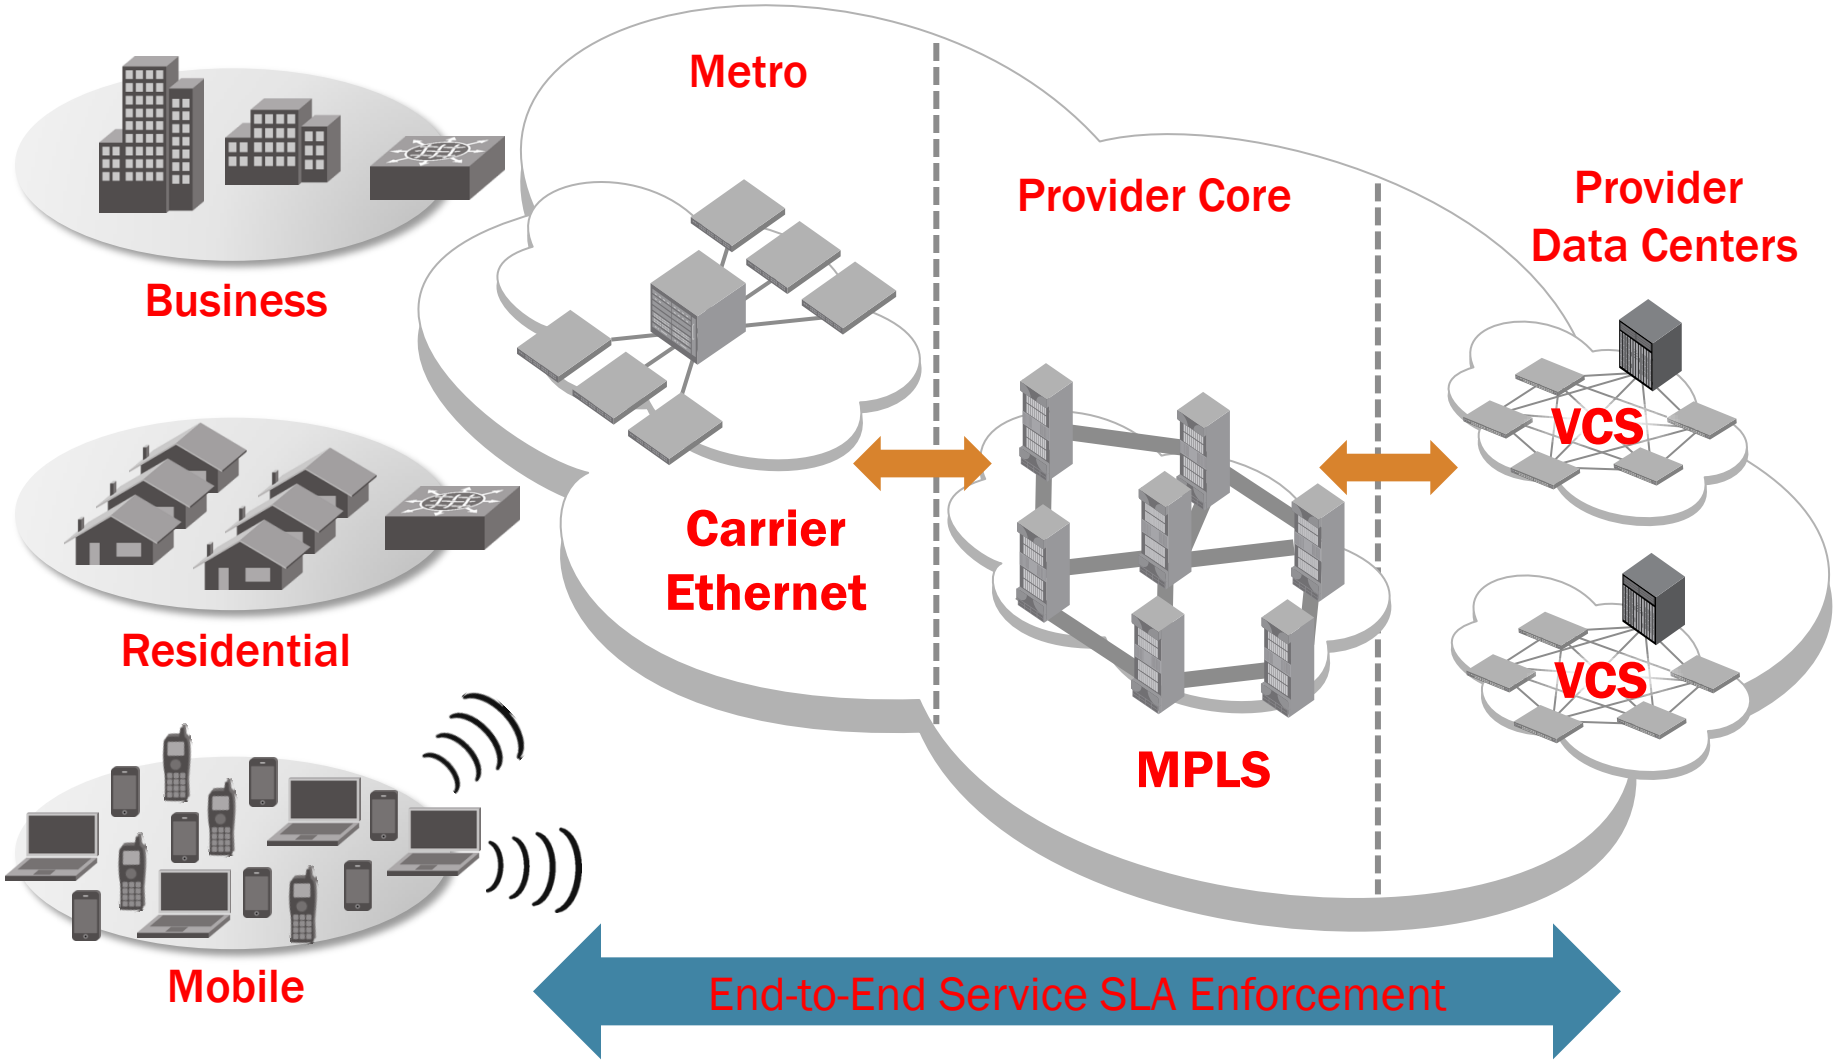

Data Center Transformation

Network Evolution

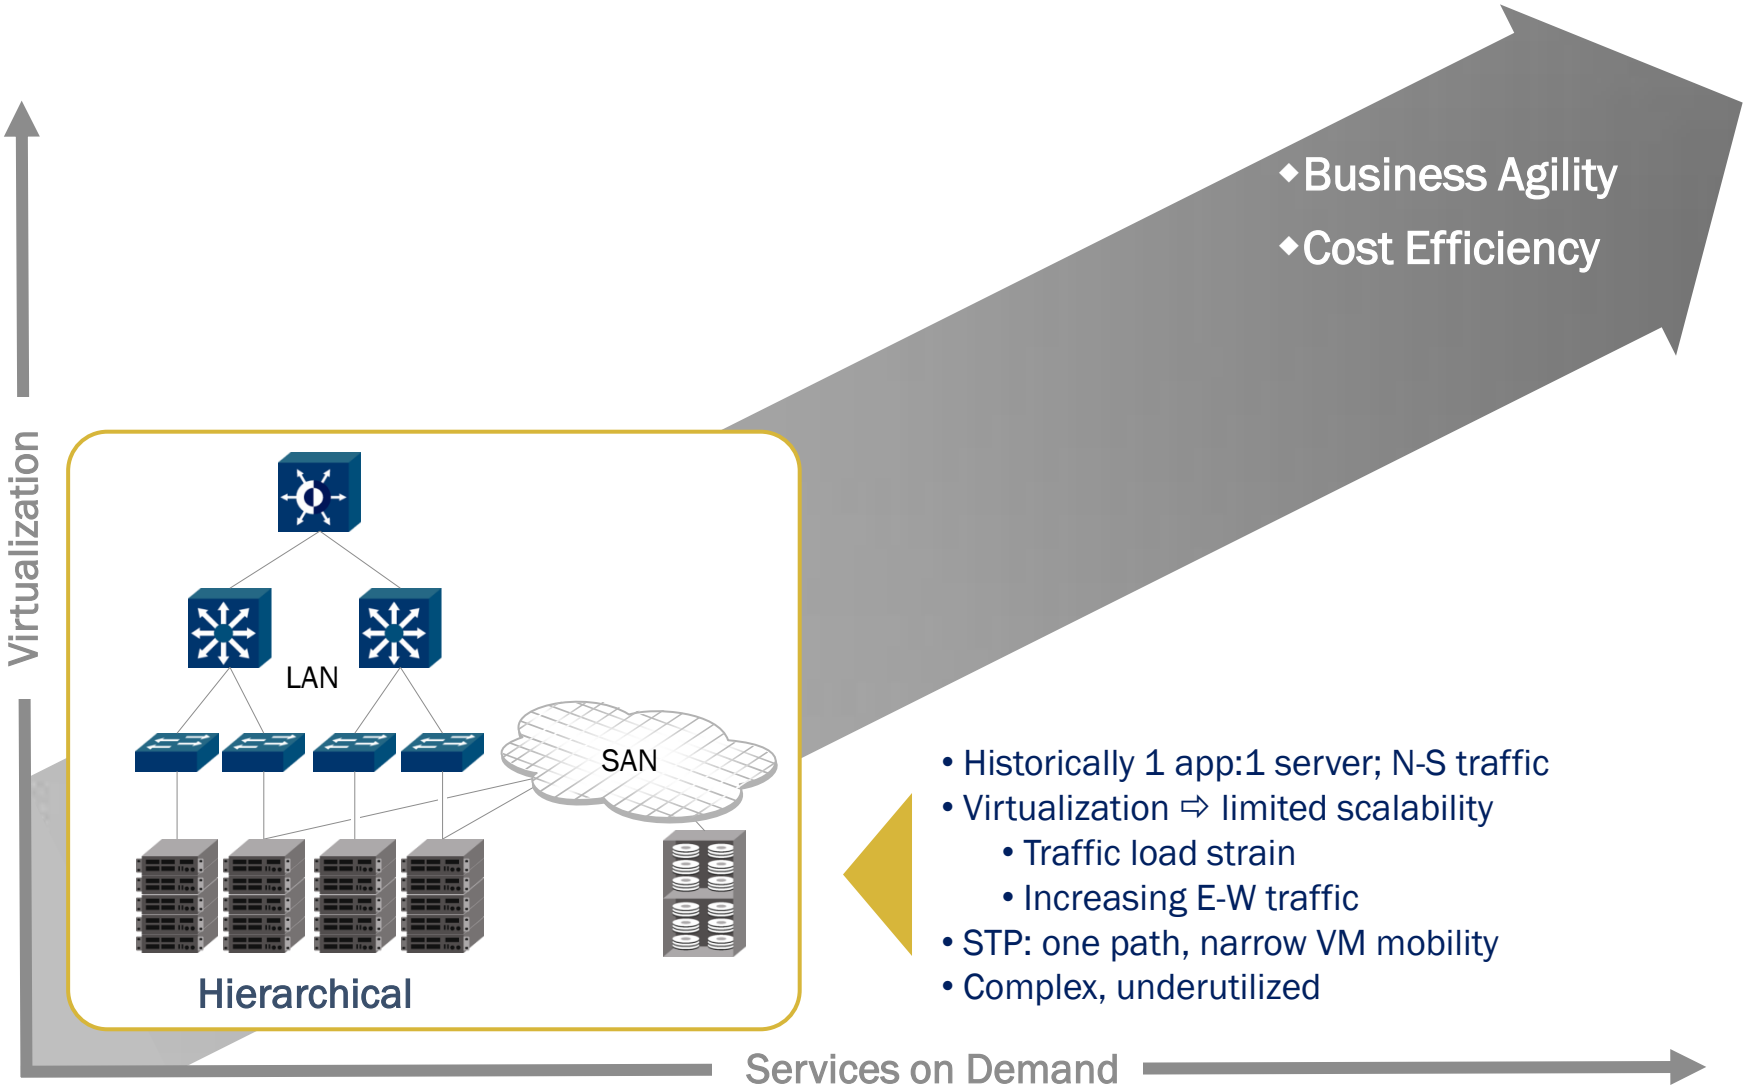

Data Center Transformation

Network Evolution

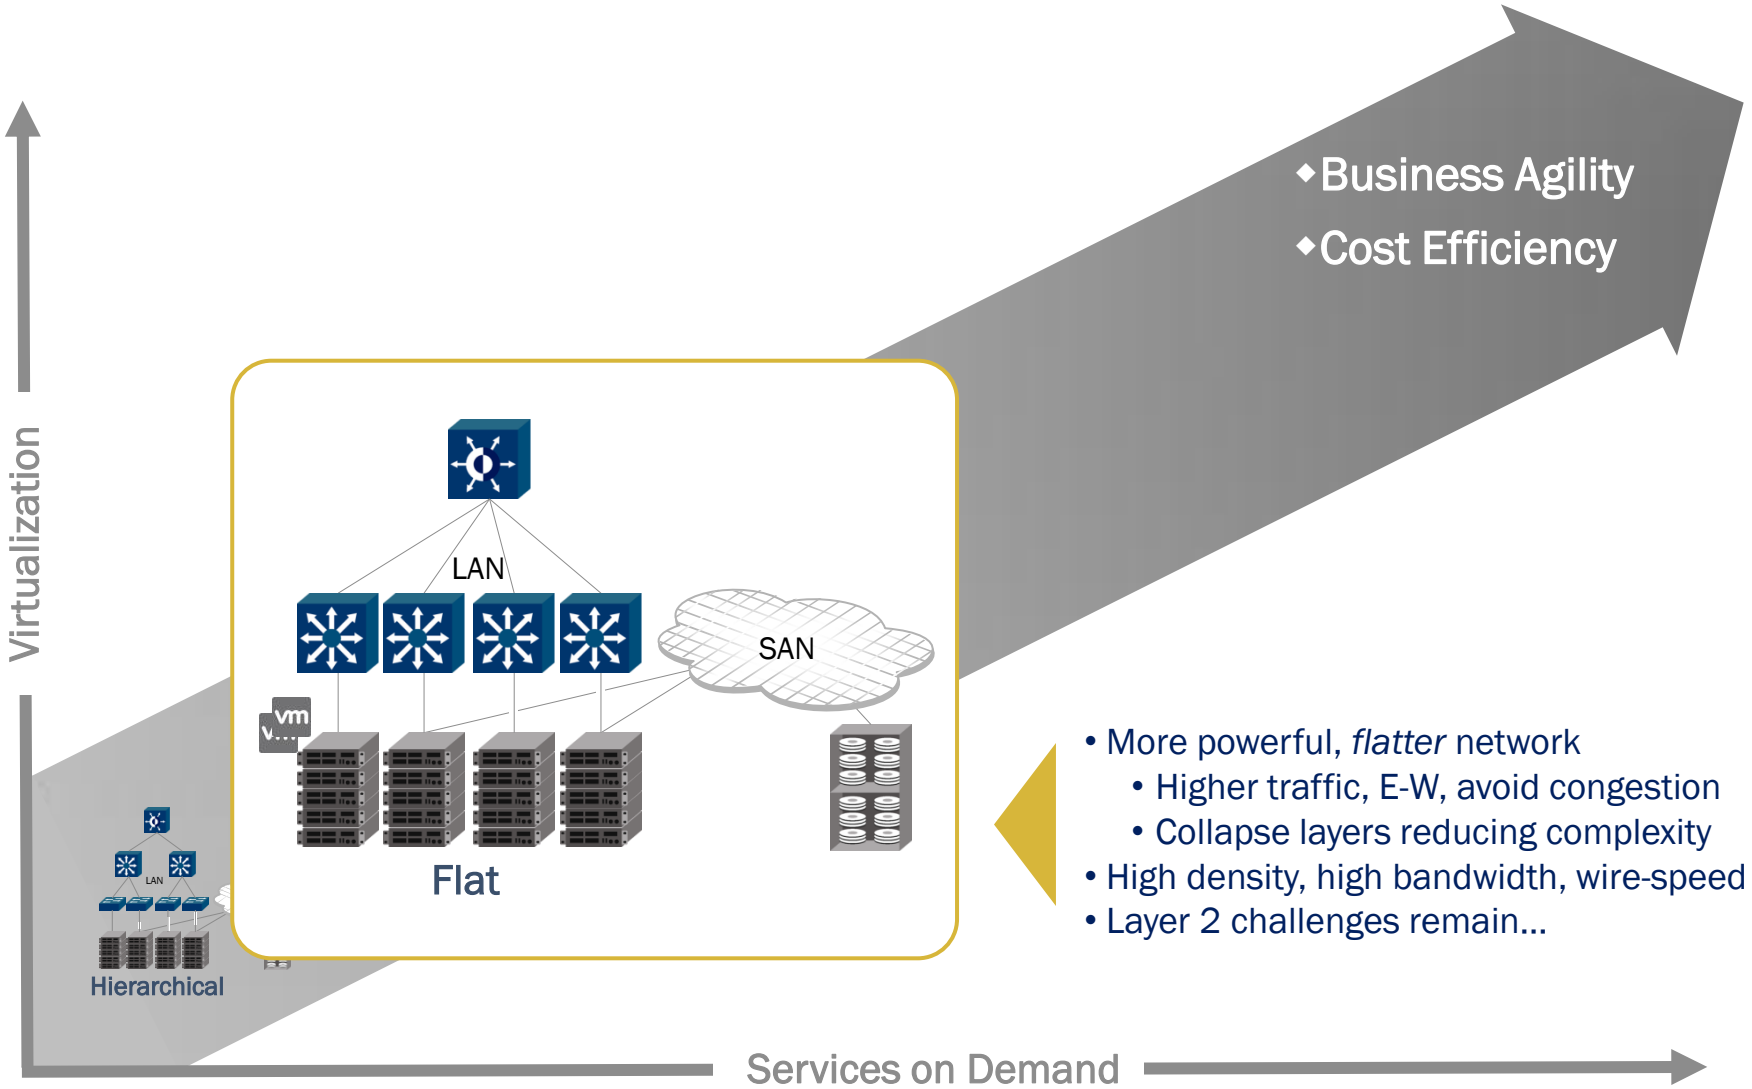

- More powerful, *flatter* network
  - Higher traffic, E-W, avoid congestion
  - Collapse layers reducing complexity
- High density, high bandwidth, wire-speed
- Layer 2 challenges remain…

Data Center Transformation

Network Evolution

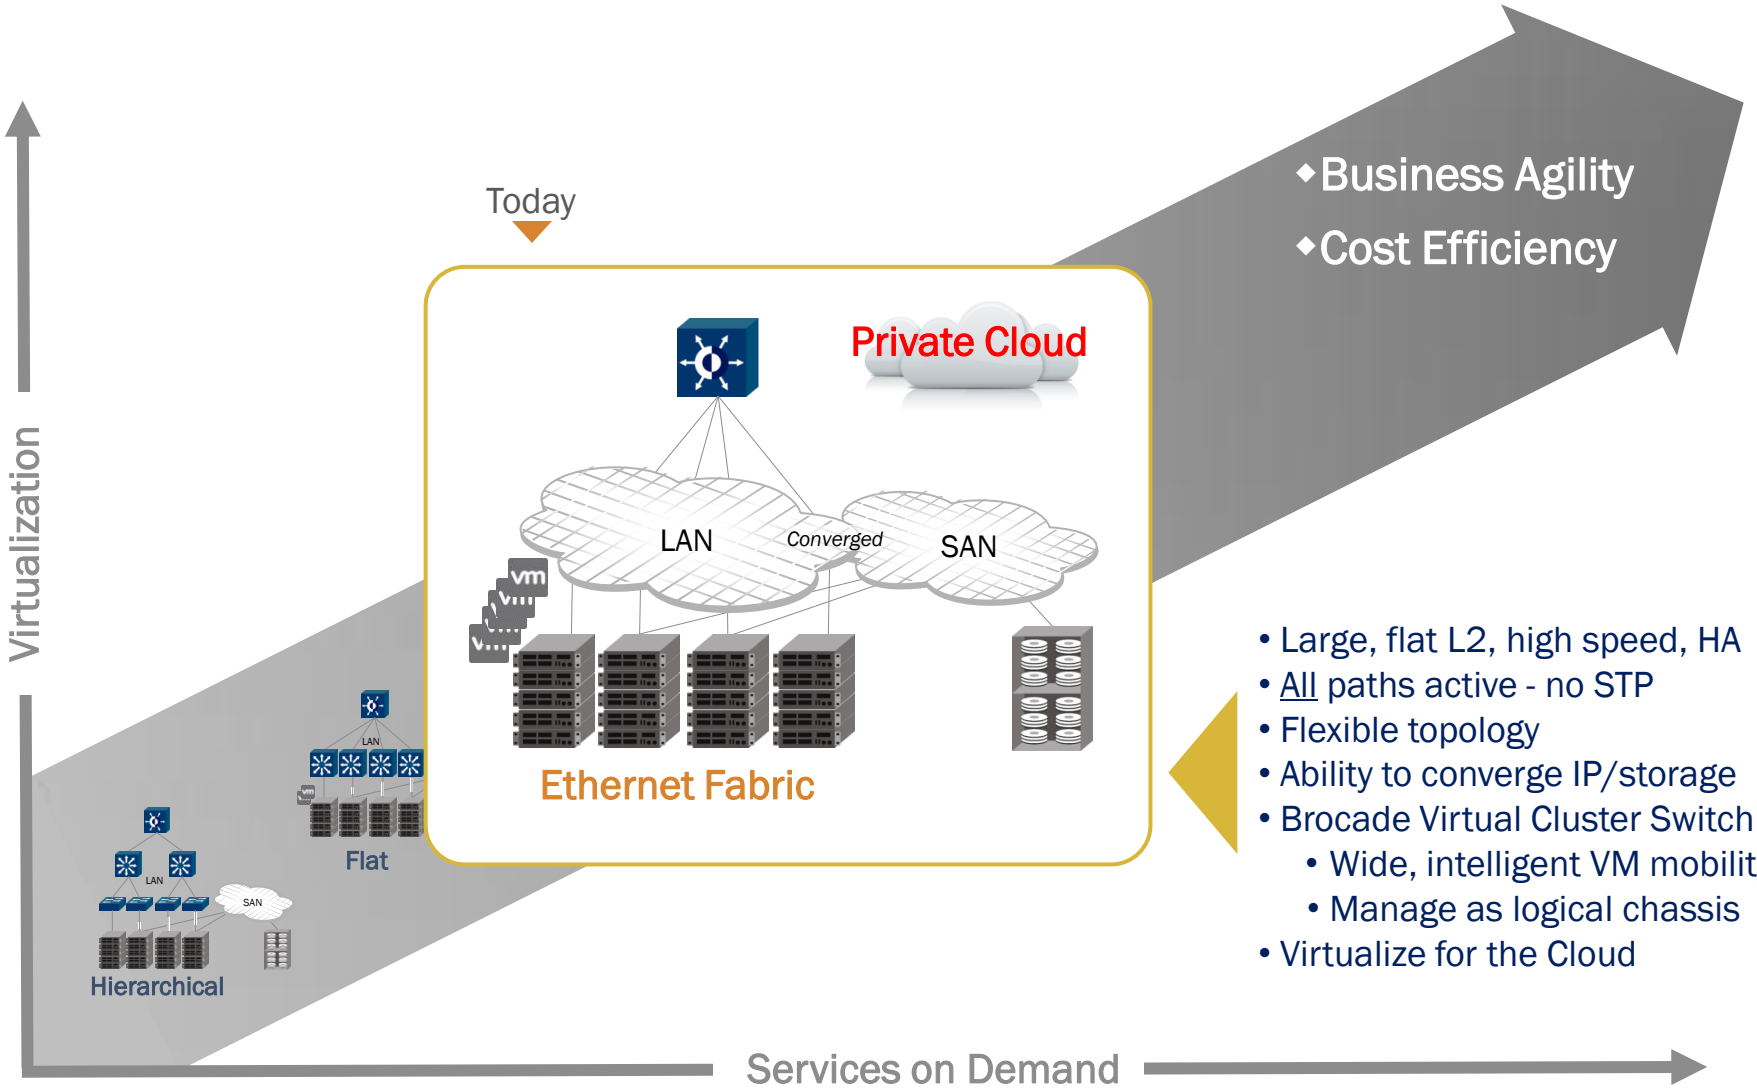

- Large, flat L2, high speed, HA
- All paths active - no STP
- Flexible topology
- Ability to converge IP/storage
- Brocade Virtual Cluster Switching
  - Wide, intelligent VM mobility
  - Manage as logical chassis
- Virtualize for the Cloud

Next Generation Data Centre

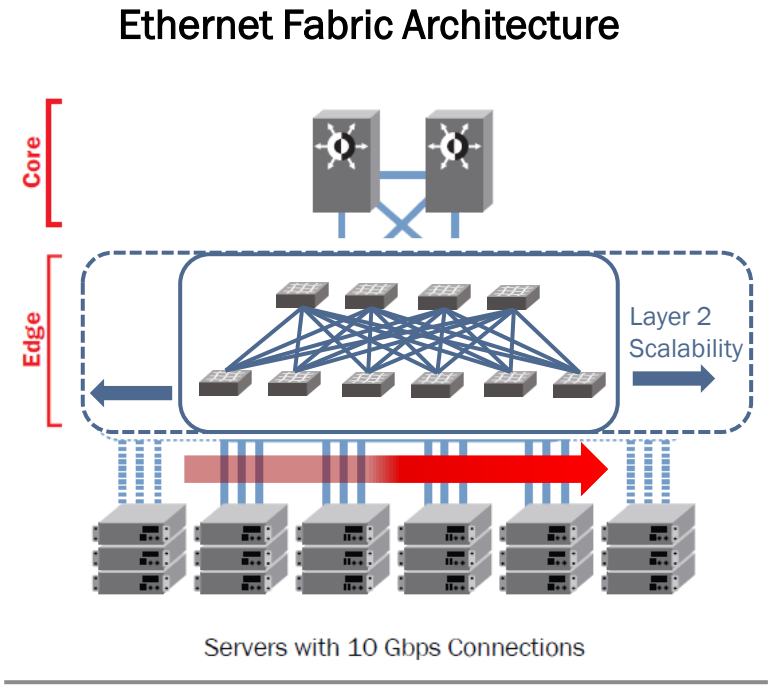

- VCS is a Ethernet fabric
- Scalable single layer 2 domain
- Optimized for East to West traffic
- Logical Chassis Nodes working together

IMAGINE IF...

- Eliminate the need for the Spanning-Tree Protocol
- High Performance & LARGE layer 2 Network
- Enhance Ethernet to support
  - Lossless transmission
  - Low latency
  - Storage convergence
- Active multipath fabric, similar to that found in SANs
- Greater network awareness of virtual app servers and mobility
- Able to manage all these devices as a single entity
- Reduce power consumption

Virtual Cluster Switching (VCS)

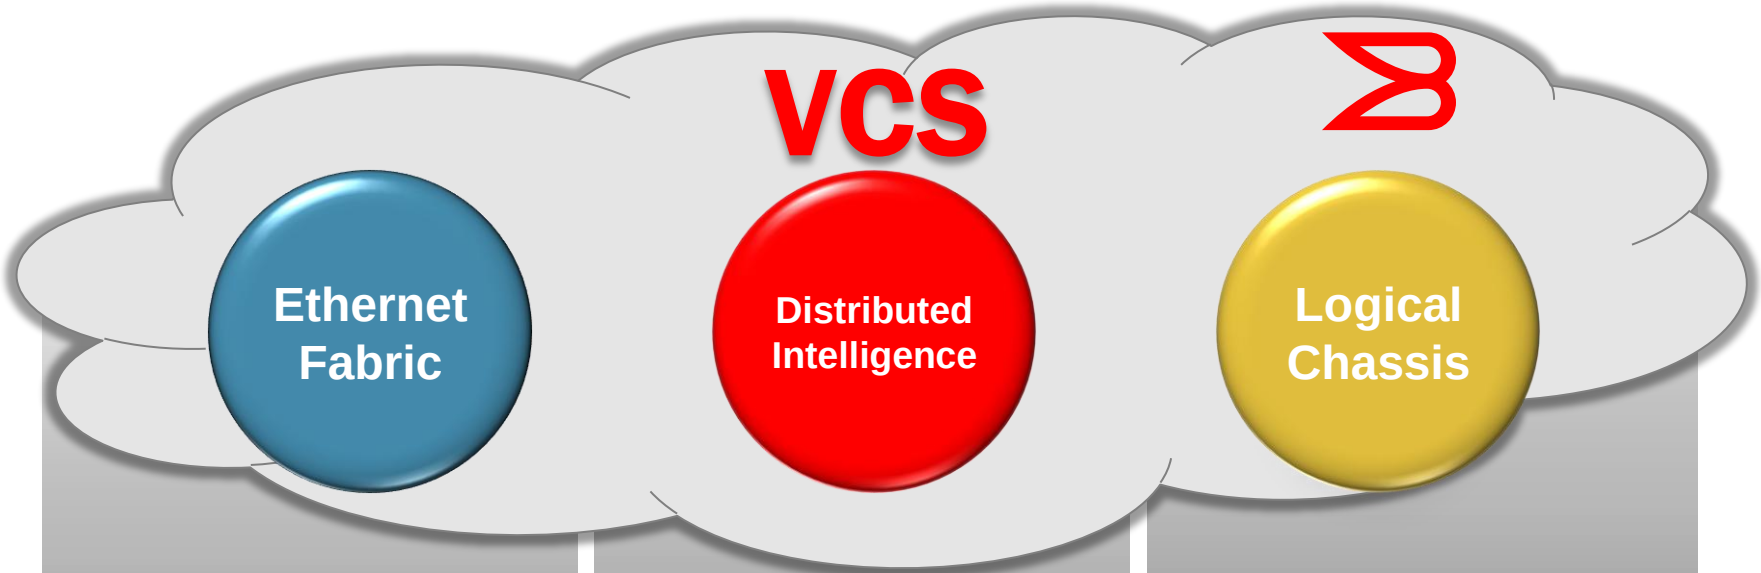

- No Spanning Tree Protocol
- Multi-path, deterministic
- Auto-healing, non-disruptive
- Lossless, low latency
- Convergence-ready

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

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

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

Ethernet Fabric Details

- Ethernet fabric
  - Layer 2 technology
- Link speed agnostic
- Data Center Bridging (DCB)
  - Lossless, deterministic
  - Priority-based Flow Control (PFC)
  - Enhanced Transmission Selection (ETS)
  - Data Center Bridging Exchange (DCBX)
- Transparent Interconnection of Lots of Links (TRILL)
  - Active multi-path
  - Multi-hop routing
  - Highly available, rapid link recovery
- LAN/SAN Convergence Ready
  - FCoE and iSCSI traffic
- Standards-based
  - Extends existing Ethernet infrastructure

Ethernet Fabric Details

Transparent Interconnection of Lots of Links (TRILL)

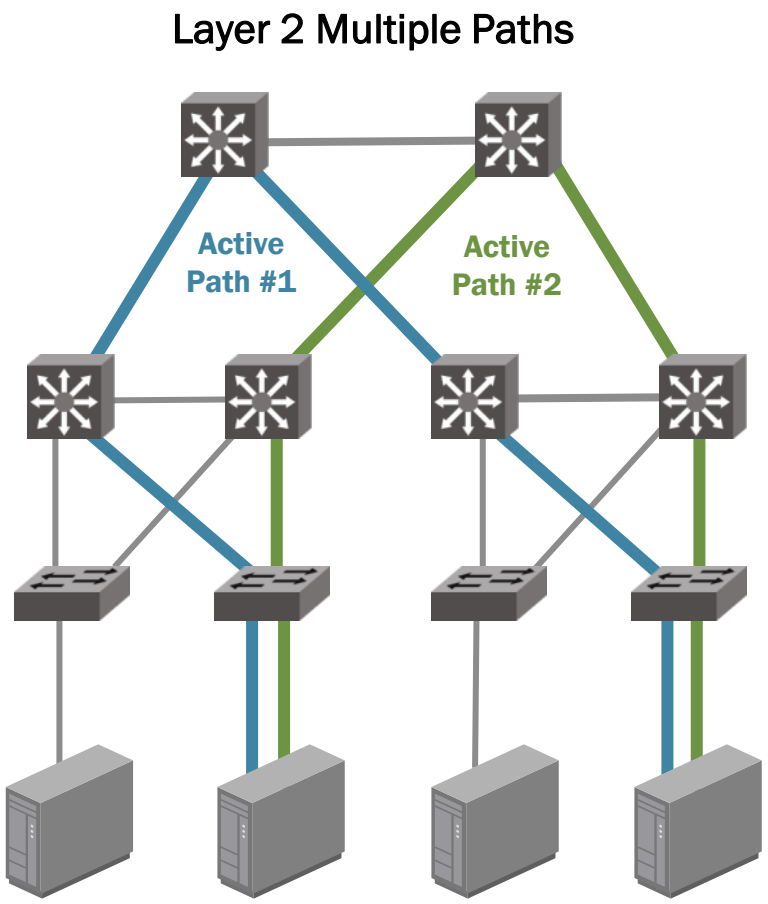

- Multi-path Layer 2 switching
  - All paths are active and traffic is distributed across all paths
  - Fully utilize all fabric bandwidth
- Establishes shortest paths through the Layer 2 fabric
- Uninterrupted response to link failures
- Backward-compatible and connects into existing infrastructures
- Delivers multiple hops for all traffic types (including FCoE)

Logical Chassis Details

- Fabric auto-configures
  - No fabric configuration necessary
- Fabric behaves/managed as a single logical chassis
  - Aggregation (or Core) layer sees one logical element
  - Fabric members act like a blade in a chassis
- Logically flattens and collapses network layers
  - Fabric is self-aggregating
  - Flexible fabric topologies
- Scales without added management complexity

Logical Chassis Details

Single Logical Switch Behavior

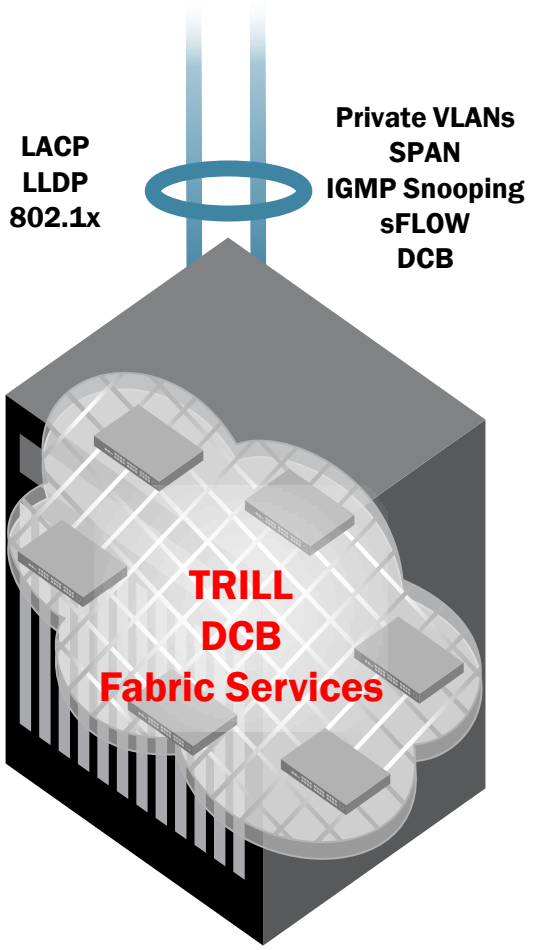

- Behaves like a transparent LAN service
  - For example, BDPUs in STP environments are passed through the fabric
- Fabric protocols used within the fabric
  - TRILL, DCB, Fabric Services, etc.
- Industry-standard protocols used to communicate outside the fabric
  - LACP, 802.1x, sFLOW, etc.

Distributed Intelligence Details

- Distributed Fabric Services
  - Fabric is self-forming
  - Information shared across all fabric members
  - Fabric is aware of all devices connected
- Masterless Control
  - Switch or link failure does not require full fabric reconvergence

- Shared Port Profiles information
  - Automatic Migration of Port Profiles (AMPP)
  - Enables seamless VM migration without compromise

Distributed Intelligence Details

Automatic Migration of Port Profiles (AMPP)

Allows VM to move with the network automatically reconfiguring

1. Port Profiles created, managed in fabric;
2. Configuration is Distributed
3. Server admin binds VM MAC address to Port Profile ID
4. MAC address/Port Profile ID association ; sent to fabric

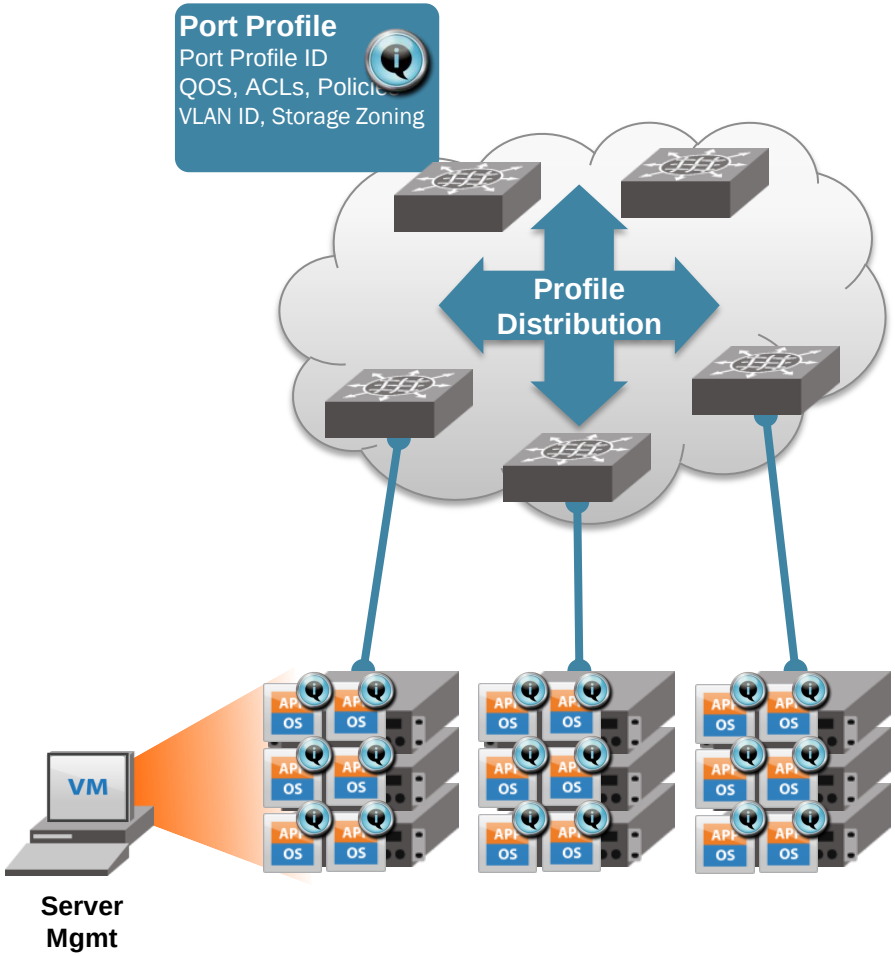

Adding a new switch to an Ethernet Fabric

Automatic Fabric Creation and Expansion

Automatic Trunk Creation

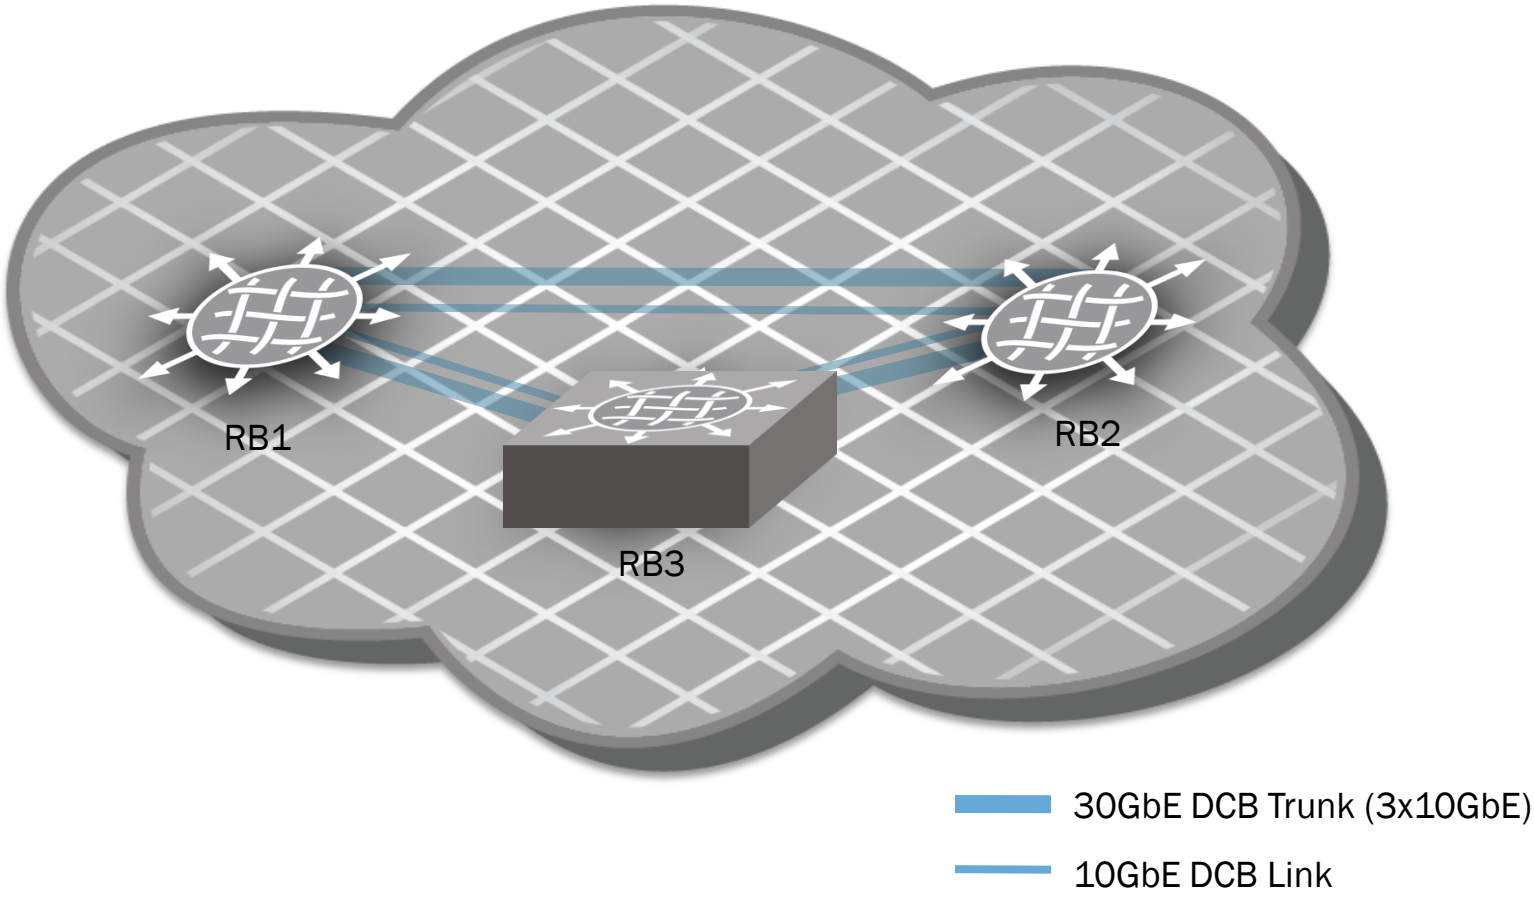

Ethernet Fabrics

Distributed MAC address tables

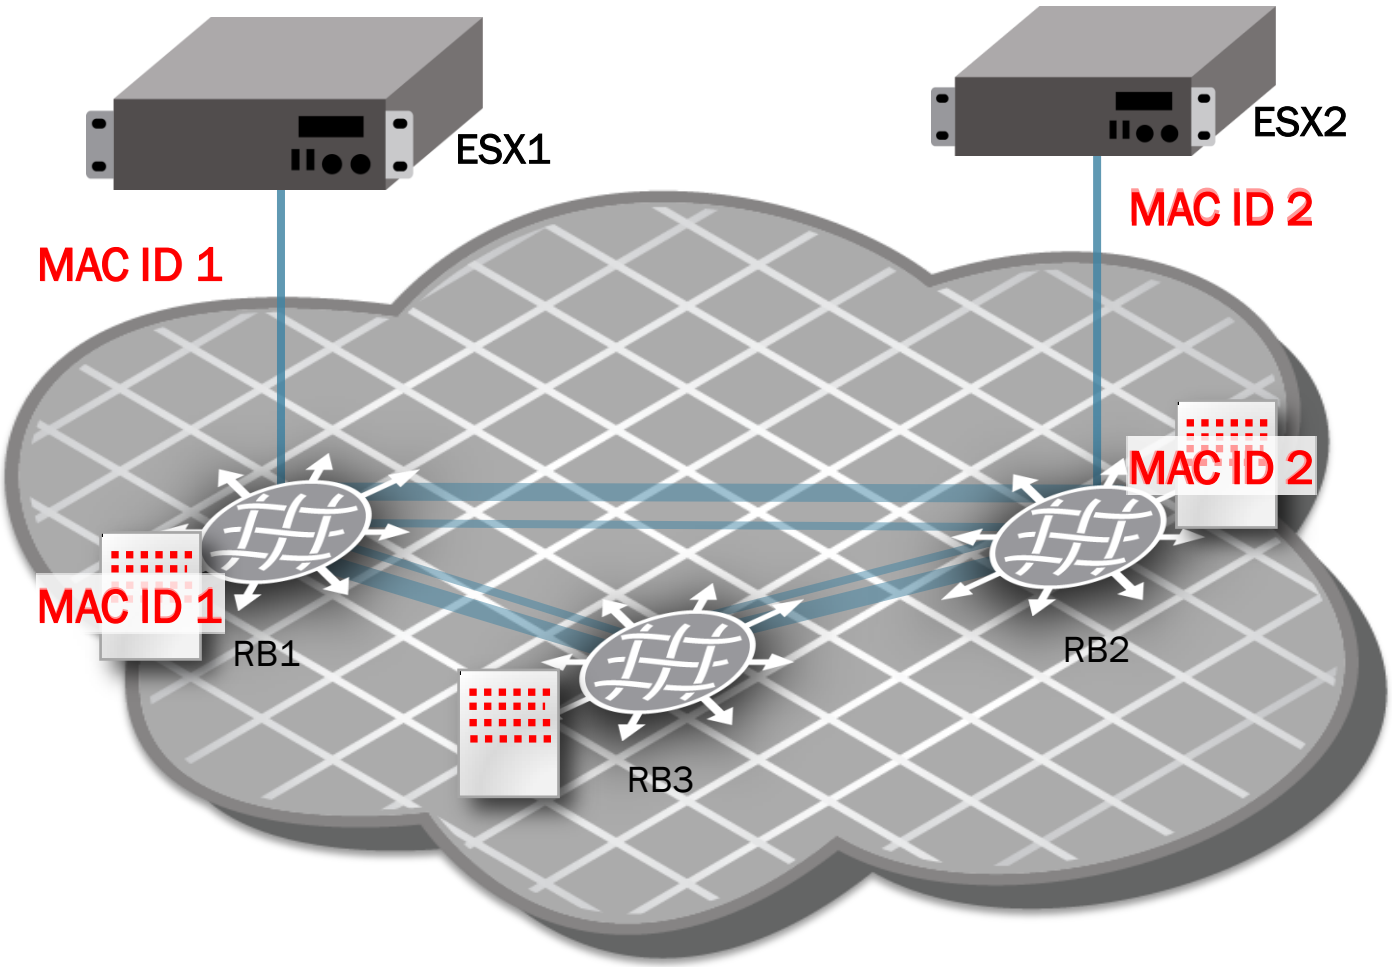

Ethernet Fabrics

Host Based LAG – LAG to different physical VDX Switches

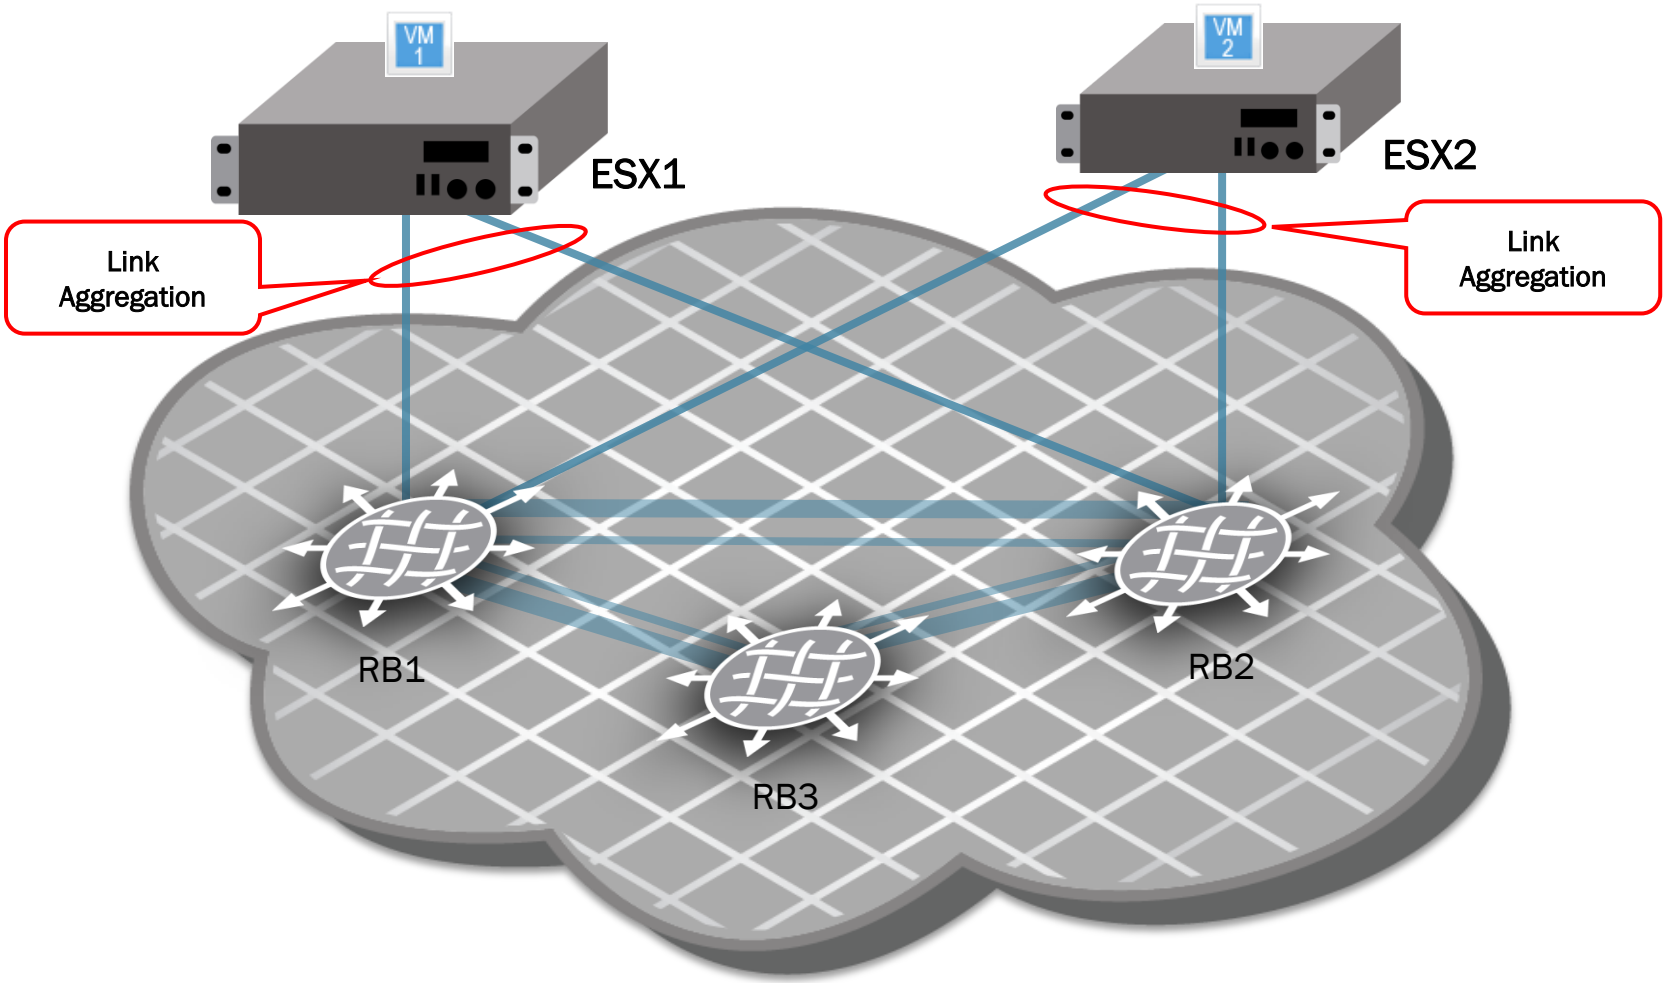

Ethernet Fabrics

Equal Cost Multi-Pathing

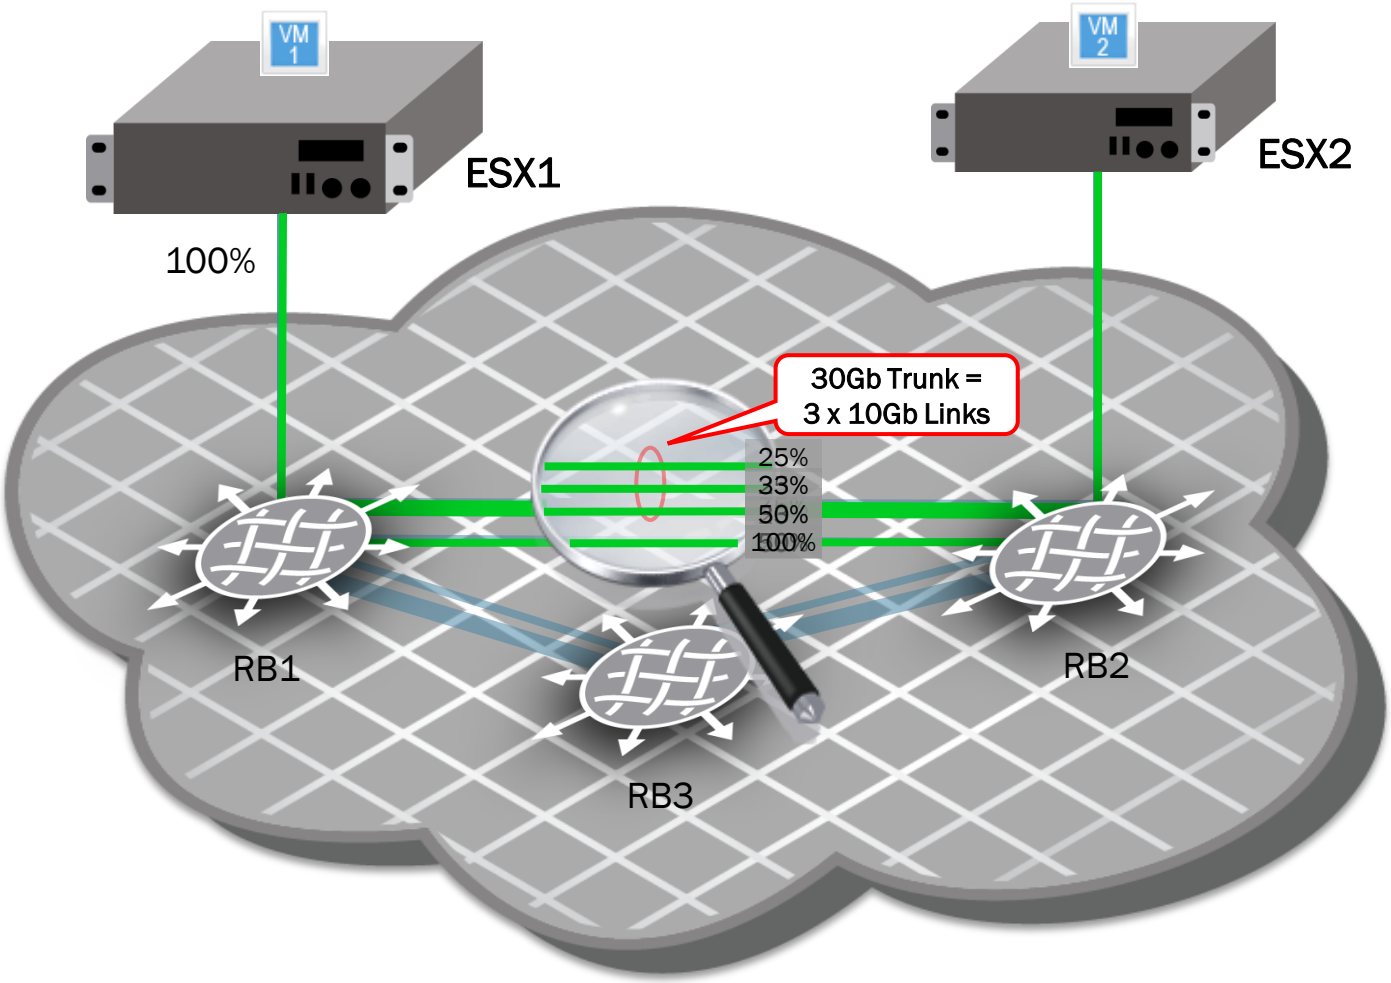

VCS – Known unicast data path (TRILL)

Trill Encapsulation

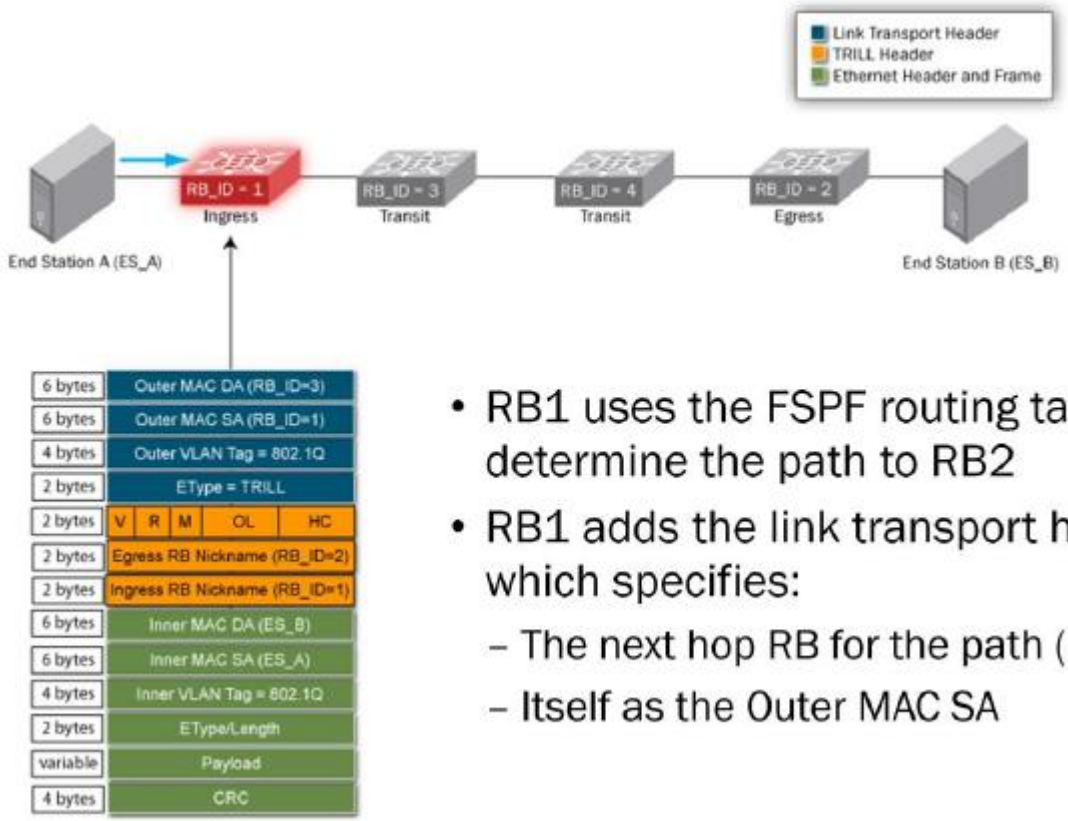

- RB1 uses the FSPF routing table to determine the path to RB2
- RB1 adds the link transport header which specifies:
  - The next hop RB for the path (RB3)
  - Itself as the Outer MAC SA

Brocade VCS Use Case example 1

1/10 Gbps Top-of-Rack Access – Architecture

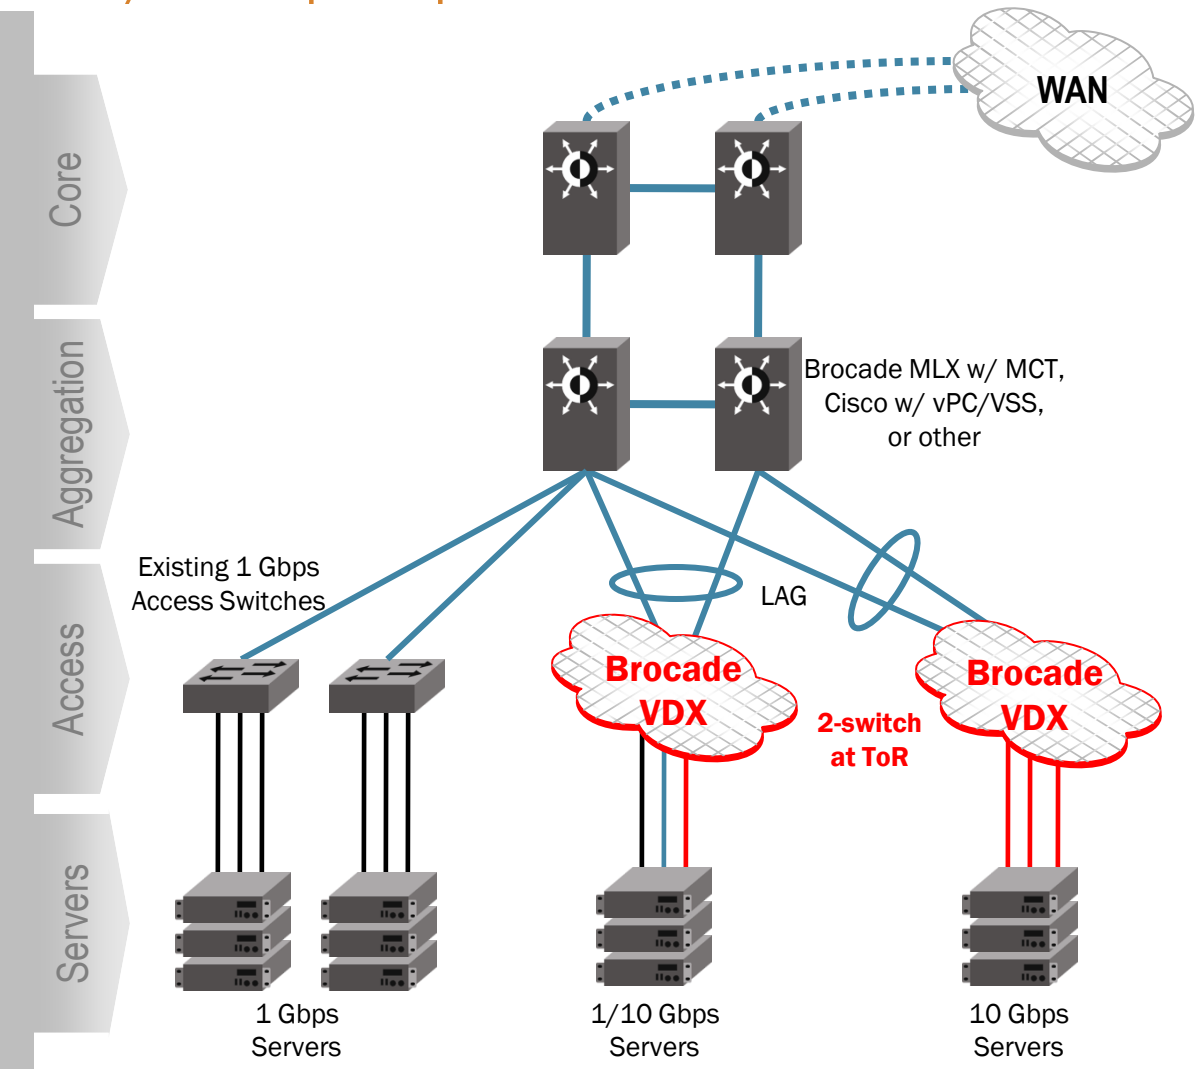

- Preserves existing architecture
  - Leverages existing core/agg
  - Coexists with existing ToR switches
- Supports 1 and 10 Gbps server connectivity
- Active-active network
  - Load splits across connections
- No single point failure
  - Self healing
- Fast link reconvergence
- High-density access with flexible subscription ratios
  - Supports up to 36 servers per rack with 4:1 subscription

Brocade VCS Use Case example 1

1/10 Gbps Top-of-Rack Access – Topology

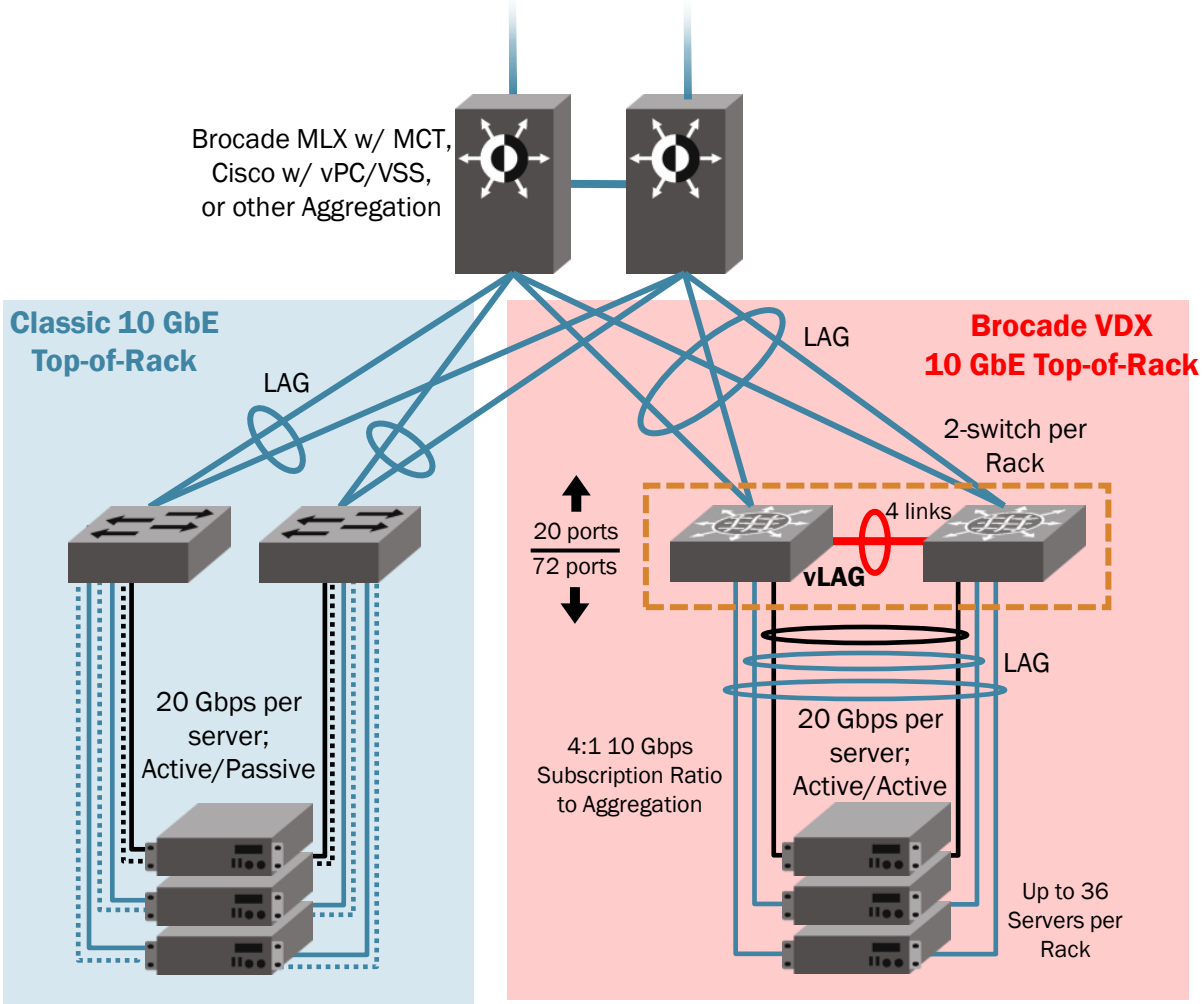

- Active/Active multi-homed server connections
  - Servers only see one ToR switch
  - Half the server connections
- Reduced switch management
  - Half the number of logical switches to manage
- Unified uplinks
  - One virtual LAG per VCS

	Classic ToR	VCS ToR
Utilization	Active/ Passive	Active/ Active
Connections per Server	4	2
Logical Switches per Rack	2	1
LAG per Rack	2	1

Brocade VCS Use Case – example II

1/10 Gbps Access; Collapsed Network – Architecture

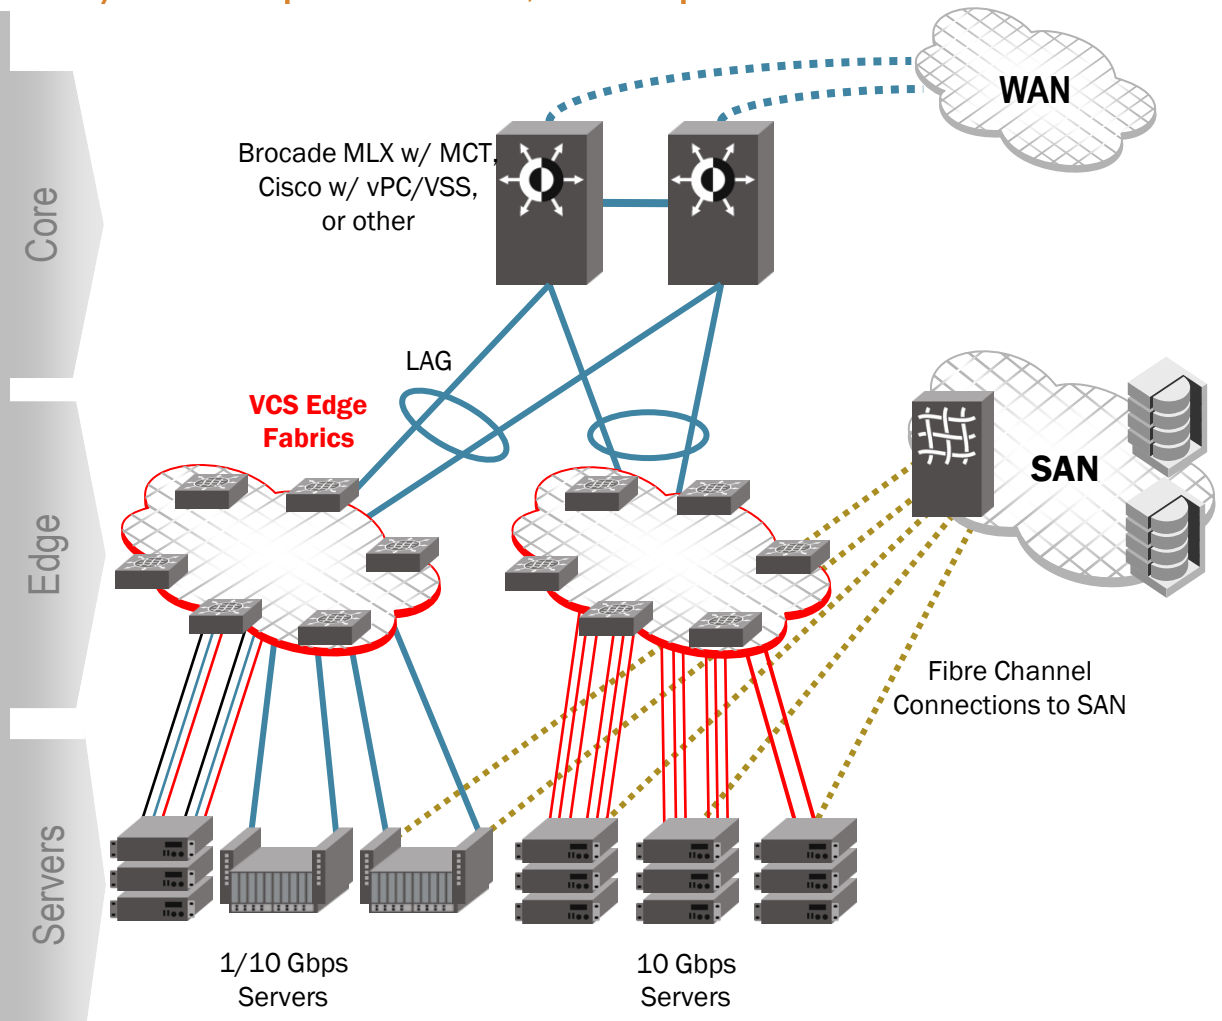

- Flatter, simpler network design
  - Logical two-tier architecture
  - Ethernet fabrics at the edge
- Greater Layer 2 scalability/flexibility
  - Increased sphere of VM mobility
  - Seamless network expansion
- Optimized multipath network
  - All paths are active
  - No single point failure
  - STP not necessary

Brocade VCS Use Case – example II

1/10 Gbps Access; Collapsed Network – Topology – Clos Fabric

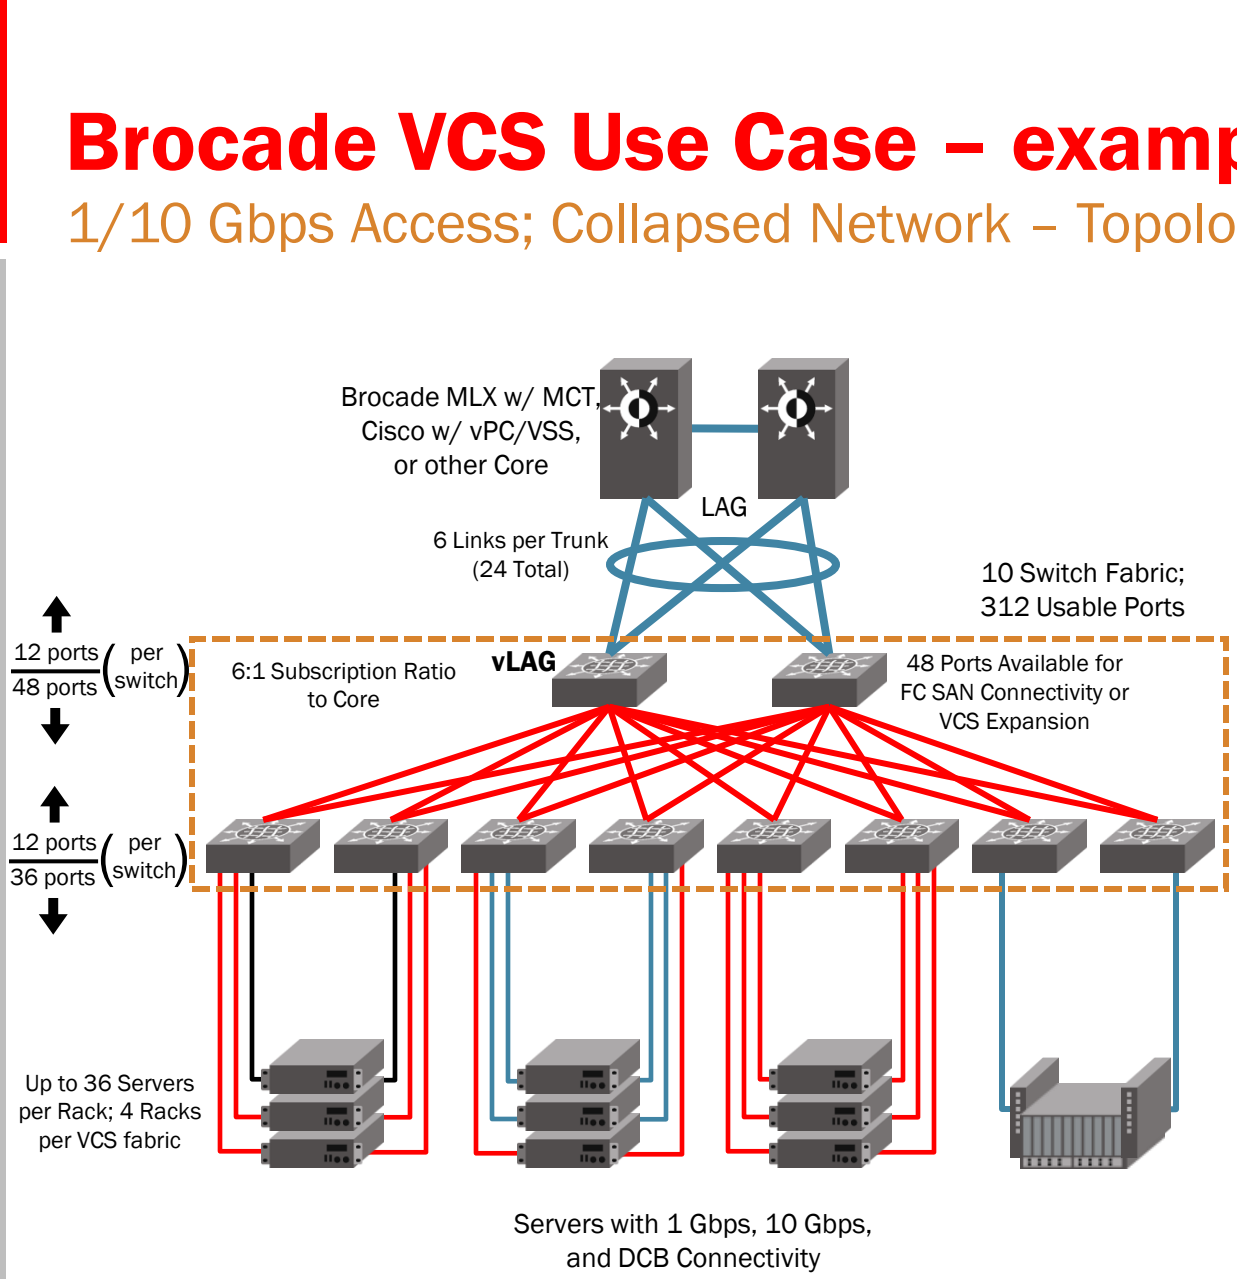

Scale-out VCS edge fabric

- Self aggregating, flattens the network
- Clos Fabric topology for flexible subscription ratios
- 312 usable ports per 10-switch VCS fabric
- Supports 144 servers in 4 racks, all with 10 Gbps connections

Drastic reduction in management

- Each VCS managed as a single logical chassis

Enables network convergence

- DCB and TRILL capabilities for multihop FCoE and enhanced iSCSI

Brocade VCS Architecture

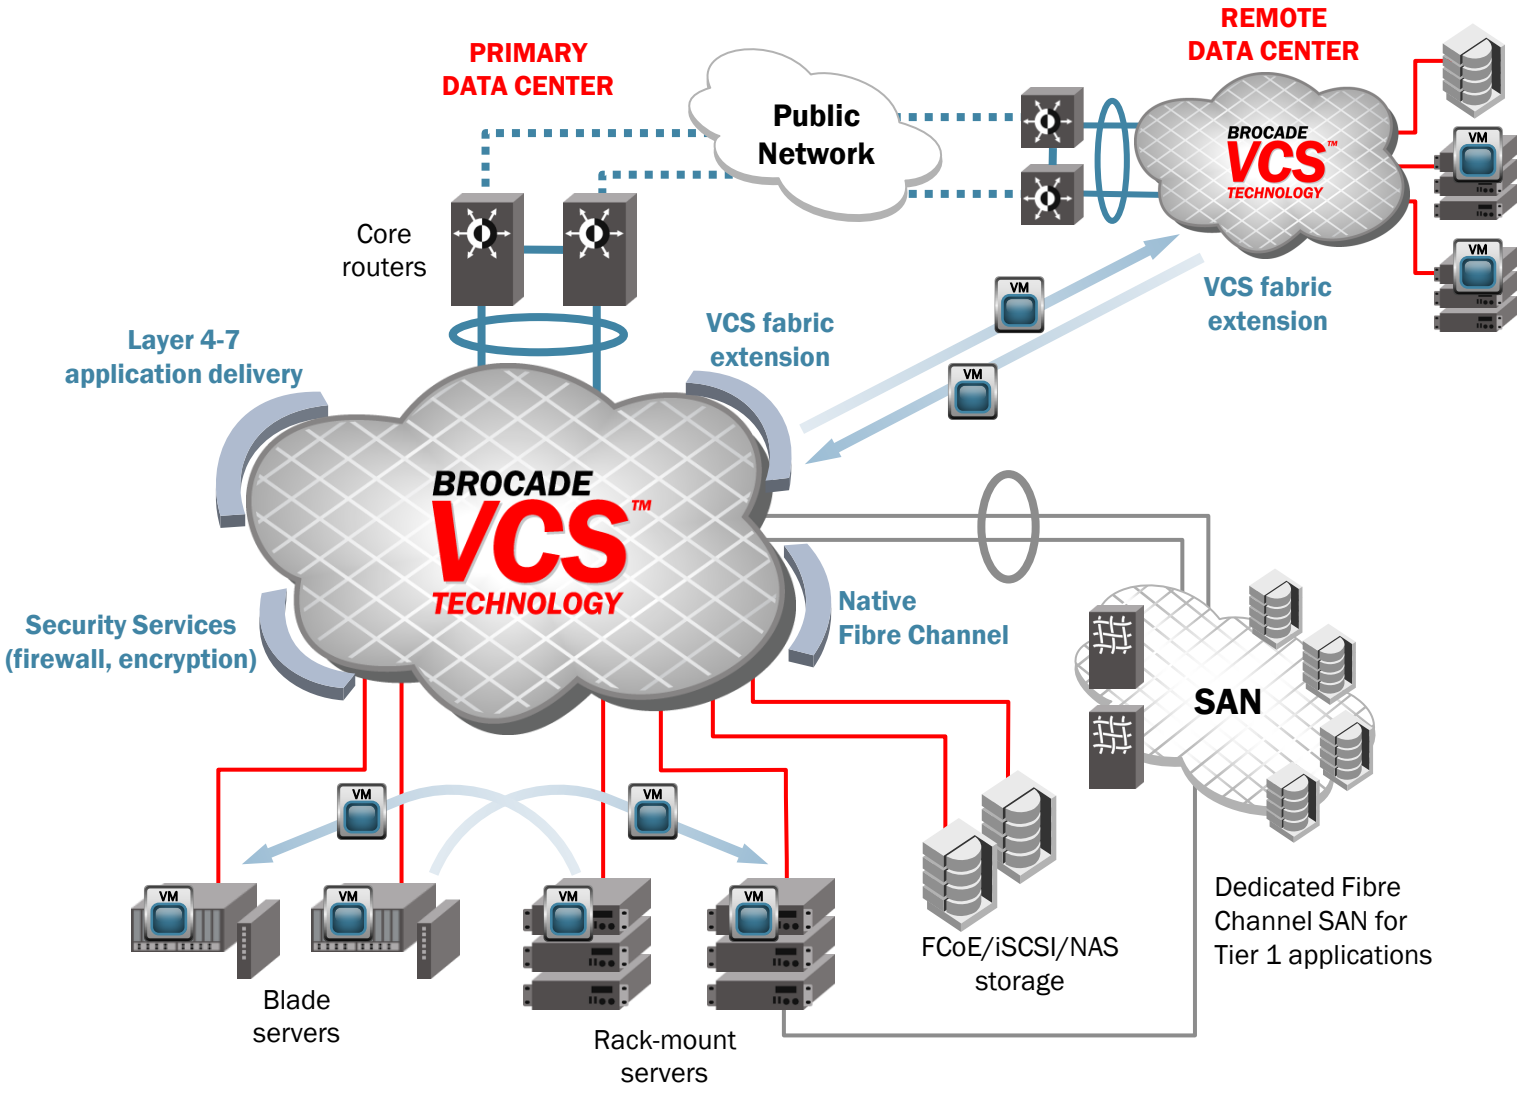

Brocade's Answer
VDX 6720 Ethernet Fabric Switches

- **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 VCS Architecture

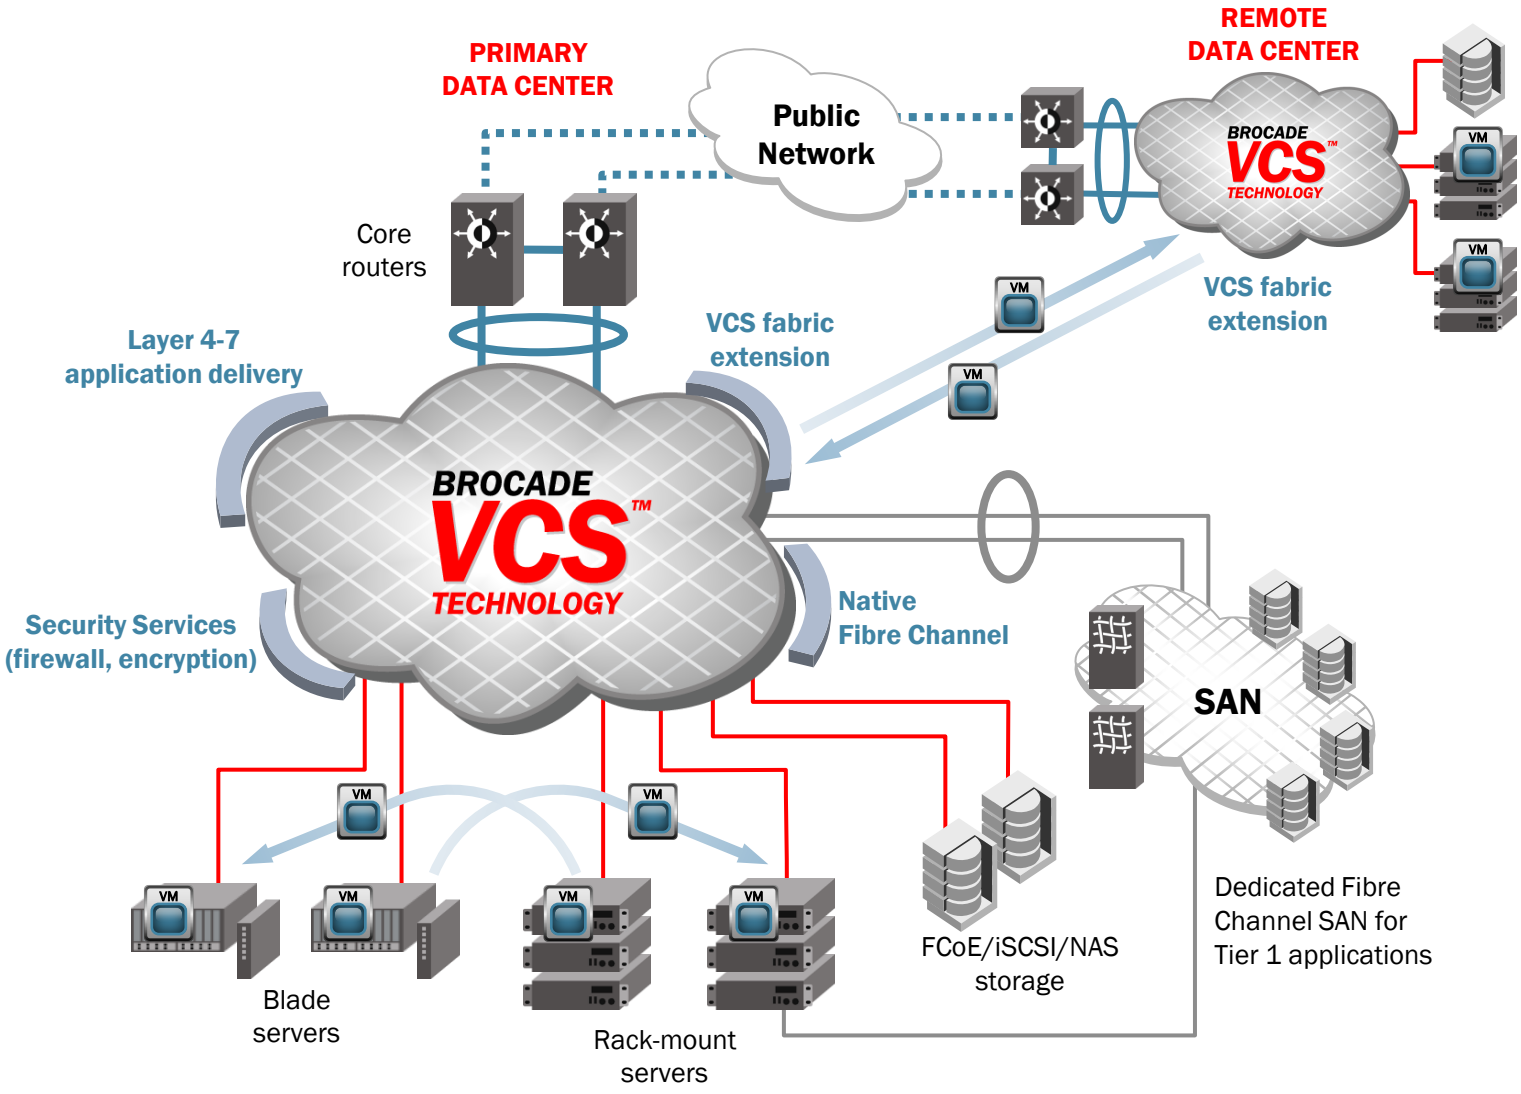

Brocade VDX Product Family

Delivering virtual cluster switching

BROCADE VCS

A new family of Ethernet Fabric switches

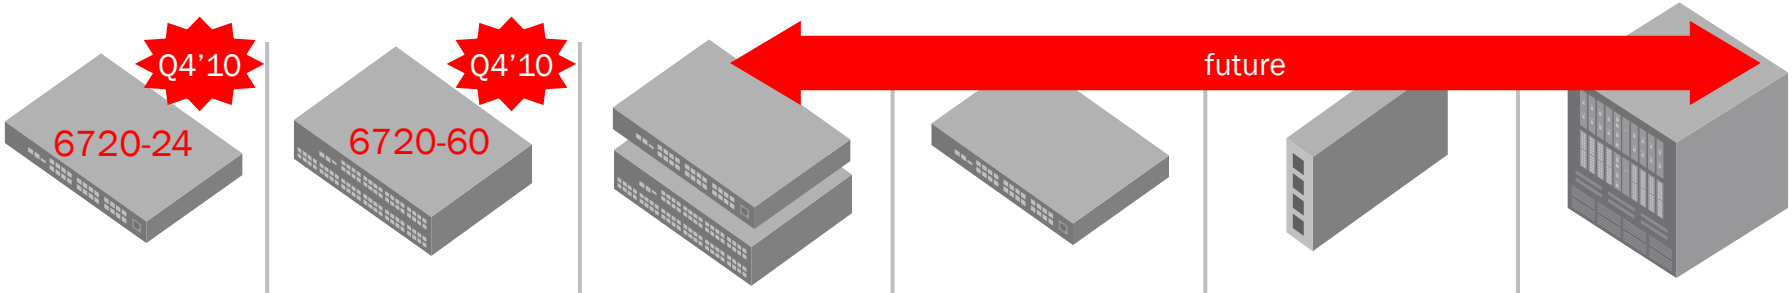

24-port	60-port	24- and 60-port	48-port	VCS in blade server chassis	Wire-speed chassis with VCS
1/10 GbE	1/10 GbE	1/10 GbE	1 GbE	1/10 GbE	Allows Ethernet fabrics to scale further
600 ns latency	High density	Fibre Channel ports for connectivity to SAN	High-density 1 GbE VCS connectivity	Blade switches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

Virtual Cluster Switching (VCS)

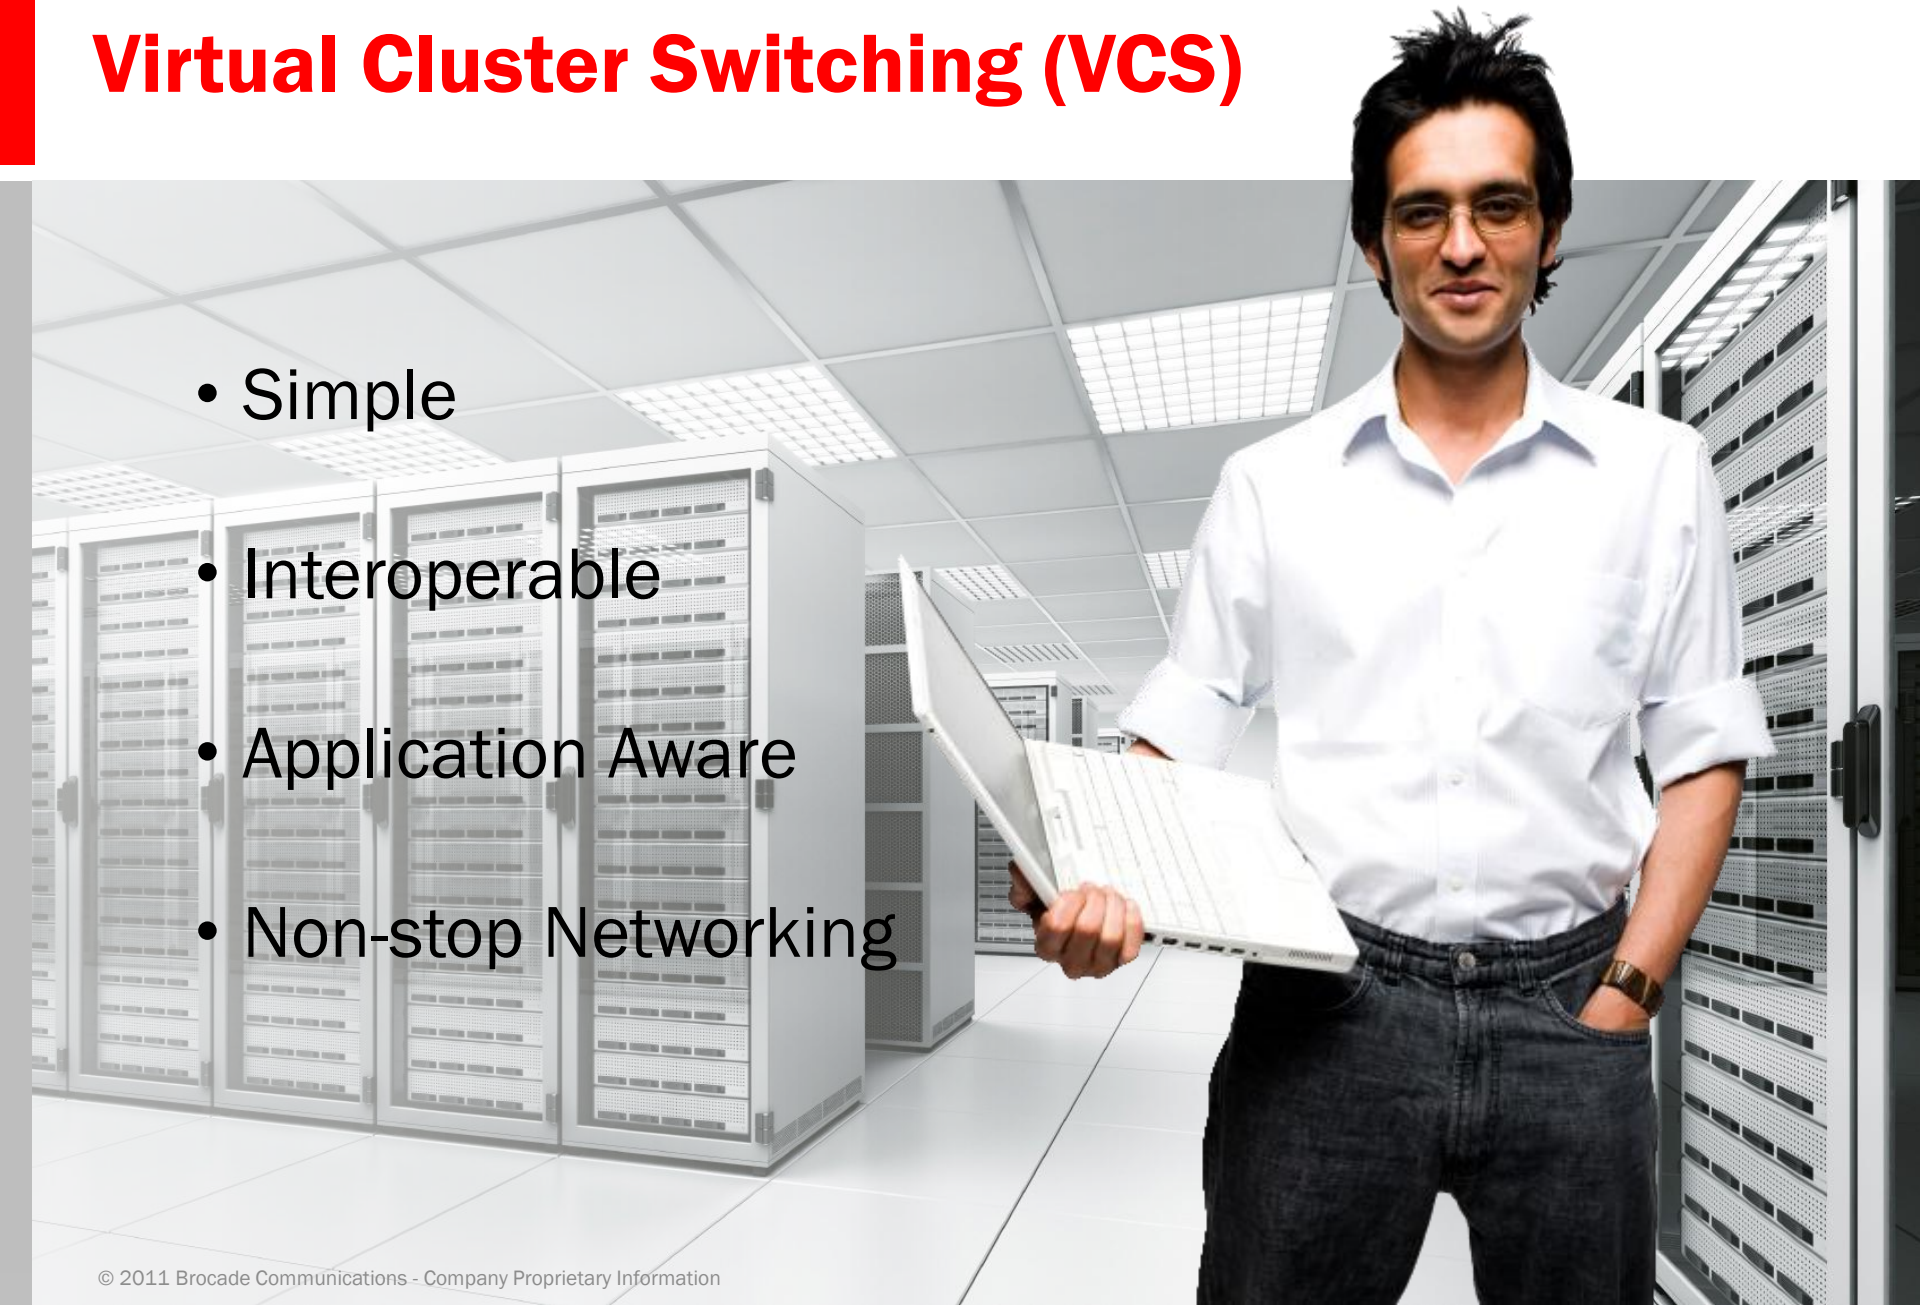

- Simple
- Interoperable
- Application Aware
- Non-stop Networking

BROCADE

WHEN YOU THINK NETWORKS, THINK BROCADE

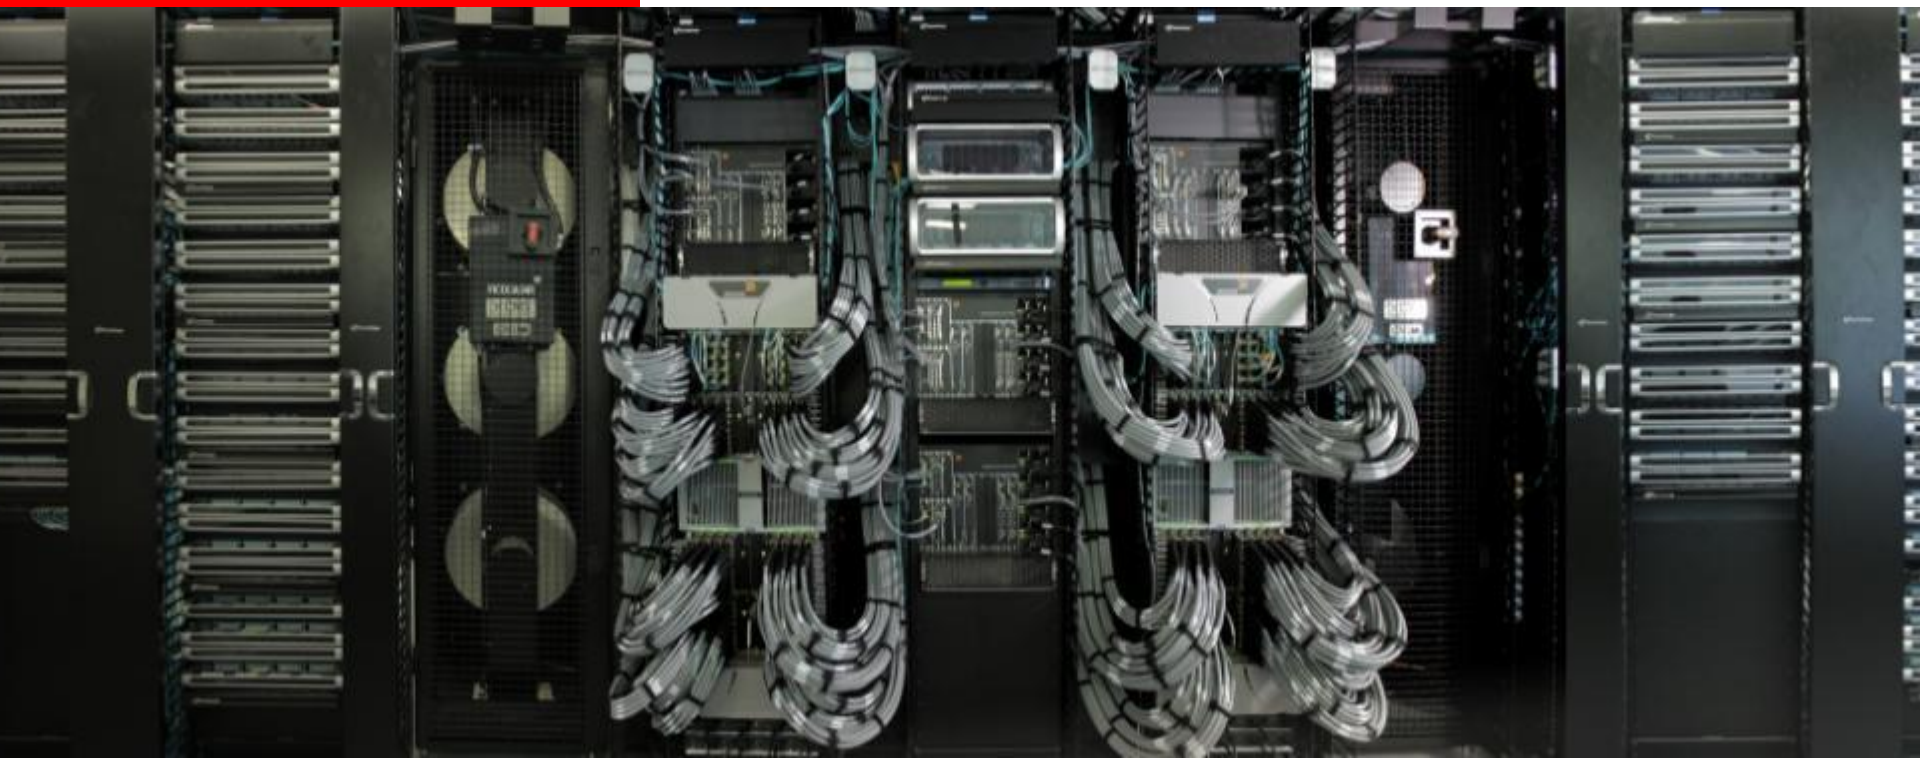

Thank You