



DataCenter Architecture & Design

Julio Alba (jalba@brocade.com)

IP SE España

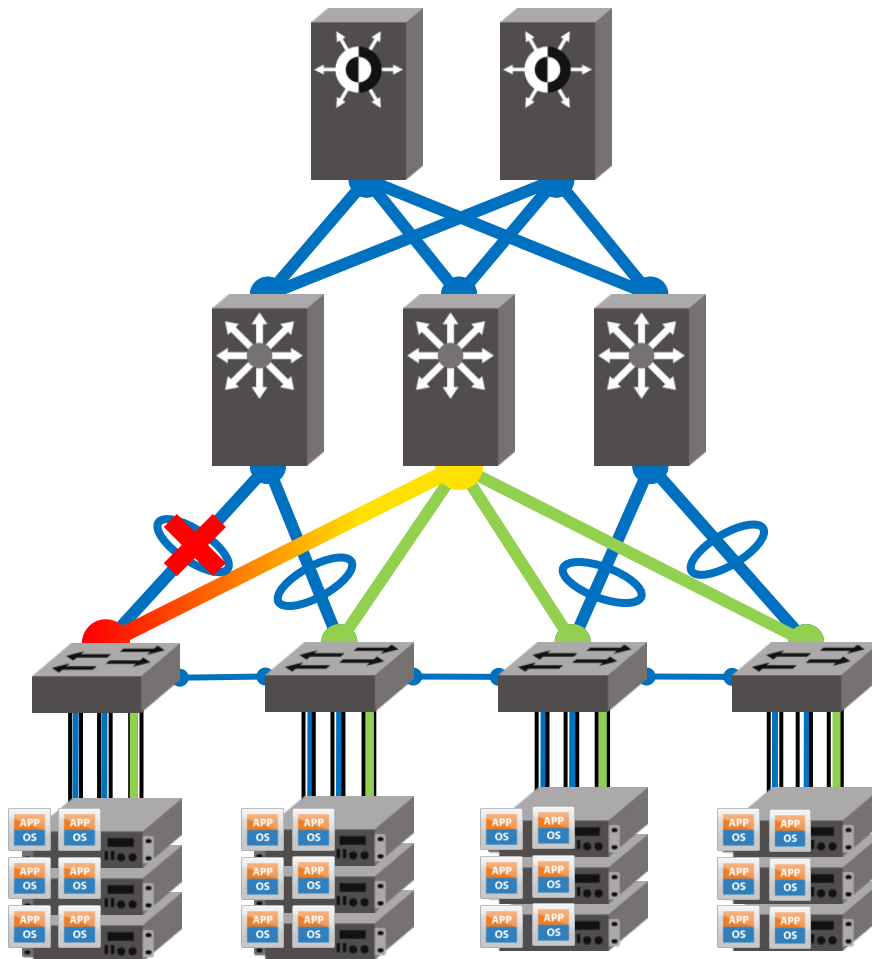
October 2010

Addressing Today's Challenges



Scaling Virtual Server Environments

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

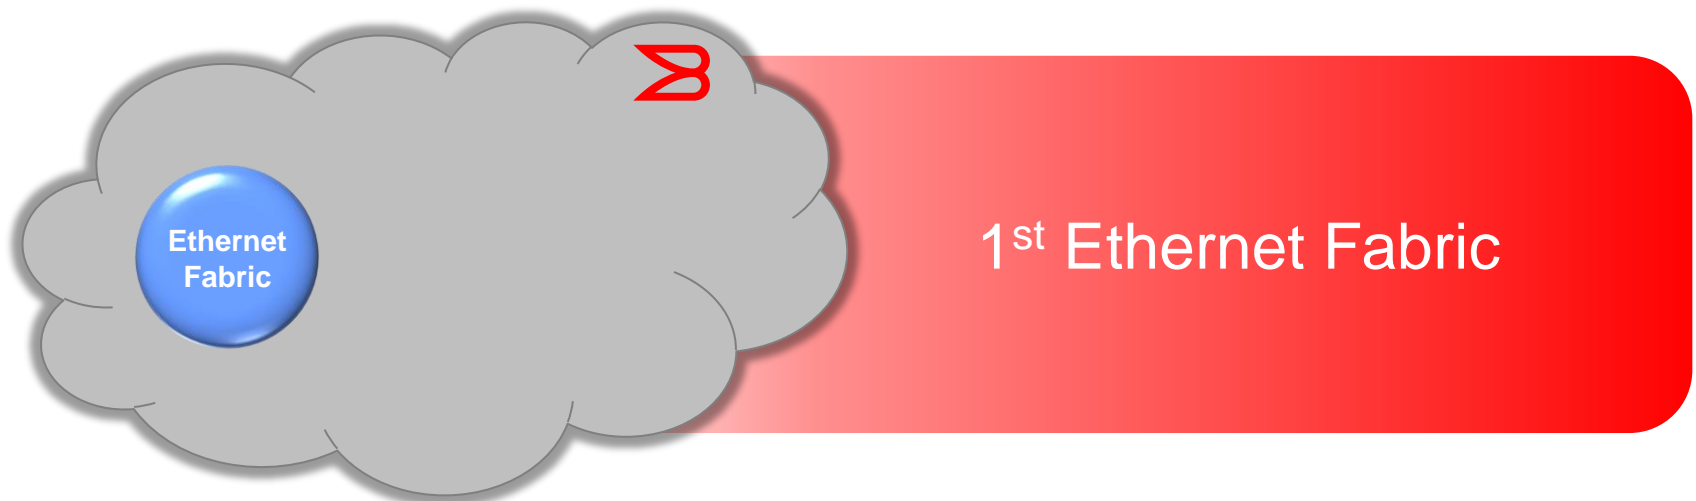
Layer 3 as an alternative

Greater complexity; higher cost

VM mobility limited to rack

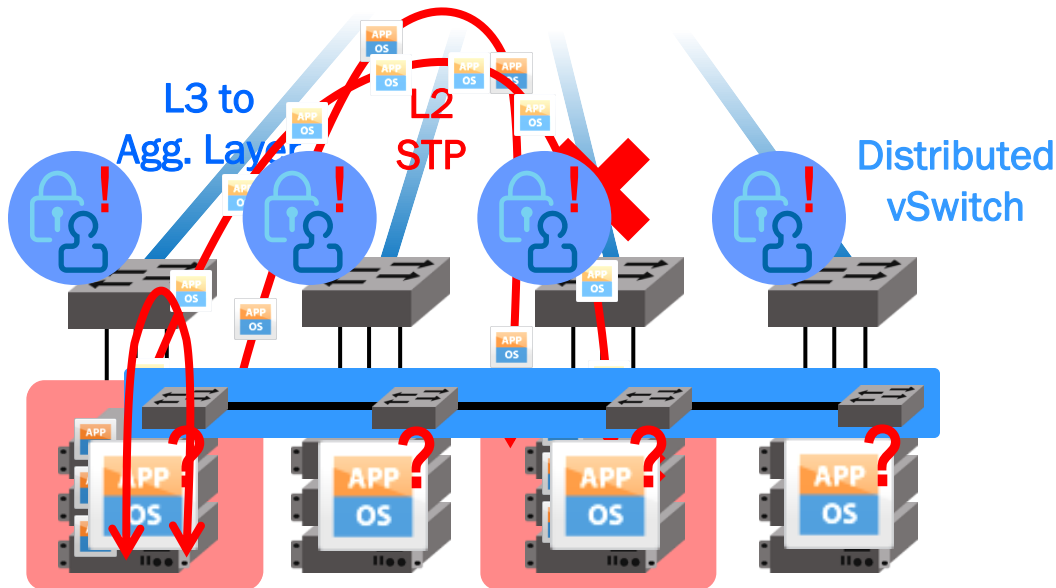
Imagine if...

- There was no requirement for STP in Layer 2 networks
- All paths in the networks were utilized with traffic automatically distributed
- Link failure did not result in a temporary outage and paths were always deterministic
- The network provided low latency, lossless transmission and could carry both IP and storage traffic, without compromise



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

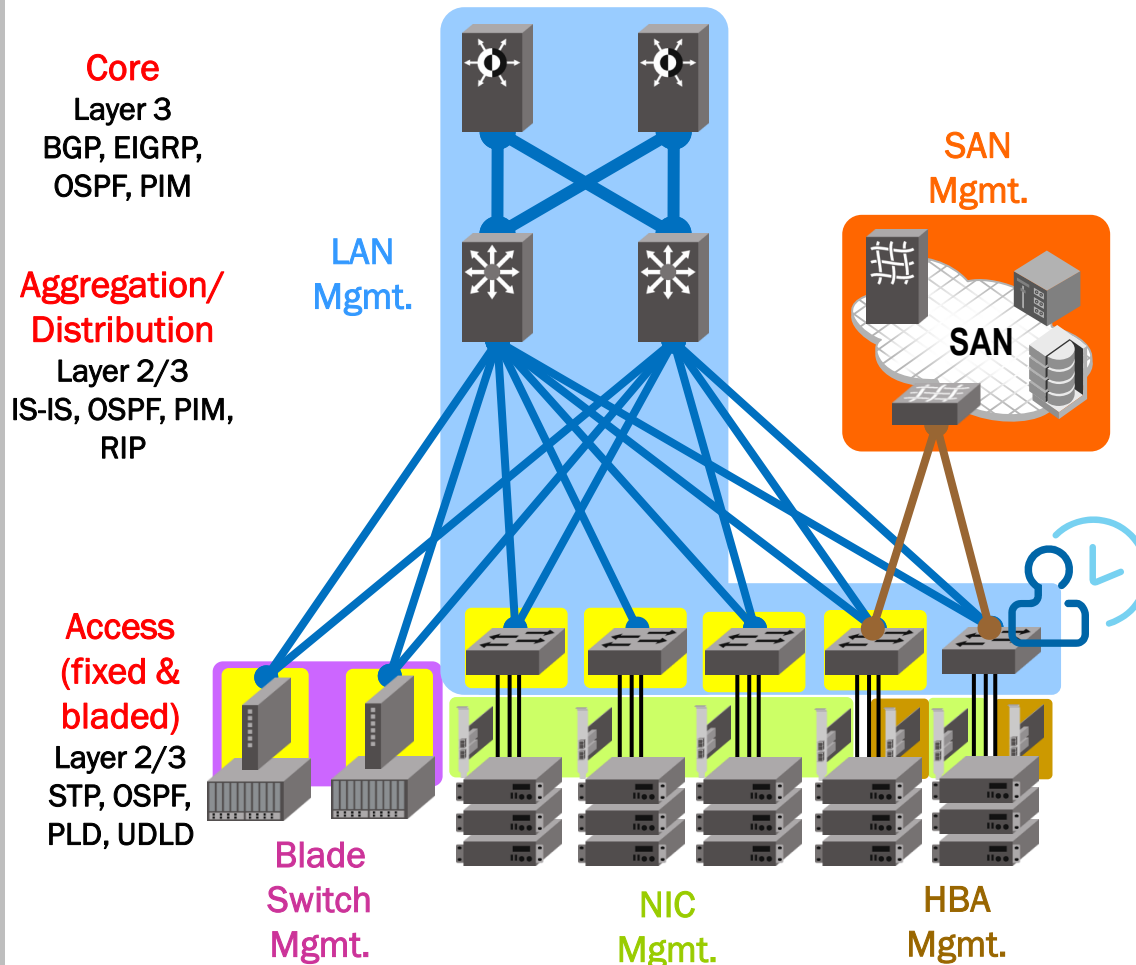
Imagine if...

- There were no physical barriers of VM migration
- Your network was aware of all VMs at all times
- Mobility did not come with a cost in compute resources
- You could leverage your entire server environment to maximize application performance and availability



Network Management

Challenges Today



Too many network layers

Utilize many L2/L3 protocols

Lots of small-form-factor switches at the edge

Each switch has to be managed

Because of the number, they need to be aggregated

Configuration time when deploying new switches

Switch has to be set up

Network settings must be configured

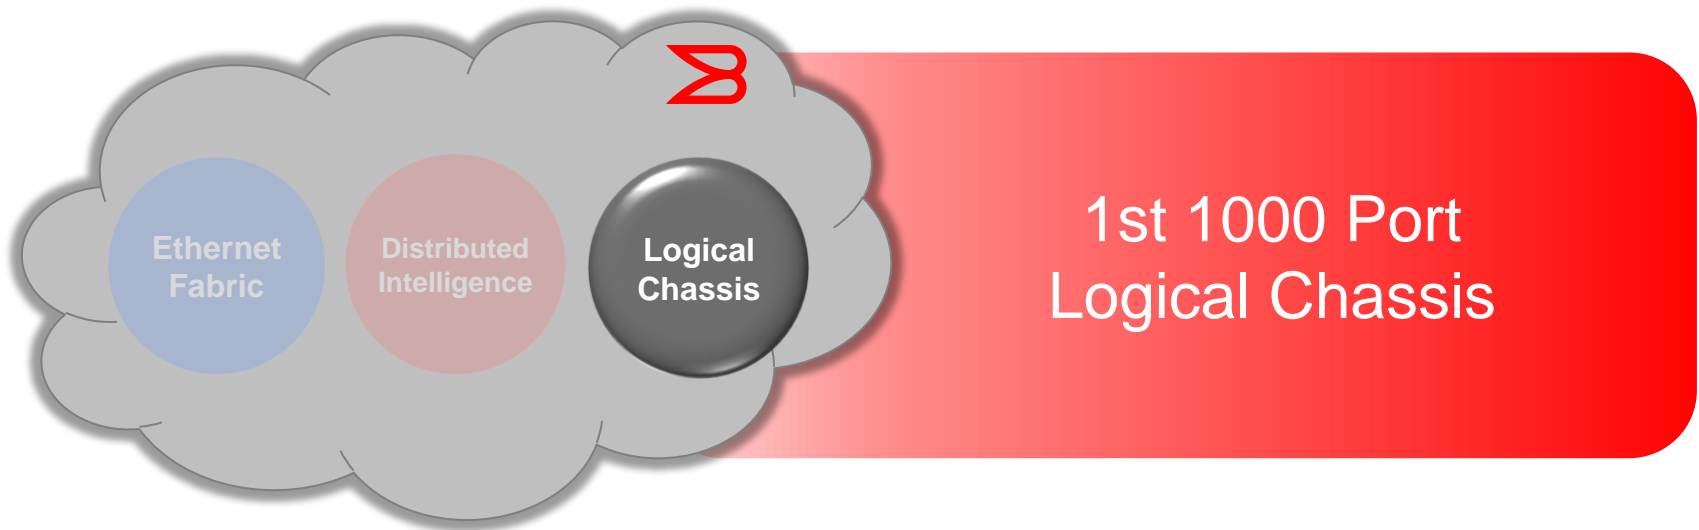
Separate management tools for LAN, SAN, NICs/HBAs

Management silos do not fit in a virtualized data center

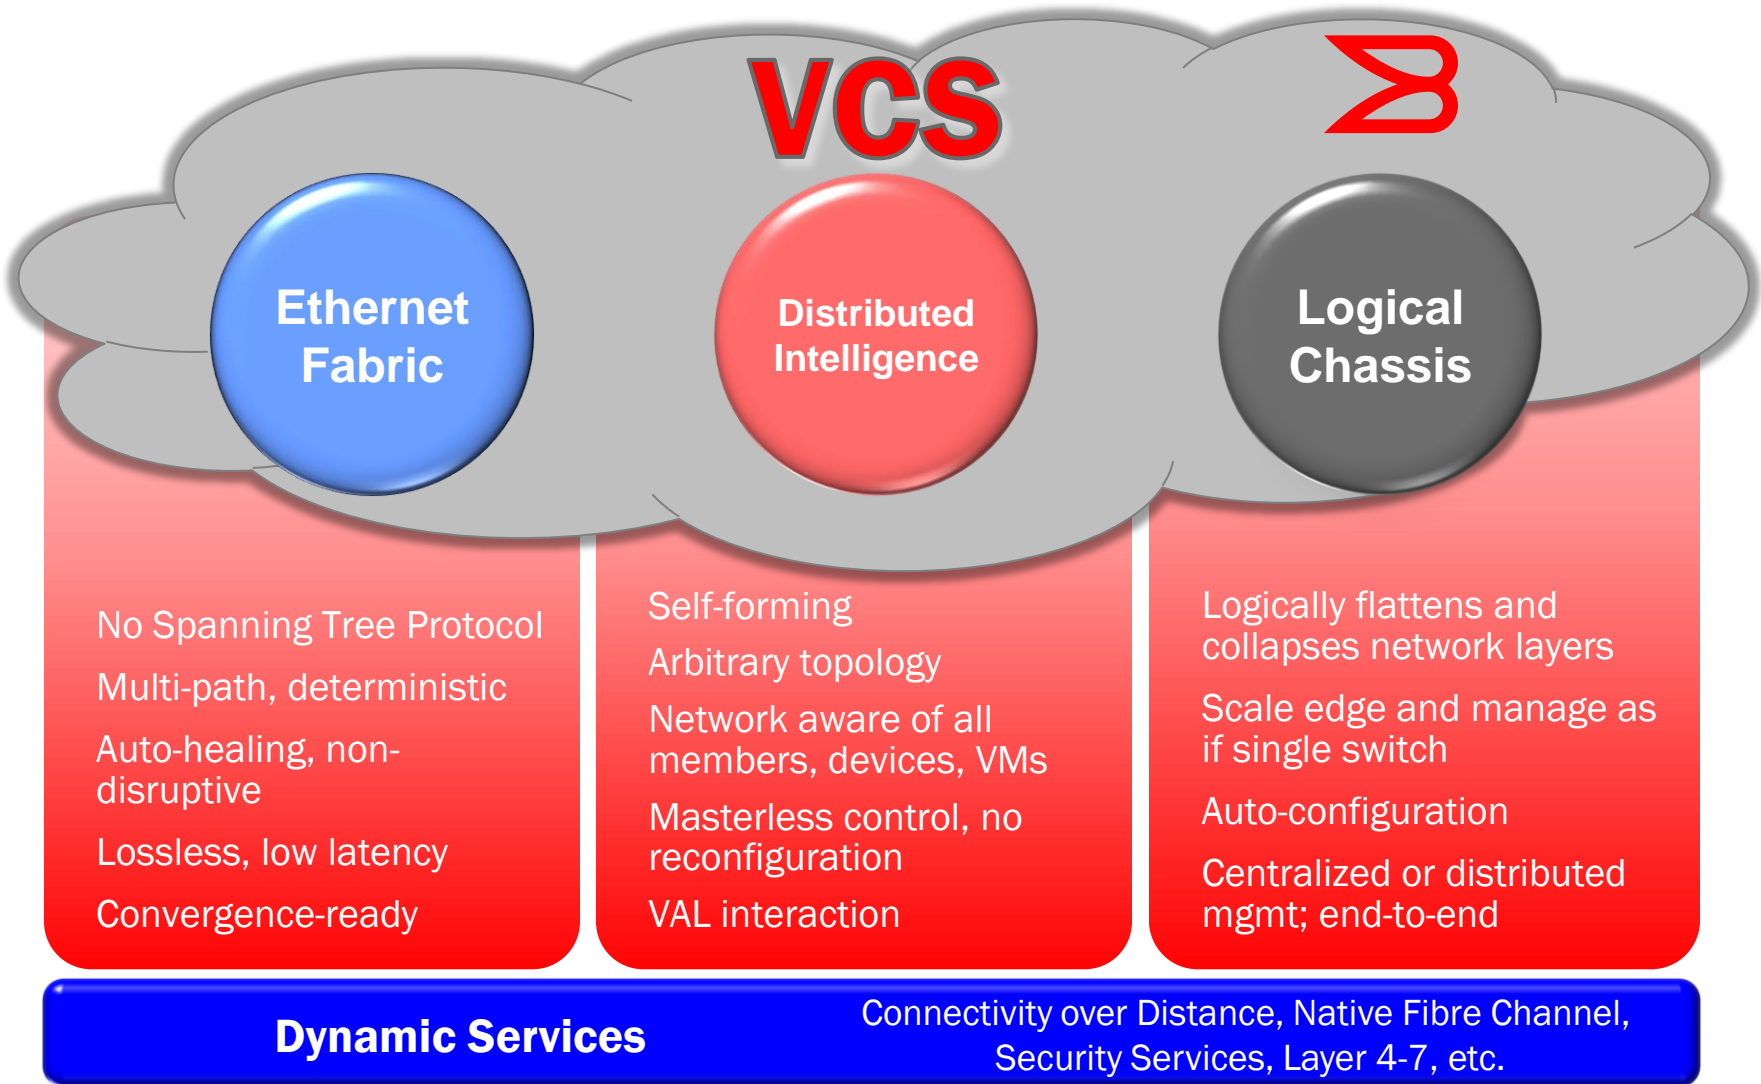
Drives up OpEx

Imagine if...

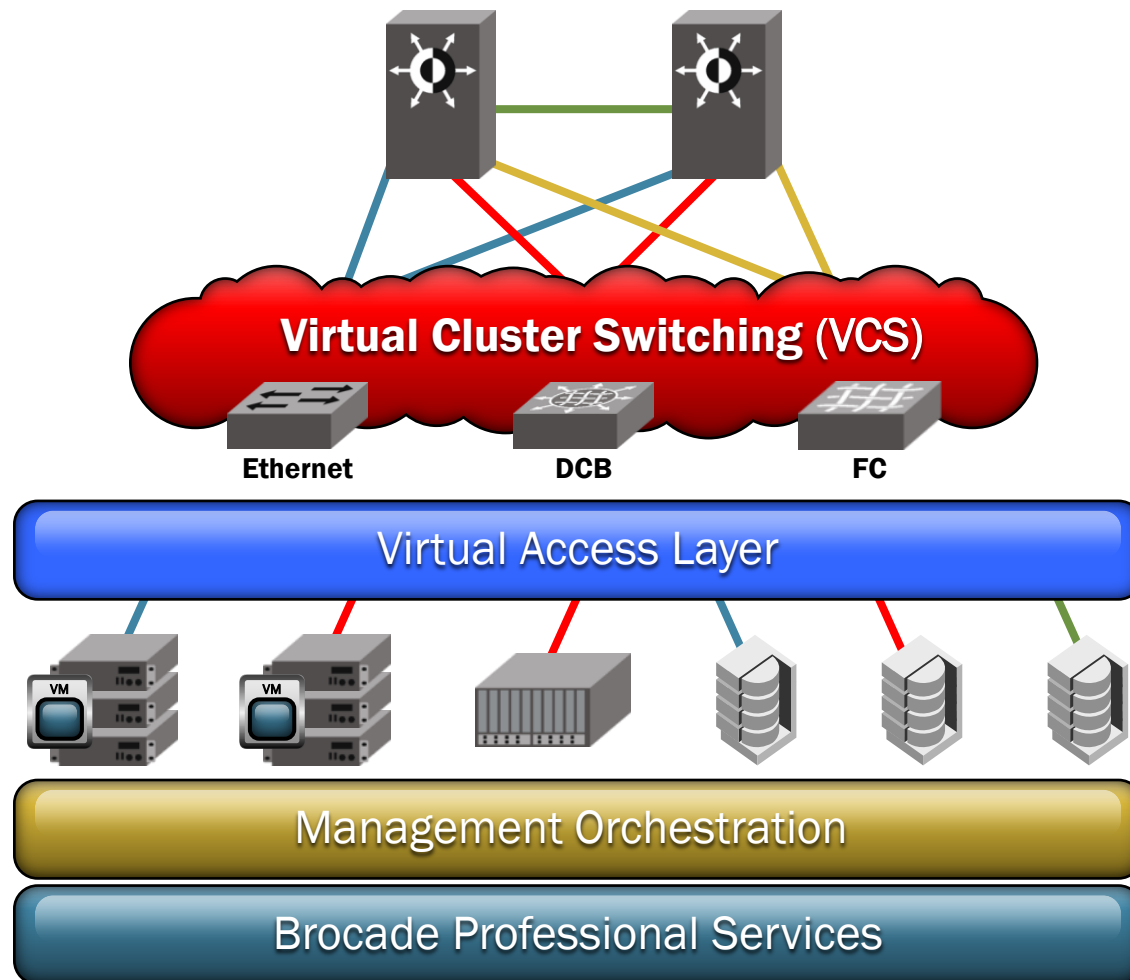
- You could logically eliminate a layer of the network
- You could connect 10, 20... edge switches and manage them as one
- You could scale the network without added complexity
- There was a common tool to manage all components of the SAN and LAN



Virtual Cluster Switching (VCS)



Data Center Network Vision



Collapse/simplify
layered architecture

Map to application
needs

Network innovation

Revolutionize Layer 2
connectivity (VCS)

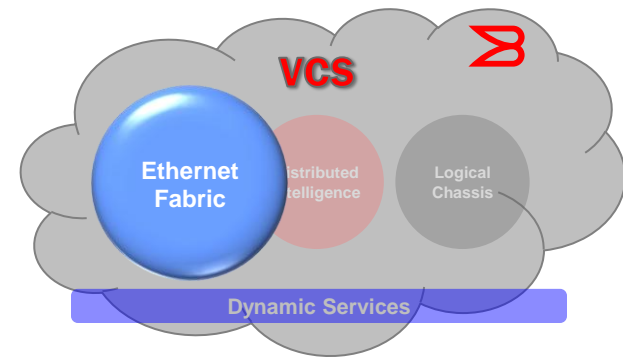
Optimize VM
performance and
mobility (VAL)

Single pane of glass
management

Virtual Cluster Switching (VCS) Details



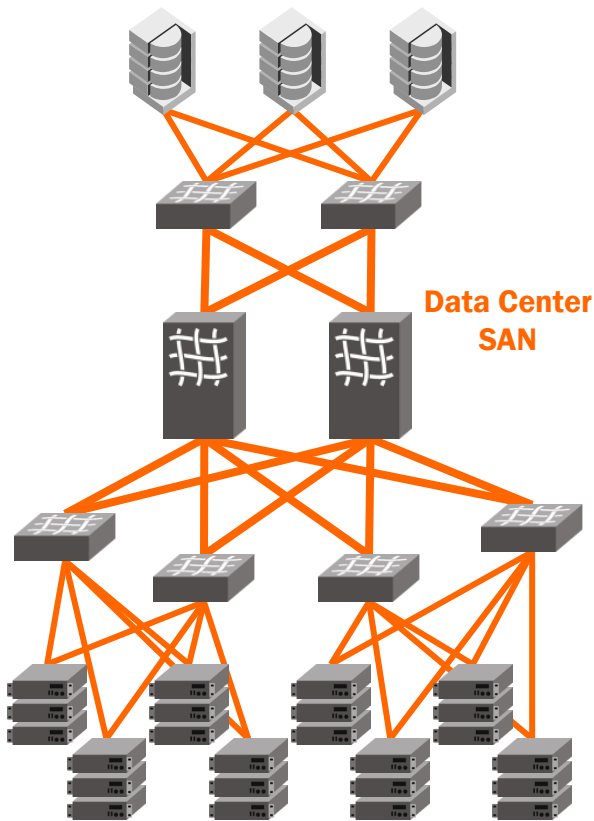
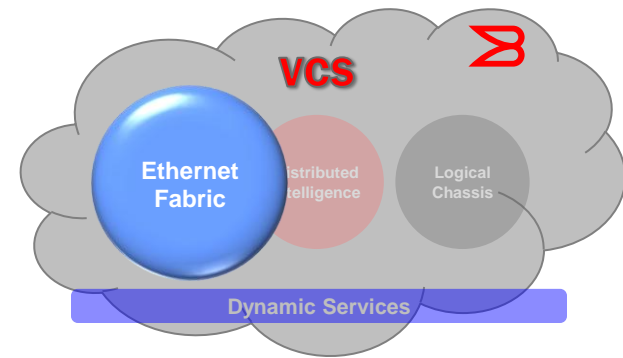
Ethernet Fabric Details



- 1st true 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, sub-250ms link recovery
- LAN/SAN Convergence Ready
 - FCoE and iSCSI traffic
- Standards-based
 - Extends existing Ethernet infrastructure

Ethernet Fabric Details

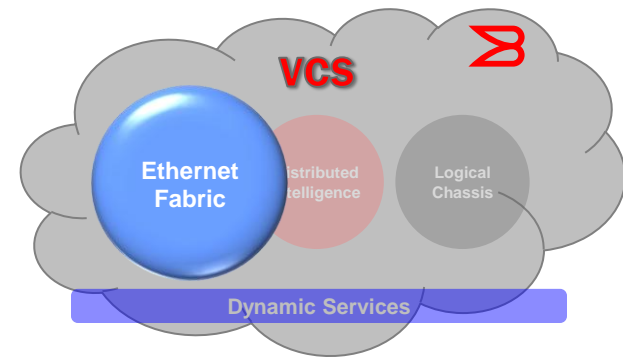
What is a Fabric?



- Common fabric attributes
 - Switched network
 - Fabric members and devices connected always know about each other
 - All paths are available for high performance and high reliability
 - Traffic travels across the shortest path
 - Traffic can be routed from fabric to fabric
- Brocade is the expert in building data center fabrics
 - In 90+% of global 1000 data centers

Ethernet Fabric Details

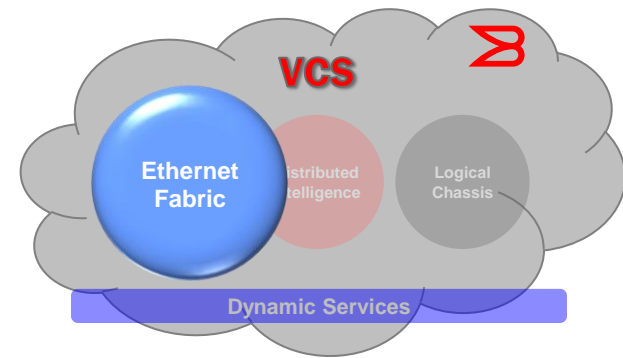
Data Center Bridging (DCB)



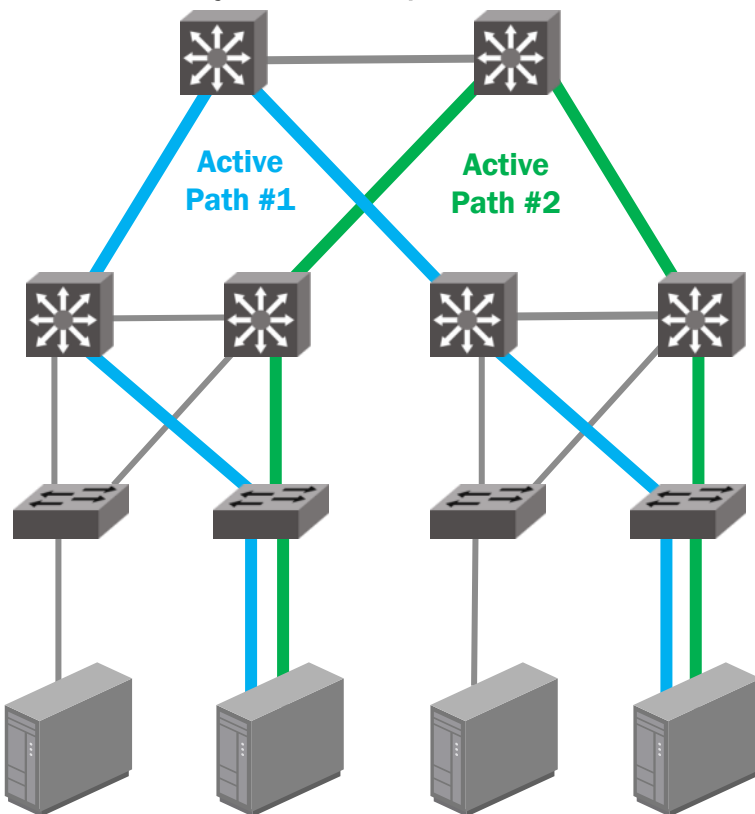
- Making Ethernet Lossless
- 802.1Qbb – Priority-Based Flow Control
 - PFC: Allows Identification and prioritization of traffic
- 802.1Qaz – Enhanced Transmission Selection/Data Center Bridging Exchange
 - ETS: Allows grouping of different priorities and allocation of bandwidth to PFC groups
 - DCBX: Discovery and initialization protocol to discover resources connected to DCB-enabled network

Ethernet Fabric Details

Transparent Interconnection of Lots of Links (TRILL)



Layer 2 Multiple Paths

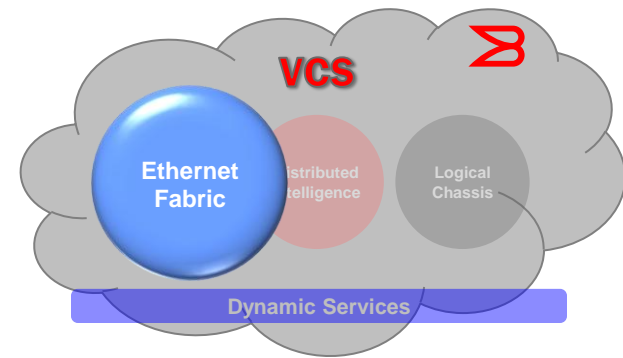


- Multi-path Layer 2 switching
 - All paths are active and traffic is distributed across all paths
 - Fully utilize all network bandwidth
- Establishes shortest paths through the Layer 2 network
- Uninterrupted response to link failures
- Backward-compatible and connects into existing infrastructures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

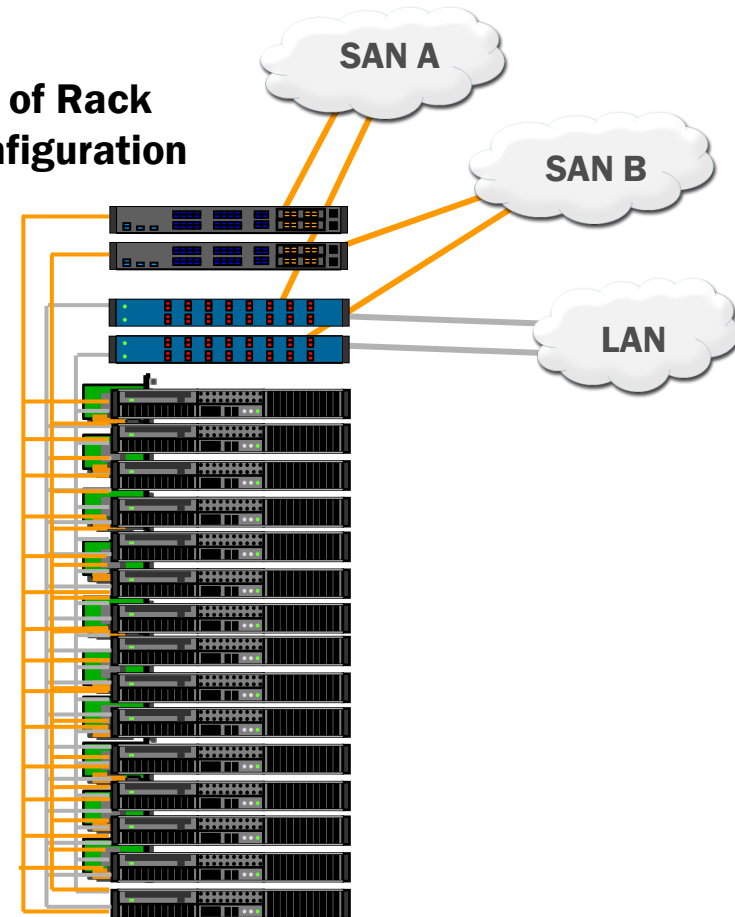
Ethernet Fabric Details

Convergence Ready

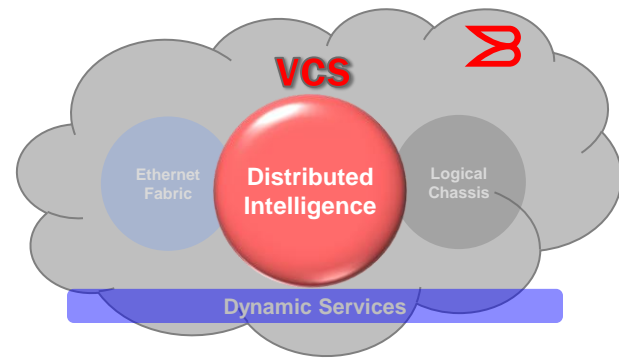
- Enables server I/O and end-to-end network convergence
Fewer cables
- Reduce number of server adapters, ports, cables and switches
Fewer adapters
- Reduce power consumption
- Increase speed and utilization of links
Fewer switches
- Simplify configurations and diagnostics



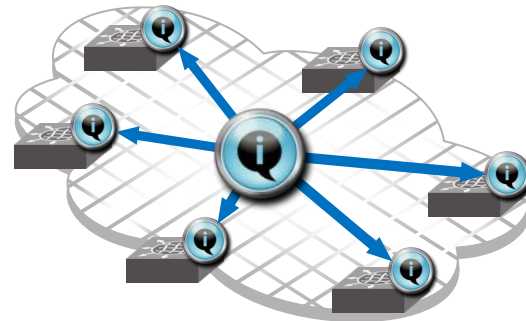
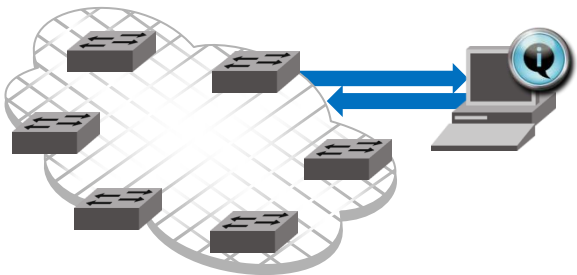
Top of Rack Configuration



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
- Optimized Virtual Access Layer
 - VEPA; frees host resources from switching and policy enforcement

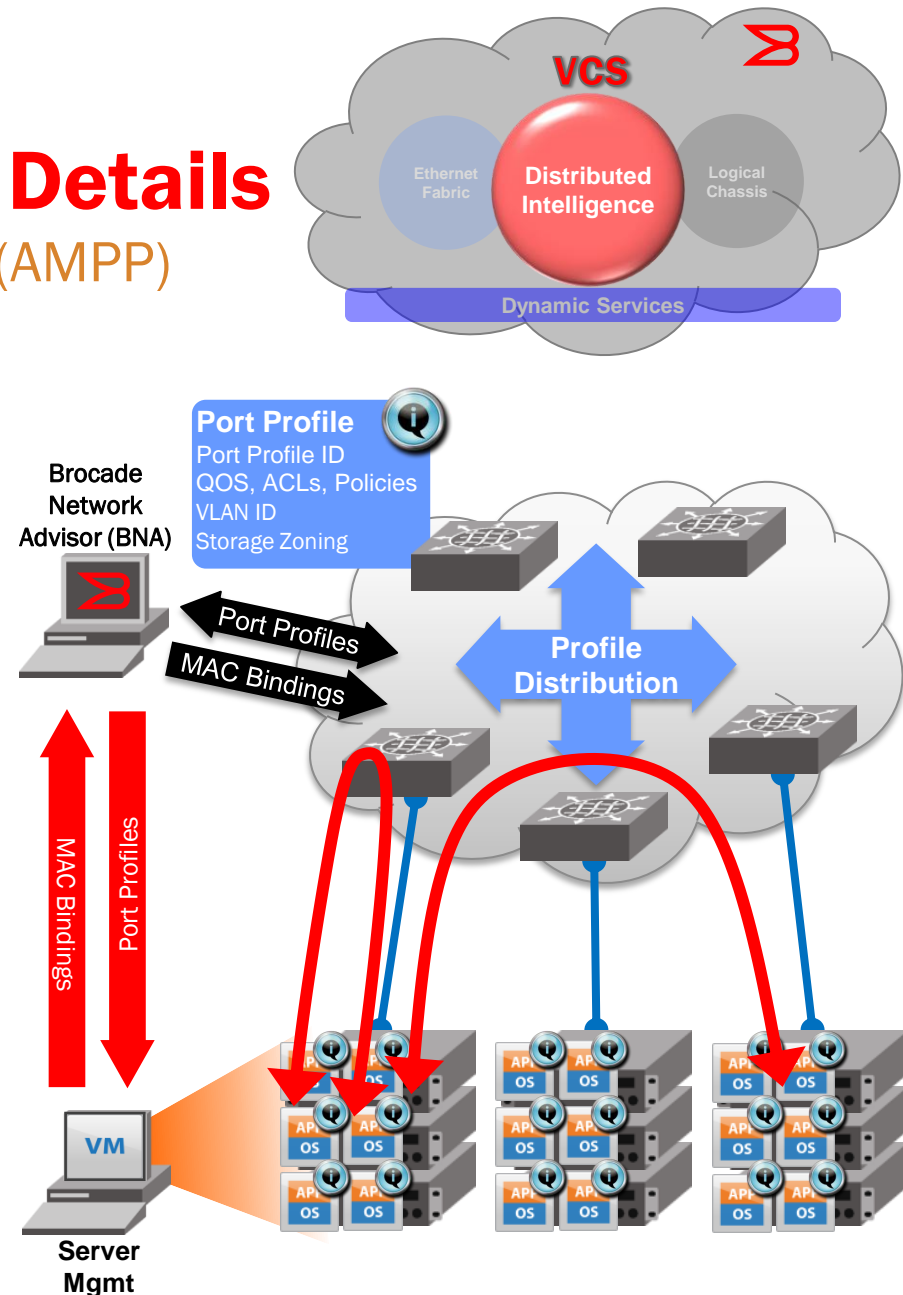


Distributed Intelligence Details

Automatic Migration of Port Profiles (AMPP)

❖ Allows VM to move with the network automatically reconfiguring

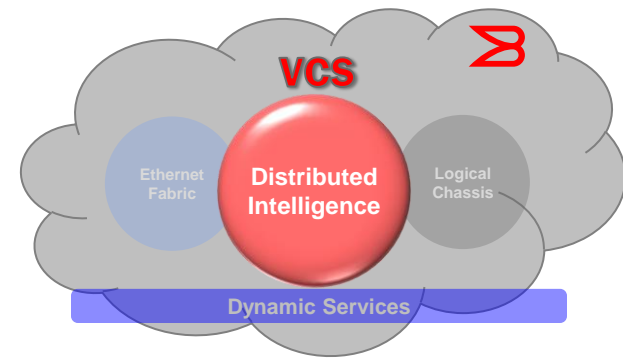
1. Port Profiles created, managed in fabric; distributed
2. Discovered by BNA; pushed to orchestration tools
3. Server admin binds VM MAC address to Port Profile ID
4. MAC address/Port Profile ID association pulled by BNA; sent to fabric
5. Intra- and inter- host switching and profile enforcement offloaded from physical servers



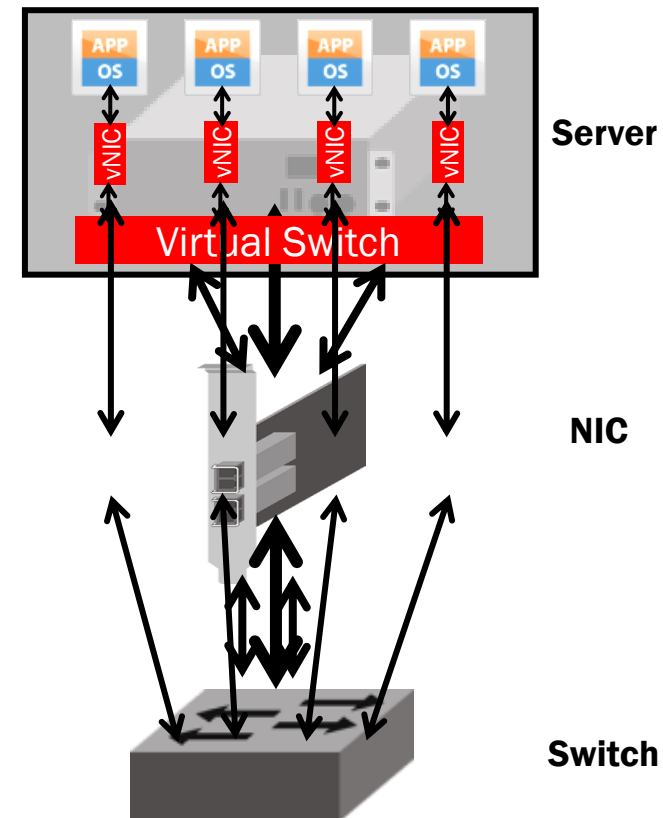
Distributed Intelligence Details

Optimized Virtual Access Layer

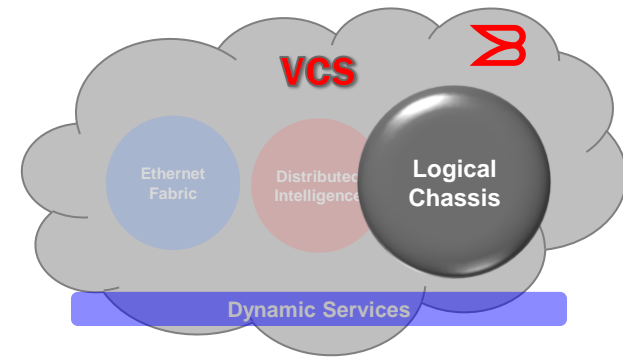
- Today, access to the network lives in the virtual hypervisor
 - Consumes valuable host resources
- Virtual switch is offloaded to the physical switch
 - Eliminates the software switch; the advantages of a distributed virtual switch plus Distributed Intelligence
 - Leverages Virtual Ethernet Port Aggregator (VEPA) technology
- Virtual NICs are offloaded to the physical NIC
 - Leverages Virtual Ethernet Bridging (VEB) technology
- Host resources are freed up for applications
 - Gives 5-20% of host resources back to applications
- VMs have direct I/O with the network
 - Network simplicity; common access across entire VCS; network is managed in the network



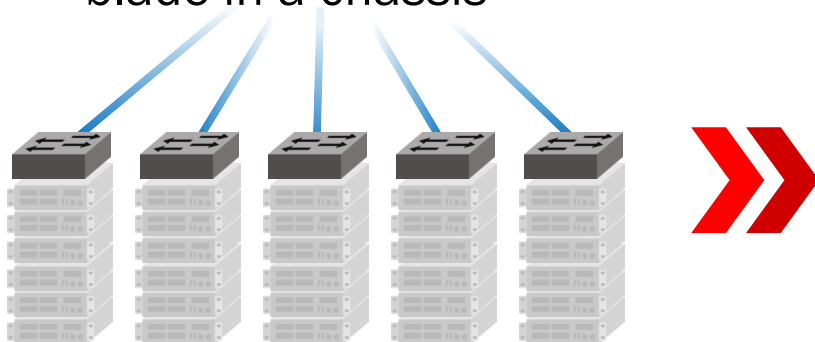
Virtual



Logical Chassis Details



- Fabric auto-configures
 - Once VCS is enabled, no configuration necessary
- Fabric behaves/managed as a single logical chassis
 - Aggregation (or Core) layer sees one switch
 - Fabric members act like a blade in a chassis

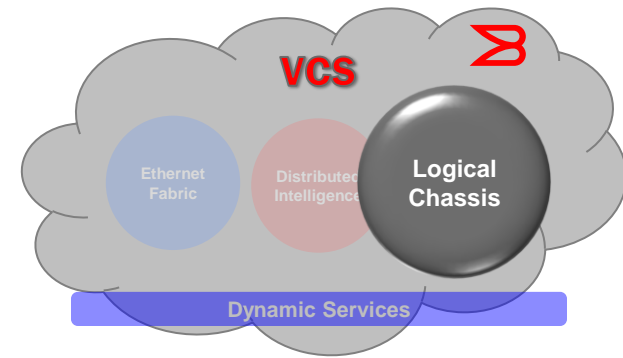
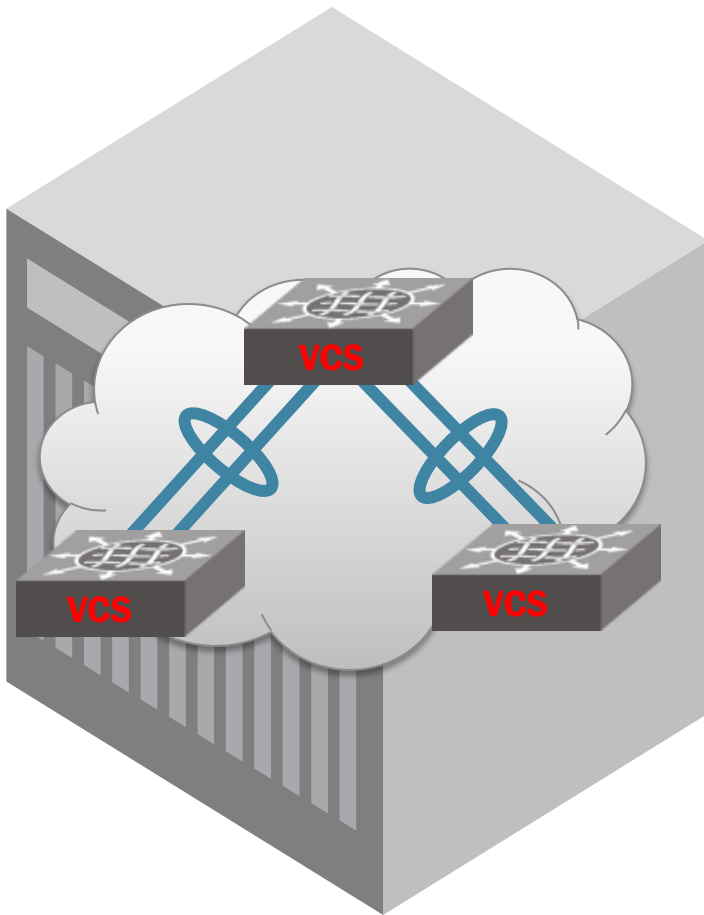


- Logically flattens and collapses network layers
 - Fabric is self-aggregating
 - Flexible fabric topologies
- Will scale to greater than 1000 device ports without added management



Logical Chassis Details


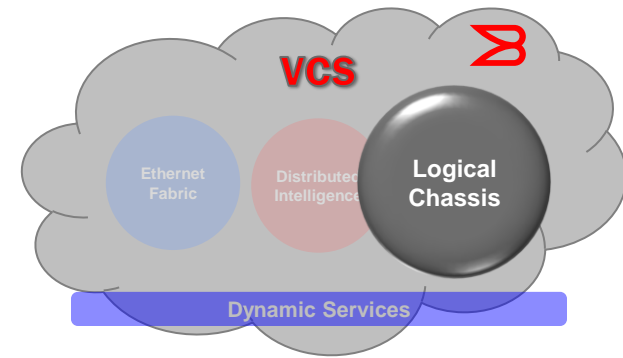
Auto-Configuration



- VCS simplified deployment, scalability, and management of the network
- Enable VCS on each switch
- Connect the switches
- Fabric automatically forms
 - Common configuration across all switches
 - vLAGs auto-configure
- Managed as a single logical chassis

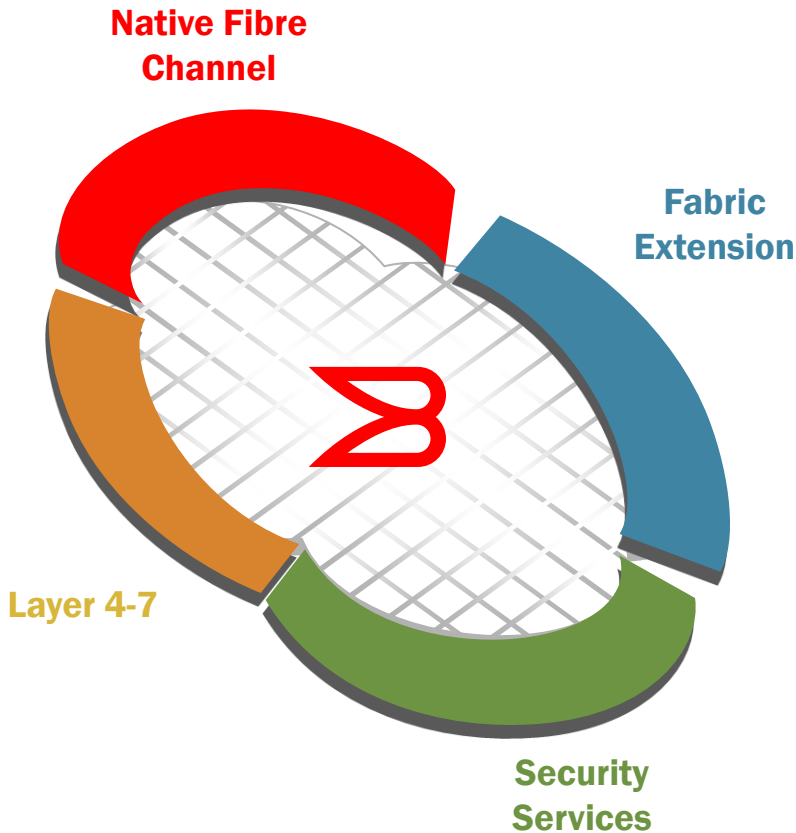
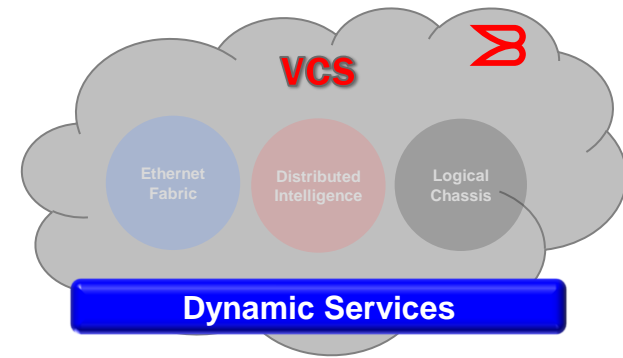
Logical Chassis Details

Single Logical Switch Behavior



- VCS behaves like a single industry-standard Ethernet switch
 - VCS fabric members are like blades in a modular chassis
- Standards-based and closed protocols used within the fabric
 - FSPF, TRILL, Fabric Services, etc.
- Industry-standard protocols used to communicate outside the fabric
 - RSTP, LACP, 802.1x, sFLOW, etc.

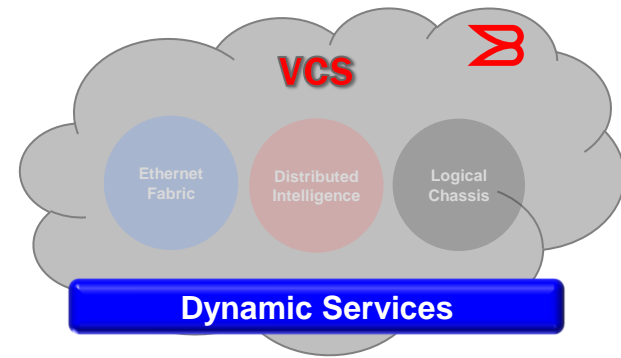
Dynamic Services Details



- Add services into the VCS fabric
 - Extends the capabilities of VCS
 - Fabric extension, native Fibre Channel, security services, layer 4-7, etc.
- Purpose-designed hardware
 - Switches with unique functionality can be added to the VCS fabric
 - Like service modules in a chassis
 - Functionality available to the entire VCS fabric

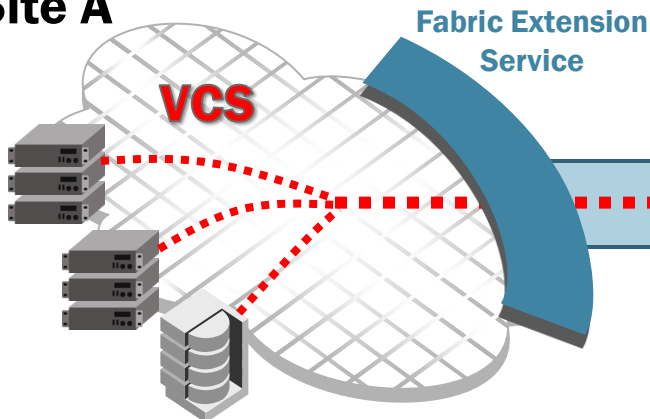
Dynamic Services Details

Data Center to Data Center Connectivity



- Dynamic Service to connect Data Centers
 - Extend the layer 2 domain over distance
 - Maintains fabric separation while extending VCS services to secondary site (e.g. discovery, distributed configuration, AMPP)
- VCS Fabric Extension capabilities
 - Delivers high performance accelerated connectivity with full line rate compression
 - Secures data in-flight with full line rate encryption
 - Load balances throughput and provides full failover across multiple connections

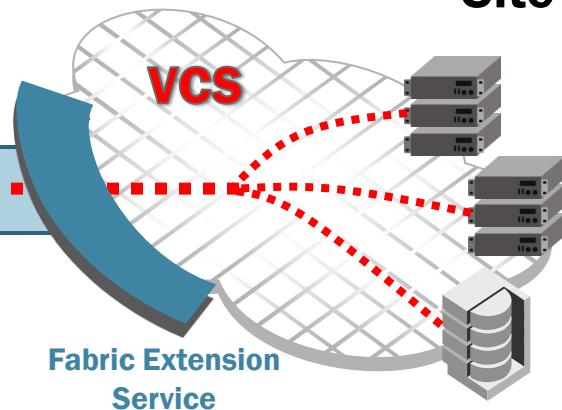
Site A



Public Routed Network

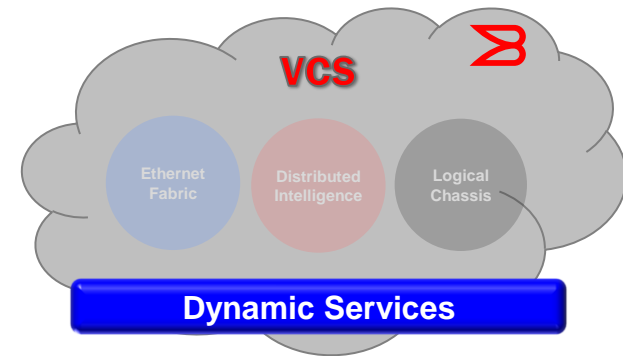
Encryption, Compression,
Multicasting

Site B

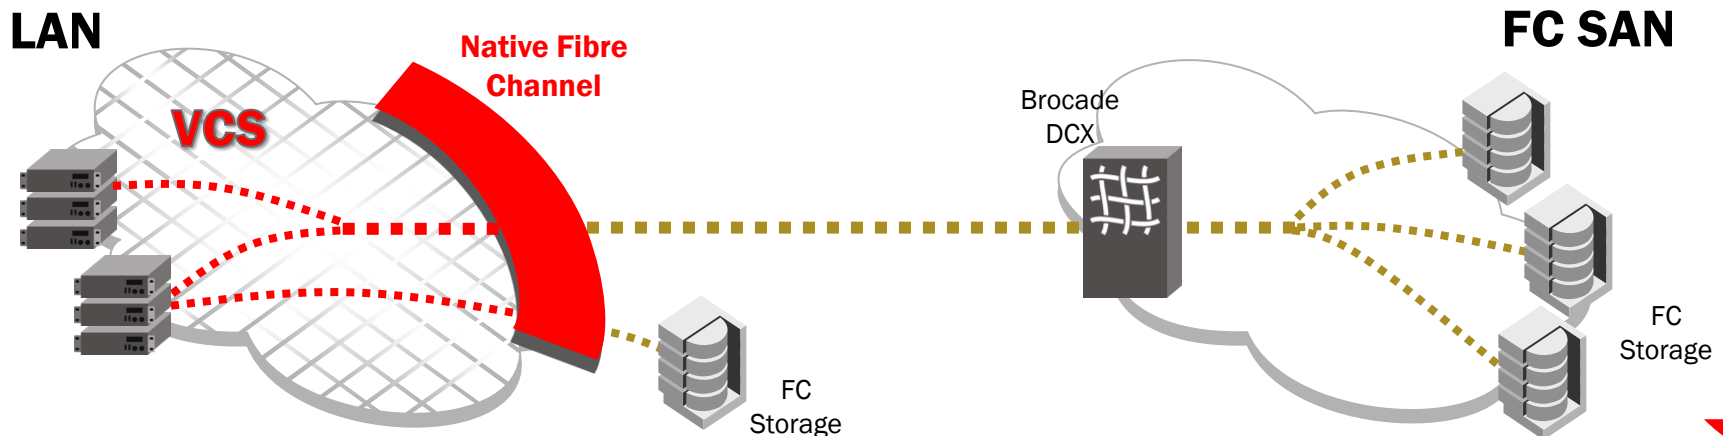


Dynamic Services Details

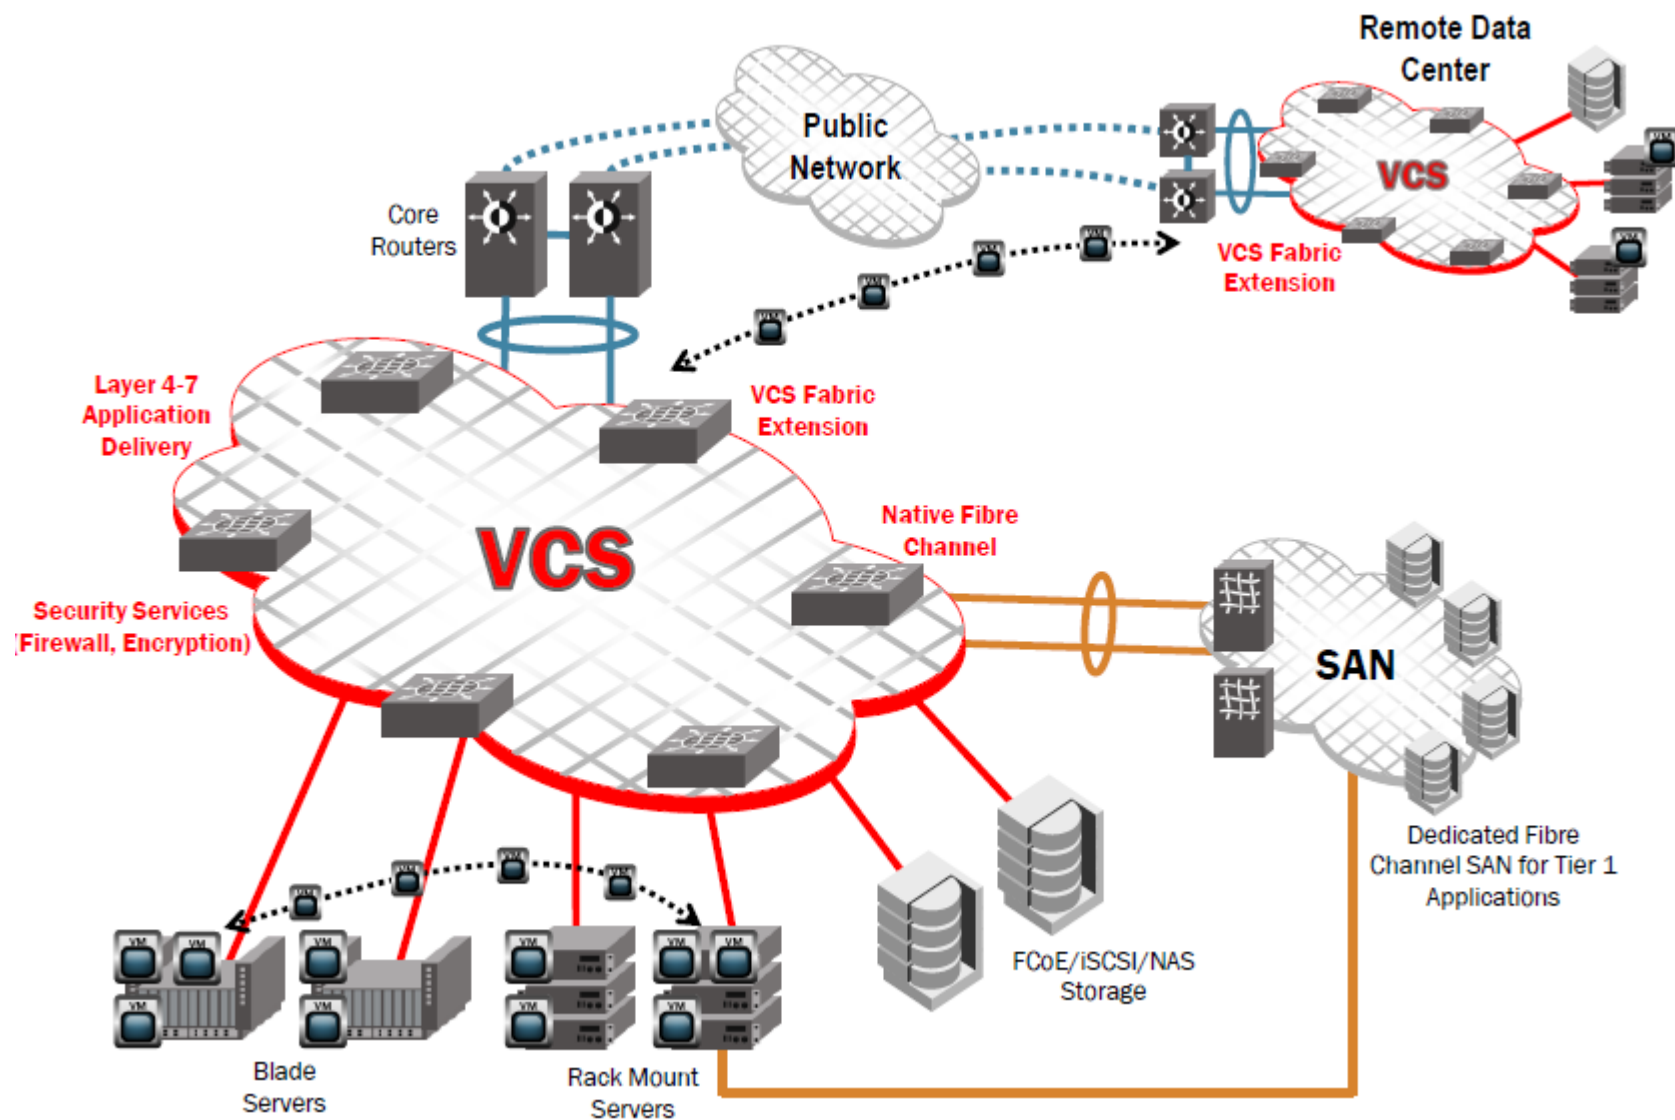
Native Fibre Channel Connectivity

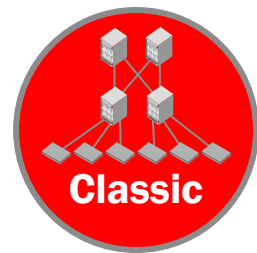


- Provide VCS Ethernet Fabric with native connectivity to FC storage
 - Connect FC storage locally
 - Leverage new or existing Fibre Channel SAN resources
- VCS Native Fibre Channel Capabilities
 - Adds Brocade's Fibre Channel functionality into the VCS fabric
 - 8 Gbps, 16 Gbps FC, frame-level ISL Trunking, Virtual Channels with QoS, etc.



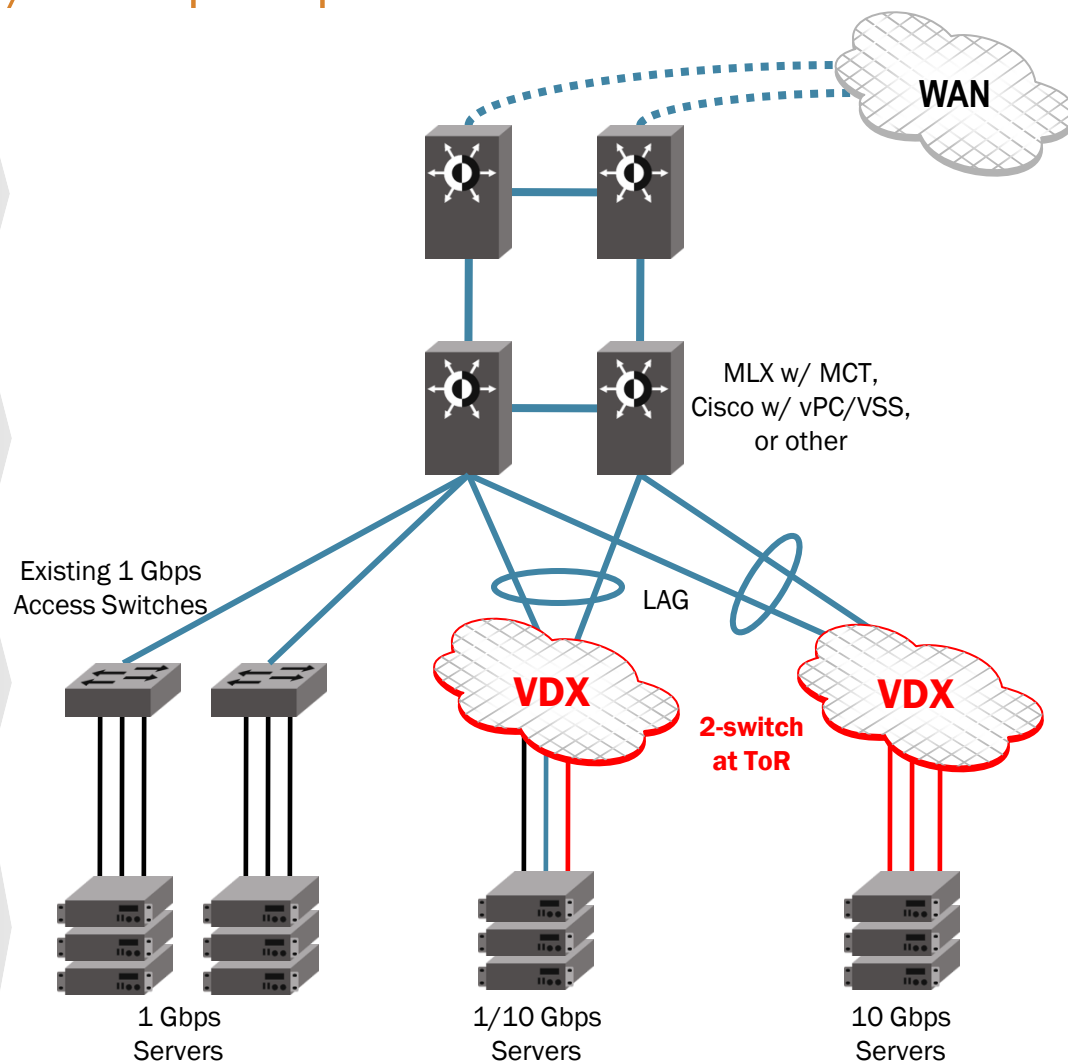
VCS architecture





VCS Use Case #1

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



Preserves existing architecture

Leverages existing core/agg

Co-exists 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

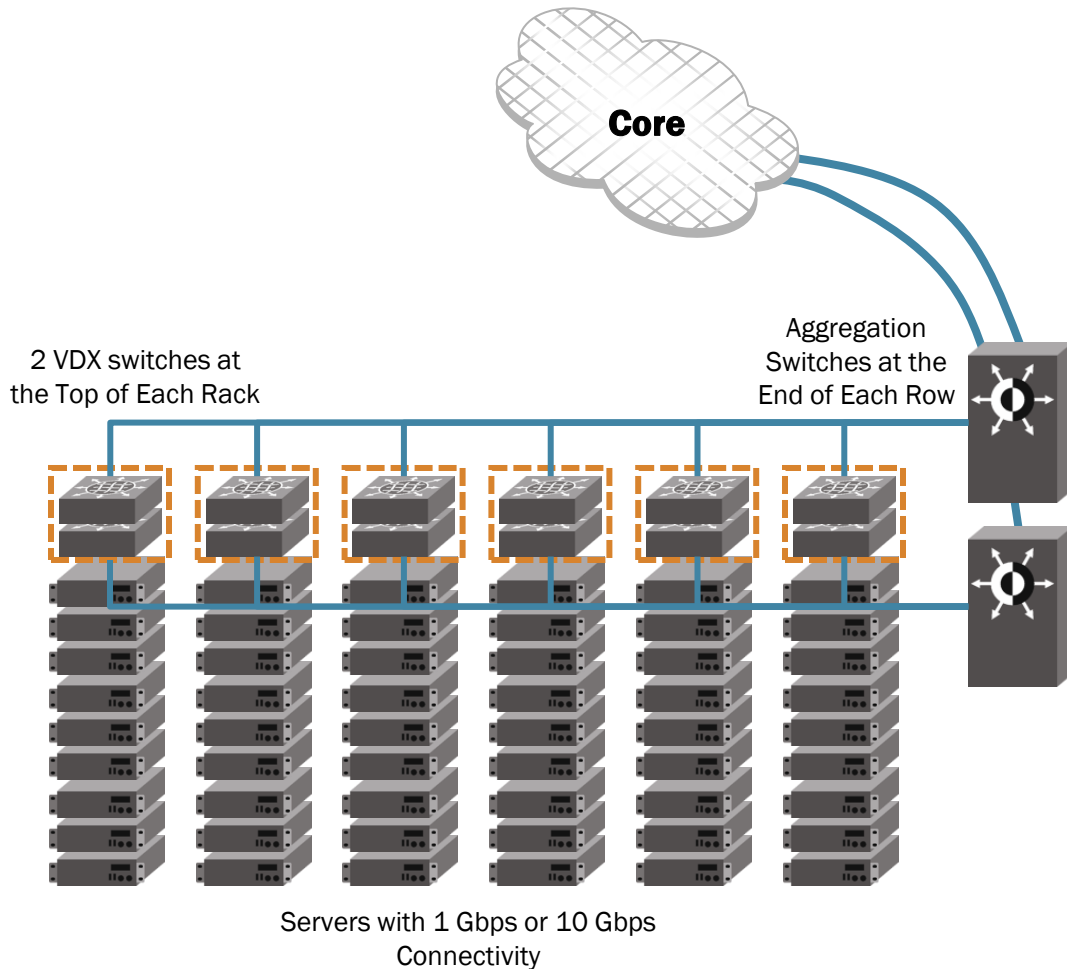
High-density access with flexible subscription ratios

Supports up to 36 servers per rack with 4:1 subscription



VCS Use Case #1

1/10 Gbps Top-of-Rack Access – Layout



Preserves existing network architecture

Leverage VCS technology in stages

2 VDX switches in each server rack

Managed as a single switch

1 Gbps and 10 Gbps connectivity

Highly available; active/active

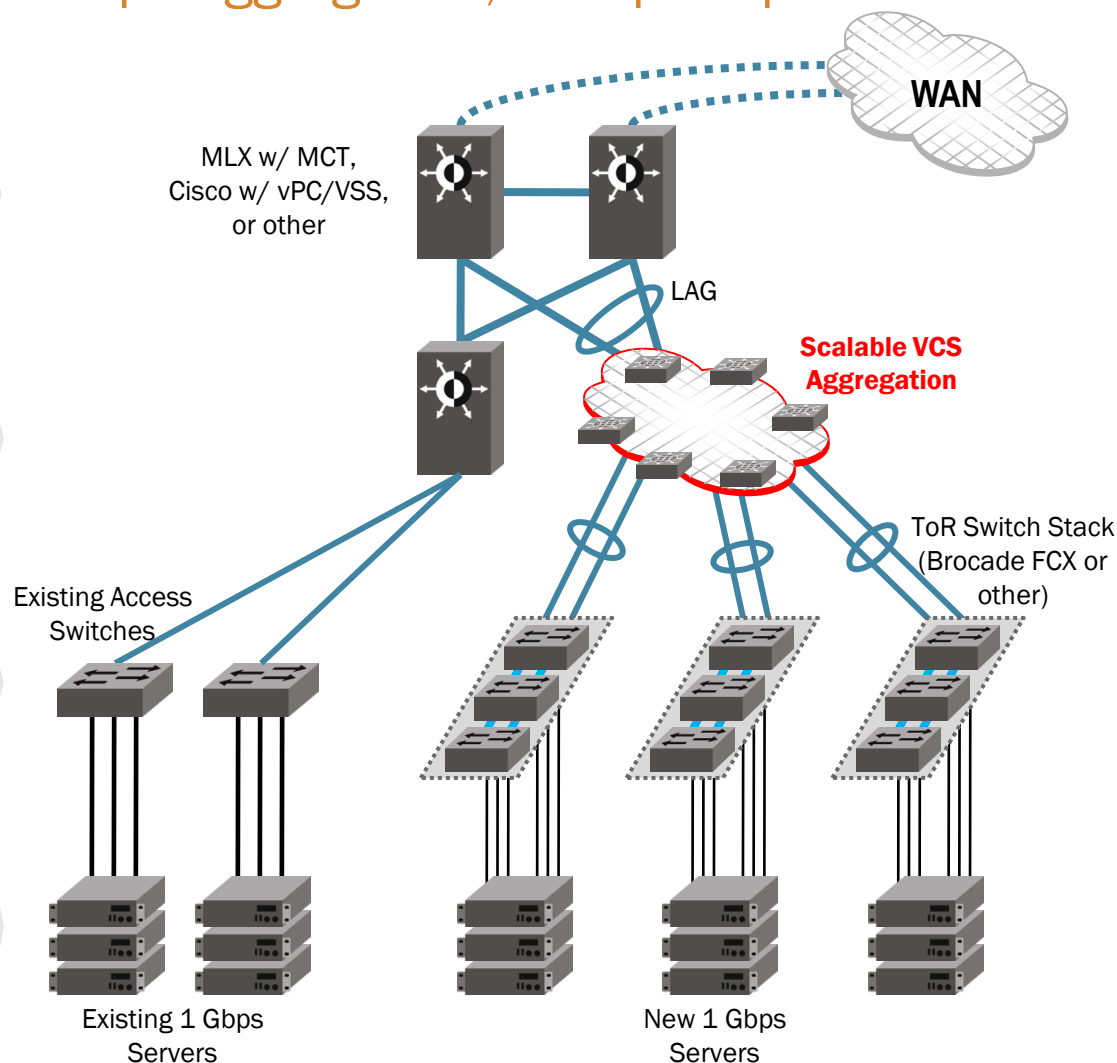
High performance connectivity to End-of-Row Aggregation

One virtual LAG to core for simplified management and rapid failover



VCS Use Case #3

10 Gbps Aggregation; 1 Gbps Top-of-Rack Access – Architecture



Low cost, highly flexible logical chassis at aggregation layer

Building block scalability

Per port price of a ToR switch

Availability, reliability, manageability of a chassis

Flexible subscription ratios

Ideal aggregator for 1 Gbps ToR switches

Supports 900 servers in 25 racks, assuming 4 NICs per server

Optimized multi-path network

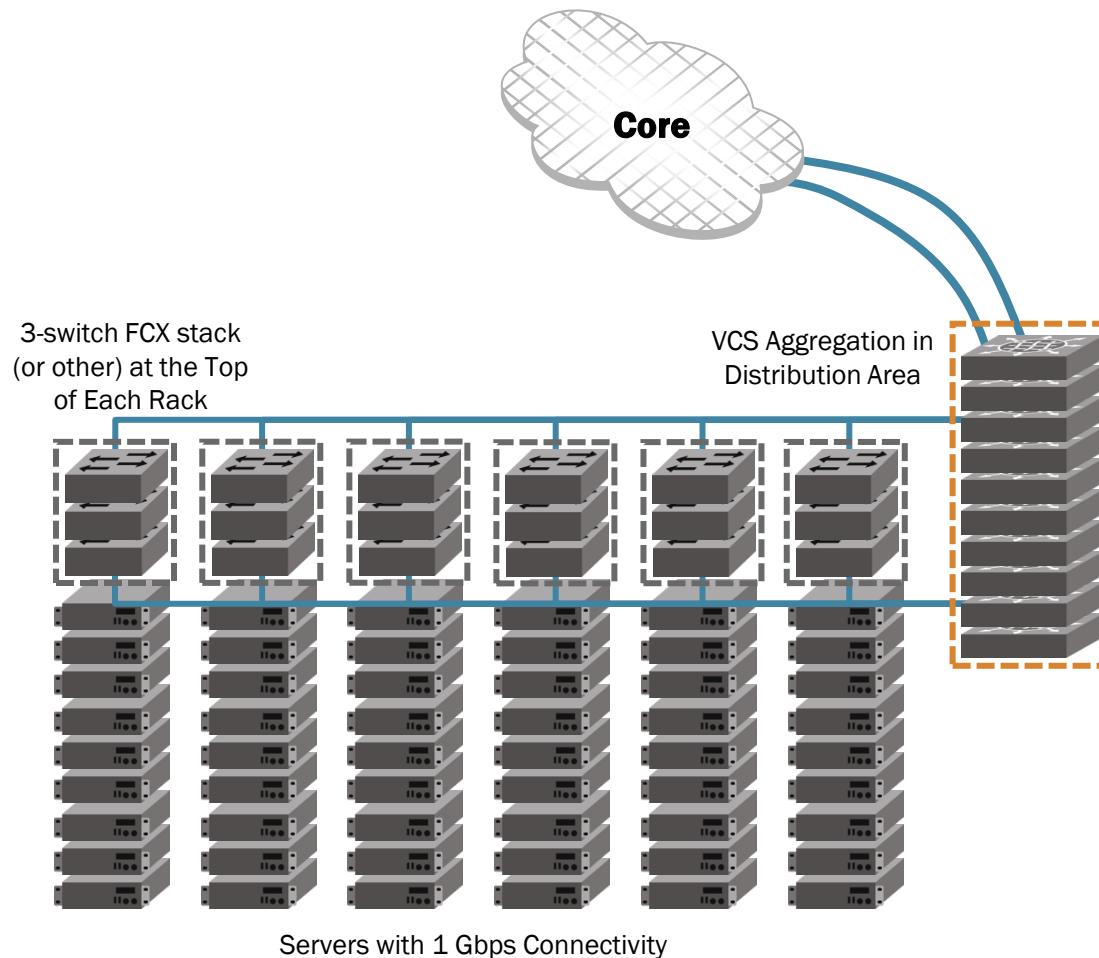
No single point failure

STP not necessary



VCS Use Case #3

10 Gbps Aggregation; 1 Gbps Top-of-Rack Access – Layout



3-switch GbE stack in each rack

Managed as a single switch

Redundancy throughout
network, without STP

High density 10 Gbps LAG to
VCS aggregation

Logical Chassis Aggregation
Router in Distribution Area

Build out aggregation as needed

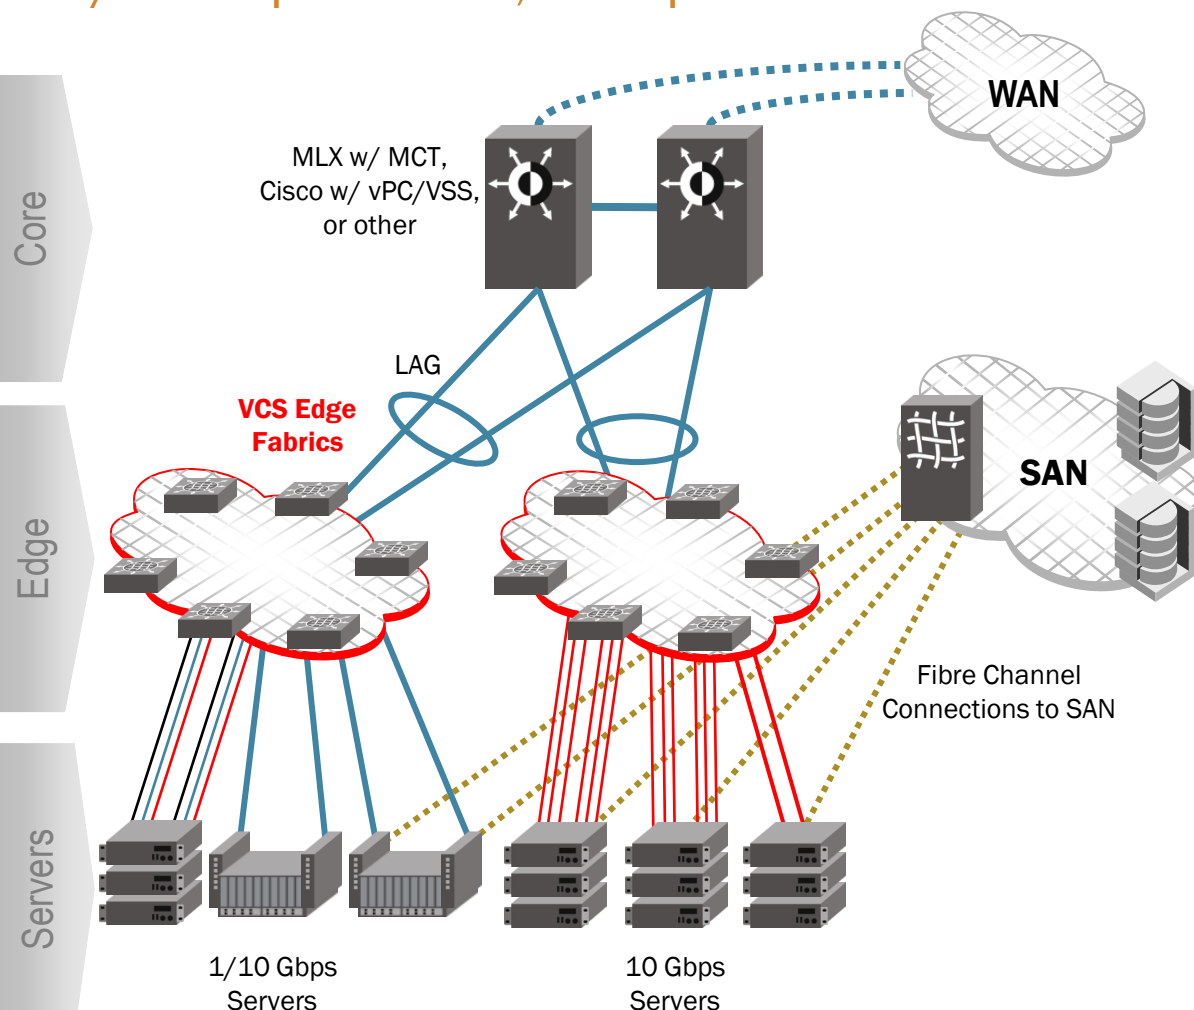
Supports 30 racks of servers

High performance, resilient
connection to Core

One LAG for simplified
management and rapid failover

VCS Use Case #4

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 multi-path network

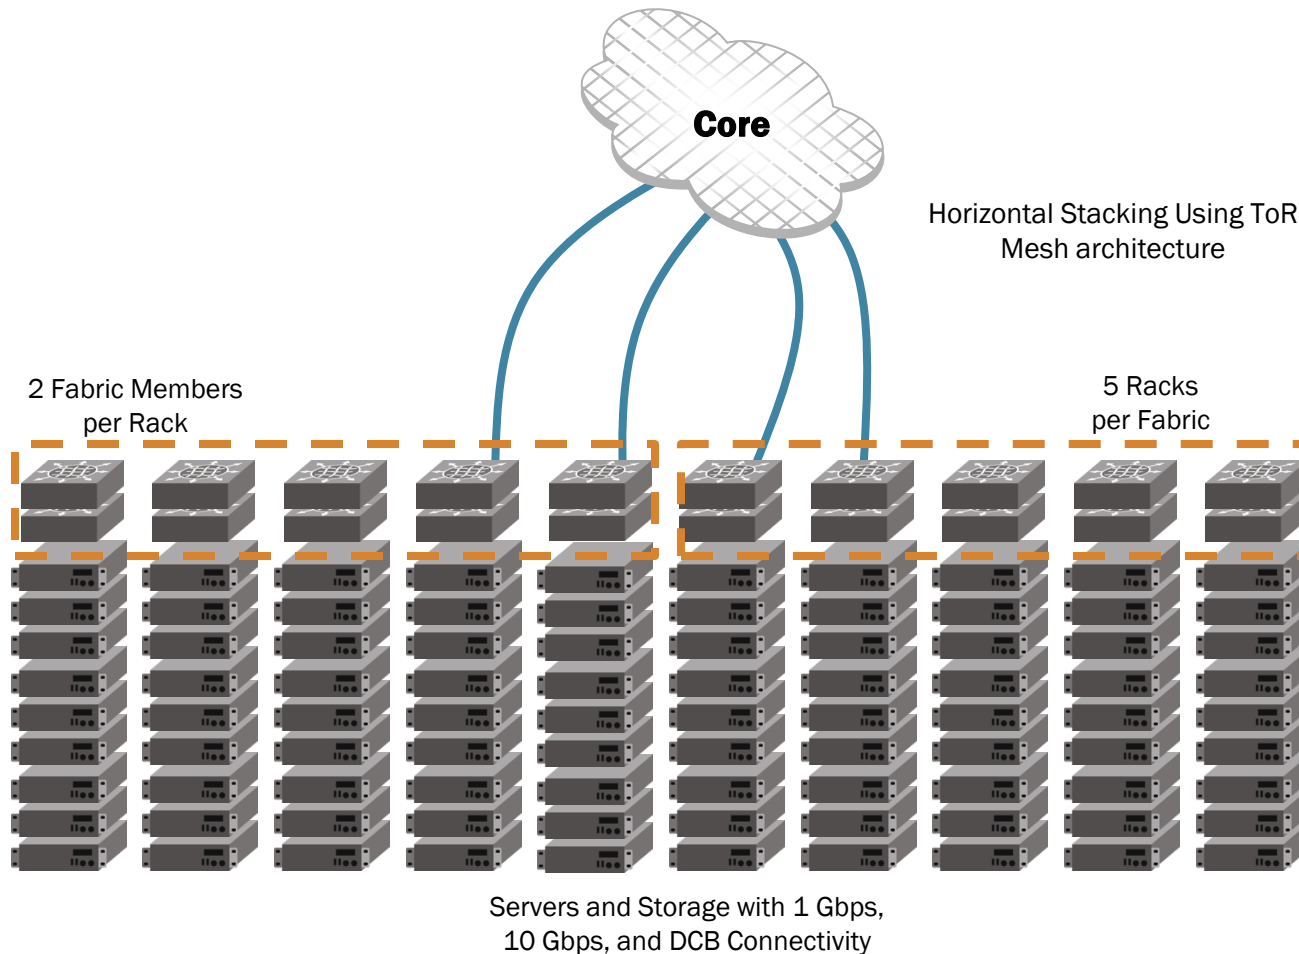
All paths are active

No single point failure

STP not necessary

VCS Use Case #4a

1/10 Gbps Access; Collapsed Network – Layout – ToR Mesh



2 VCS fabric members in each rack

Dual connectivity into fabric for each server/storage array

Low cost Twinax cabling in rack

Switches connected to peers across racks

Low cost active optical cabling between top-of-rack switches

1 VCS fabric per 5 racks of servers (assuming 36 servers per rack)

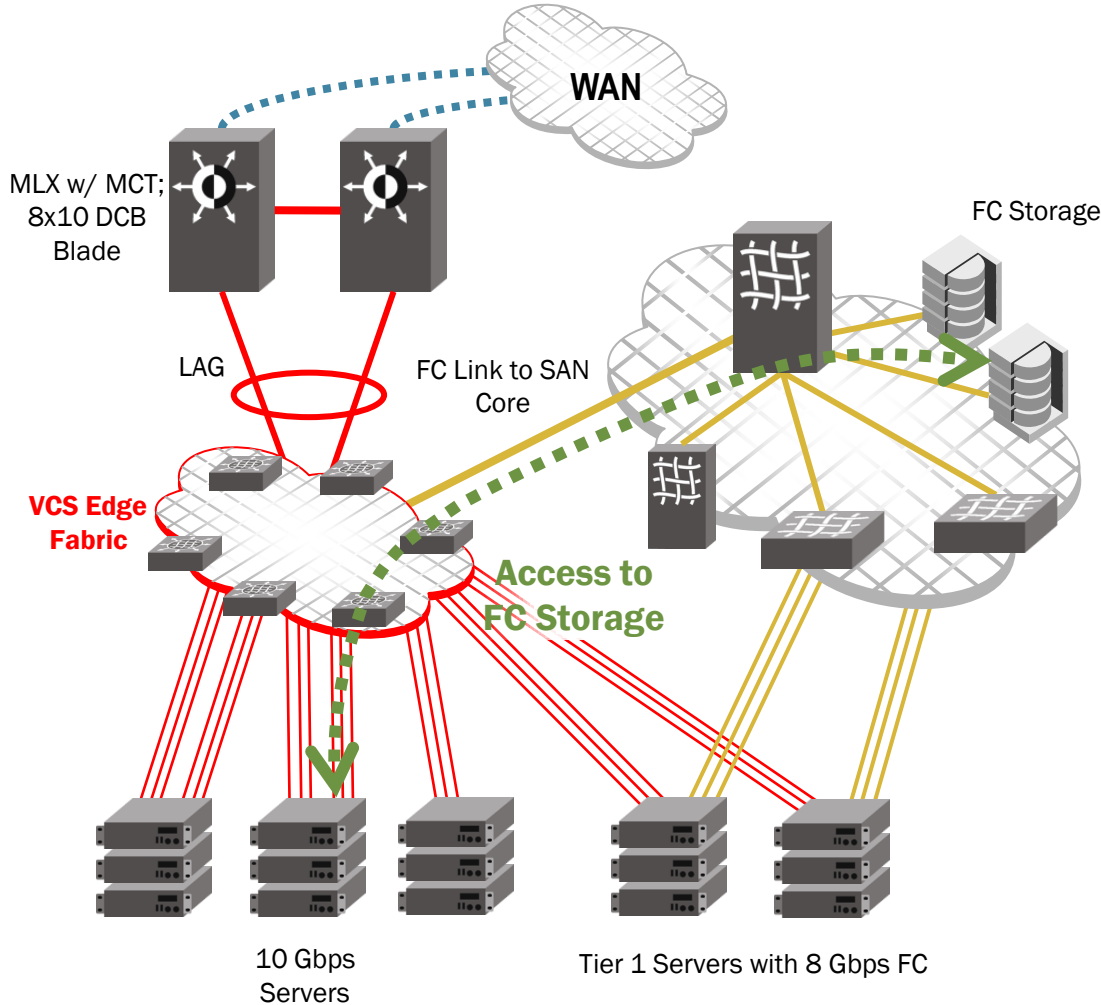
Fiber optic cabling only used for connectivity from edge VCS to core

Single virtual LAG per fabric

Reduced management and maximum resiliency

VCS Use Case #6

1/10 Gbps Access; Convergence + FC SAN – Architecture



Leverage existing resources

Connect Ethernet fabrics into Fibre Channel SAN – new servers have access to existing storage

***Enabled through future VDX switch with native FC ports

Maximum storage flexibility

Fibre Channel, FCoE, iSCSI, NAS

Deploy the right storage technology without isolating it

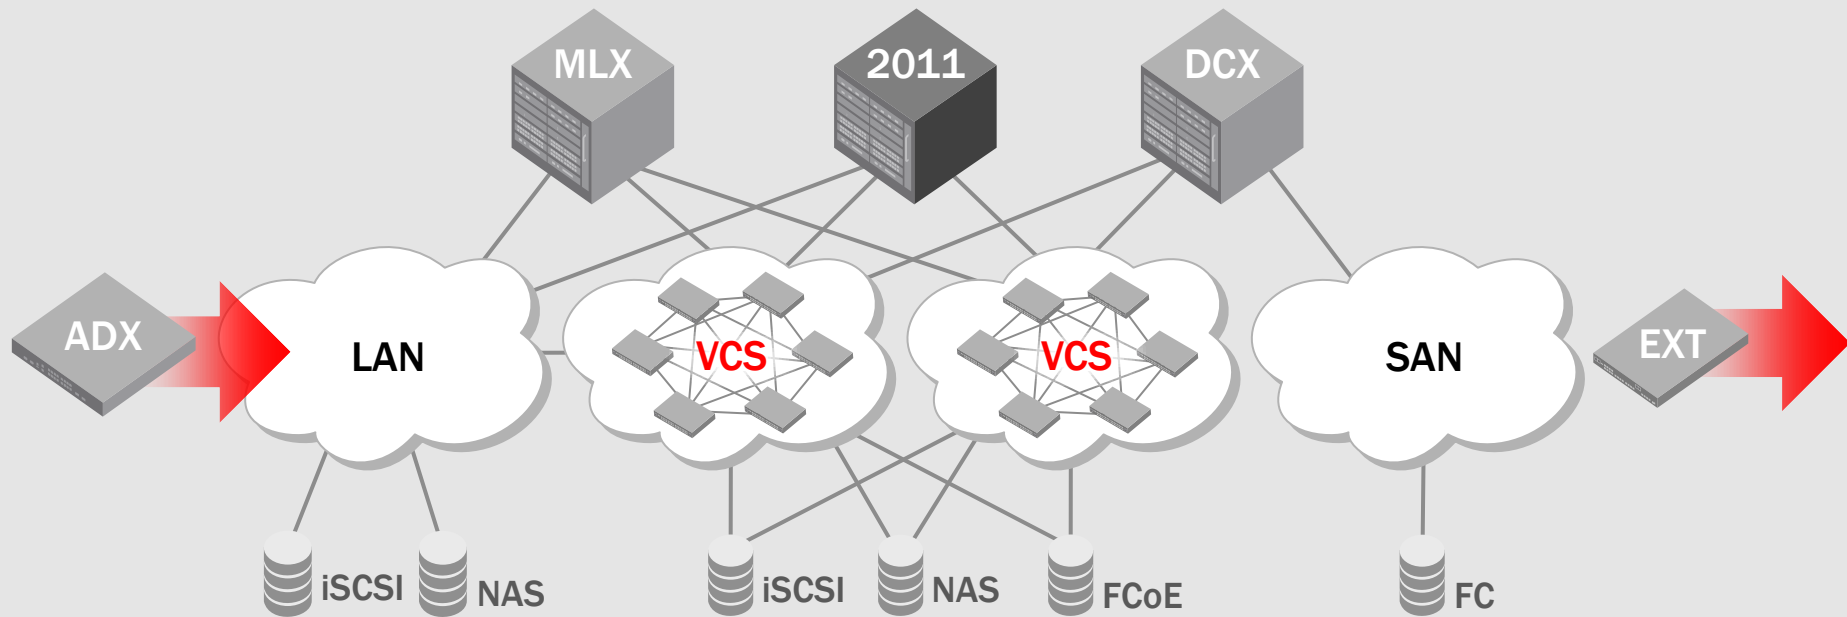
Optimal performance, availability

No single point failure

Frame-level, hardware-based trunking between fabrics

BROCADE SOLUTIONS

Scaling and Extending the Data Center



Brocade VDX 6720



Brocade VDX 6720 Data Center Switches

Product Highlights

- **Built for the Virtualized 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 with Ports On Demand
 - Non-blocking, cut-through architecture, wire-speed
 - 600 ns port-to-port latency; 1.8 us across port groups
- **Environmental Flexibility**
 - 10 Gb and 1 Gb supported on every port
 - Direct-attached copper, active 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



Data Center Access

Some Key Concepts...

ASICs, features, multihop...

- The ASIC inside of VDX 6720-24 and VDX 6720-60 is eAnvil2
 - 24 x 10/1GE ports & 8 x 8G FC Bridging ports
 - Supports TRILL, DCB, L2 and L3 features
 - **Cut-thru 600ns latency** switch chip
- Ethernet Fabric
 - vLAG (Similar to MLX MCT) – Sometime referred as internal VCS LAGs
 - AMPP
 - Hypervisor agnostic - VMWare and Hyper-V tested (but should not be any issue with others)
 - Multi-hop internal FCoE
 - Callisto-F with Native FC ports in roadmap

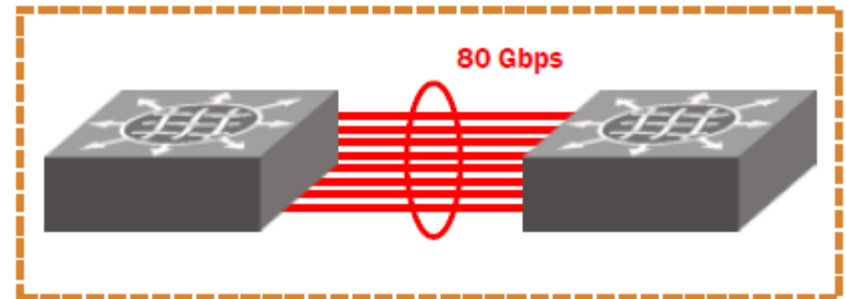


ISL Trunking

No license

Brocade ISL Trunking provides high link utilization, ease-of-use

- Frame-level, hardware-based trunking
- Frames are evenly distributed across links in the trunk
- Built into Brocade fabric switching ASIC
- ISL Trunks automatically form
 - Ports must belong to the same port group in the switch
- Once both switches are in VCS mode, multiple ISLs automatically form a trunk
- No configuration necessary



iSCSI DCB

Advantages

Provide deterministic delivery of iSCSI traffic, Maximize iSCSI throughput and
Minimize TCP re-transmissions by eliminating congestive packet loss
Use DCB Ethernet Enhancements (PFC)(ETS)

Use DCBX

- Distribute DCB configuration to iSCSI devices
- PFC and ETS –existing
- iSCSI Priority –New TLV required

Switch advertises priority to be used for iSCSI

- Priority must be PFC enabled
- Advertisement of configuration only
- Switch will not verify or enforce iSCSI device compliance
- Requires that device supports the new TLV
- DCBX 1.01 compliant TLV
- Application Protocol ID = 3260
- Priority map indicates iSCSI priority



Licenses

VDX 6720-24 (min 16 ports)

- VCS 2-node License (Default)
- VCS Multi-node License
- FCoE License
- **8 port POD(Ports on Demand) License**

VDX 6720-60 (min 40 ports)

- VCS 2-node License (Default)
- VCS Multi-node License
- FCoE License
- **10 port POD(Ports on Demand) License**

Sentinel implementation, allowing XML files to load licenses, future Temporal lic support



Brocade MLXe



Brocade MLXe Series

Product Highlights



✱ Ideal for...

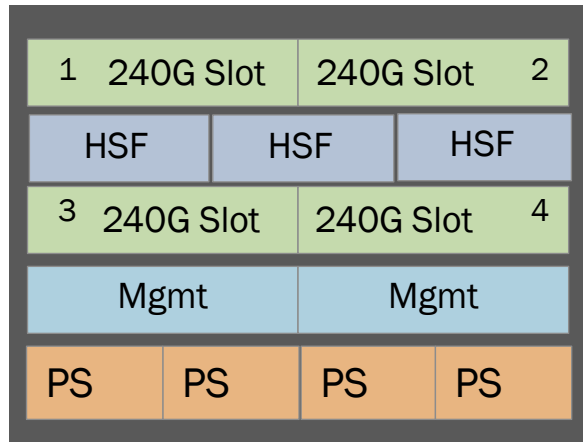
- ✓ High performance computing
- ✓ Dense data centers
- ✓ Very large-enterprise core

- ✱ 15.36 Tbps fully Distributed Architecture
- ✱ Lossless fabric
- ✱ Wire-speed routing 7.6 Gpps IPv4/IPv6/MPLS
- ✱ DCB (Data Centre Bridging) ready
- ✱ Industry-leading 100G port density
- ✱ Carrier-grade QoS
- ✱ NEBS L3 compliant
- ✱ Rear exhaust on all models
- ✱ High Availability design:
 - ✓ Redundant management modules
 - ✓ Redundant switch fabrics
 - ✓ Redundant power supplies & fans
 - ✓ Hitless failover
 - ✓ Hitless software upgrades
 - ✓ Node/link failover of <200msec

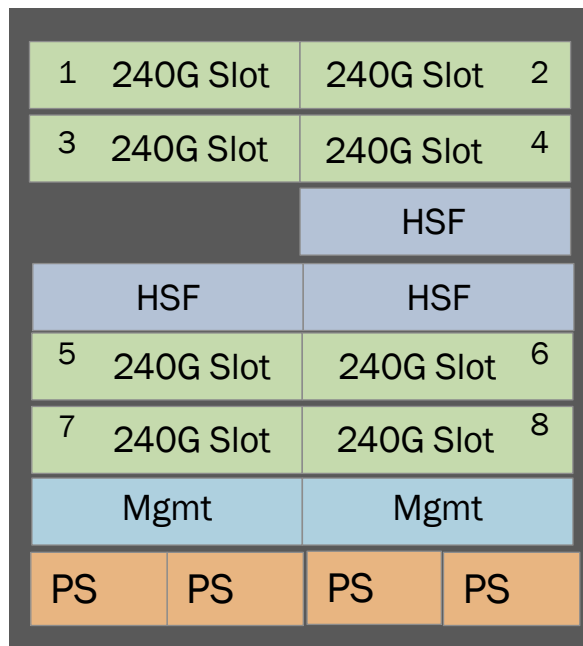


Brocade MLXe

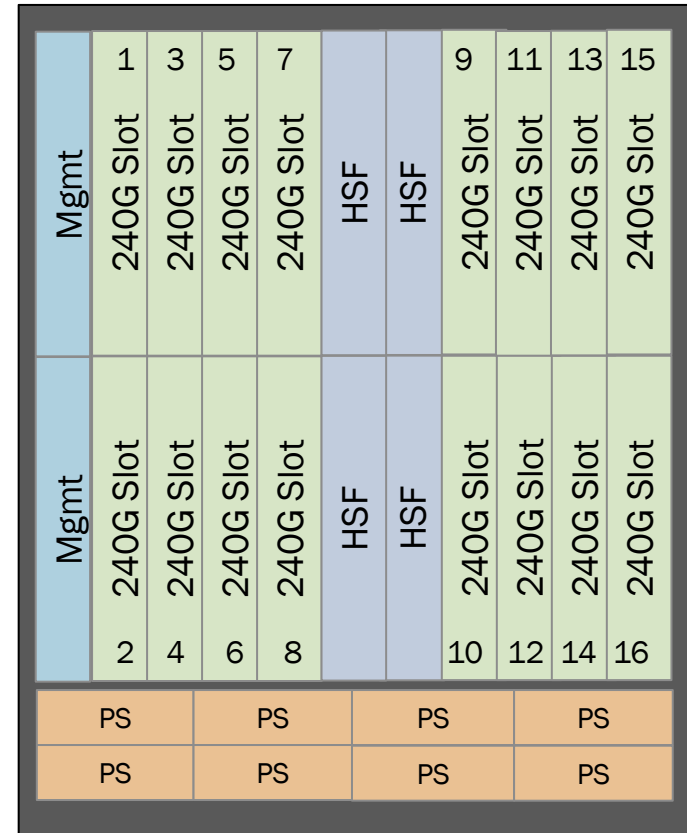
MLXe Layout



MLXe-4



MLXe-8

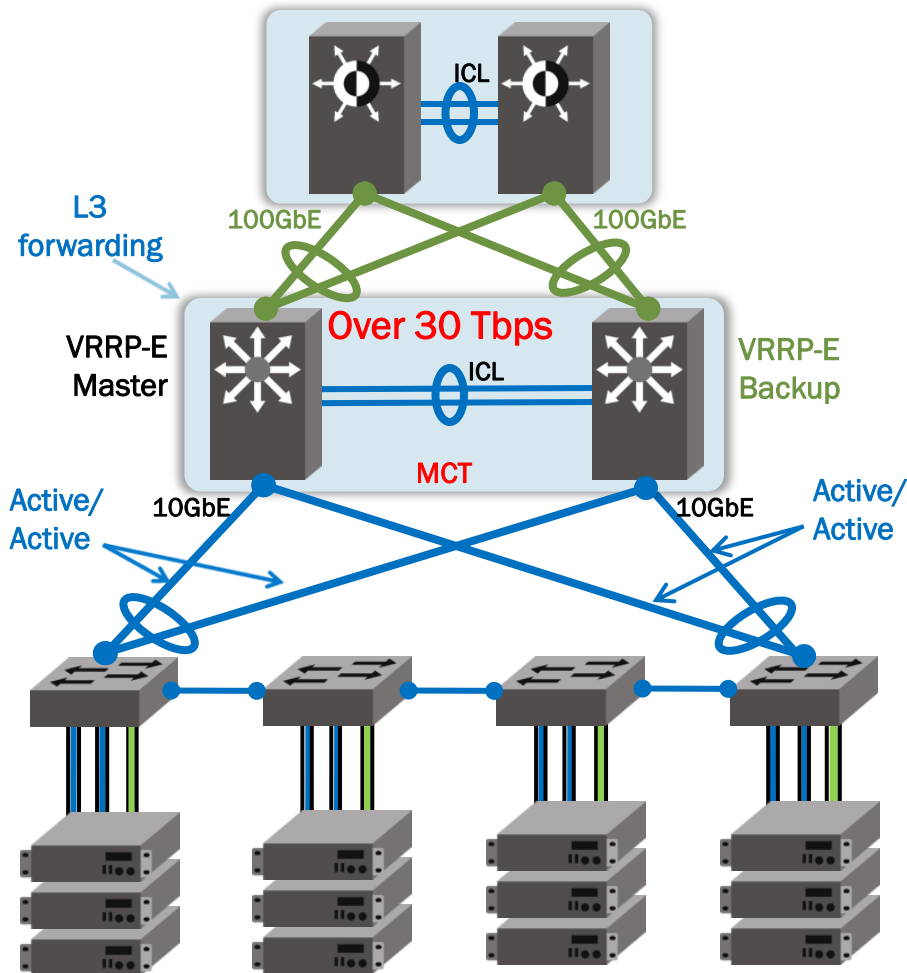


MLXe-16



Brocade MLXe Series

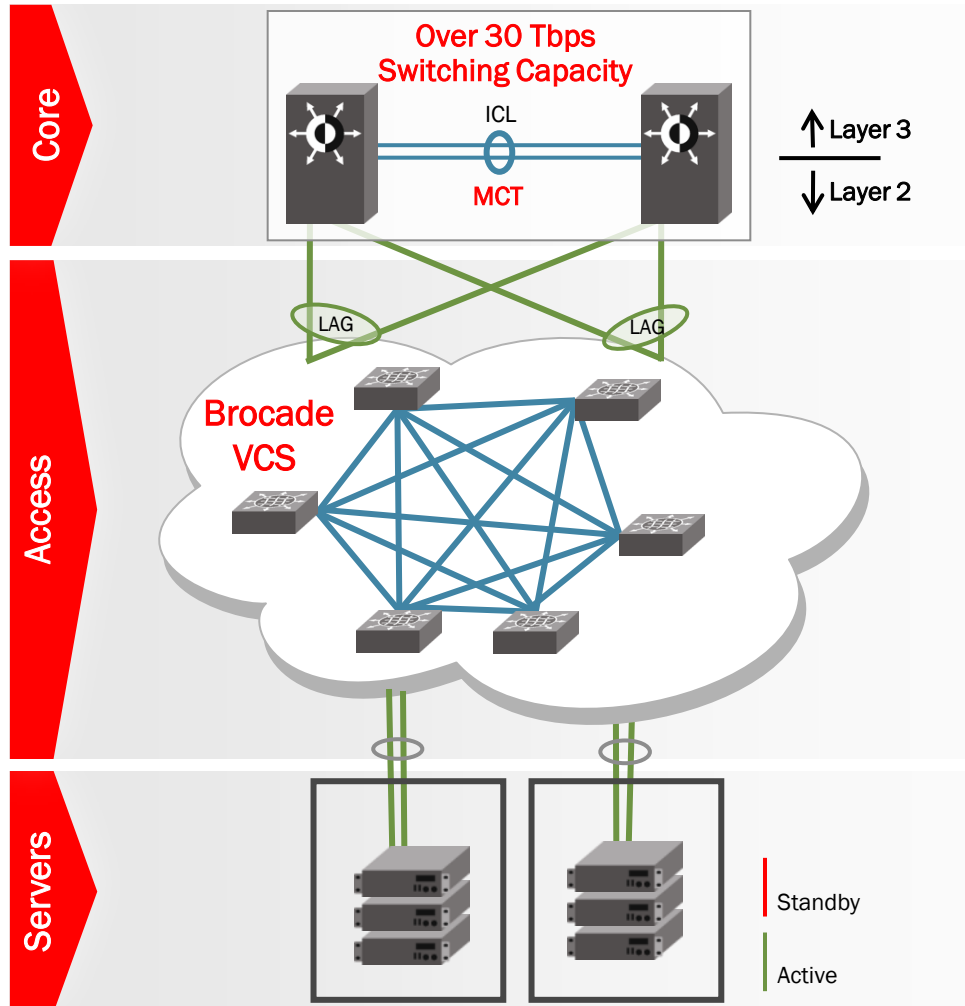
Simplified architecture & Operational Efficiency



- All links active (vs. active/passive) and forwarding Layer 2/3 traffic
- Highest resiliency High resiliency: < 200ms link or node failover
- Over 30 Tbps switching capacity in Multi-Chassis for investment protection
- 32x100G wire-speed ports for demanding networks

Brocade MLXe Series

Massively scalable Virtual Cluster Switching (VCS) Fabric

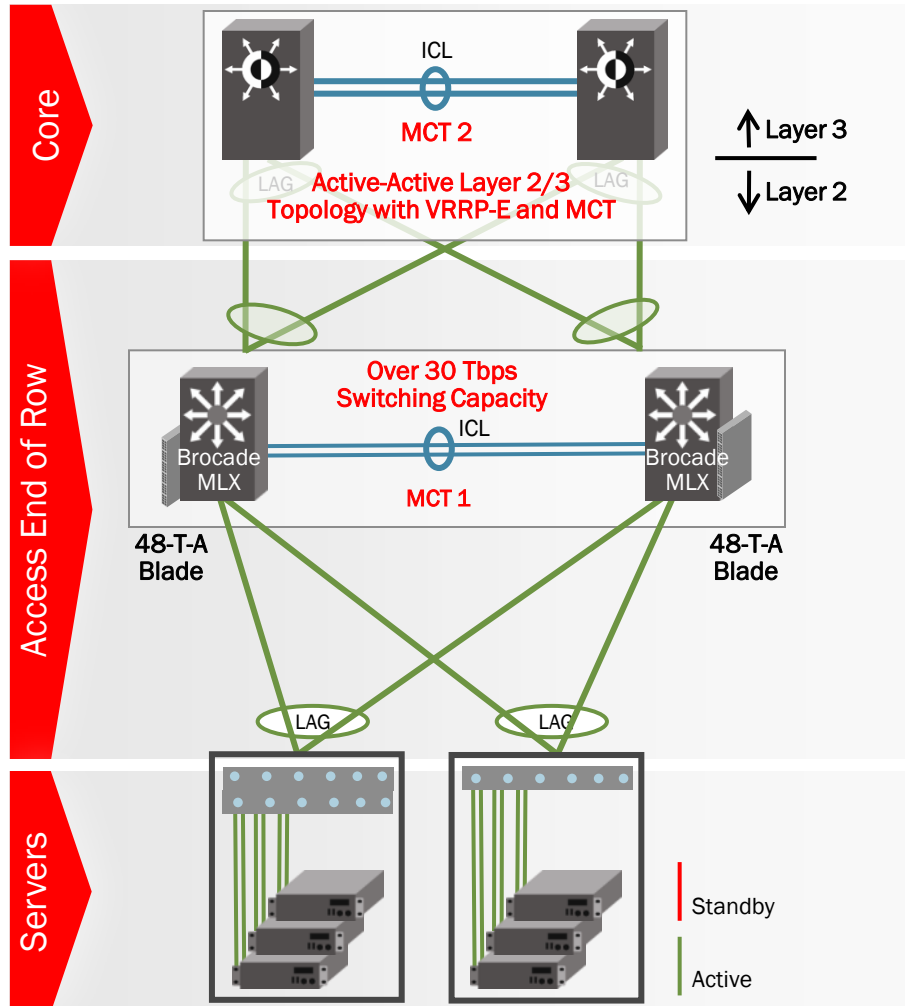


- **Easy to configure and manage**
 - Simplified architecture and scalable multipath Layer 2 domain
 - No Spanning Tree Protocol (STP) within Brocade VCS cloud or core
 - VM-aware infrastructure
- **Five times improvement**
 - Multipathing Layer 2 within Brocade VCS cloud and core
- **Two times improvement**
 - Full use of all network links
- **Investment protection**
 - Future-proof and convergence-ready DCB/FCoE architecture
 - Pay-as-you-grow model for server access



Brocade MLXe Series

Collapsed VToR architecture with MLX-48T



- **Easy to configure and manage**

- Fewer network elements to manage
- **No STP** from server to core
- Single pane NMS

- **Five times improvement**

- Deterministic sub-200 ms link and node failover from server to core
- Collapsed access/aggregation for lower device failure and higher MTBF

- **Two times improvement**

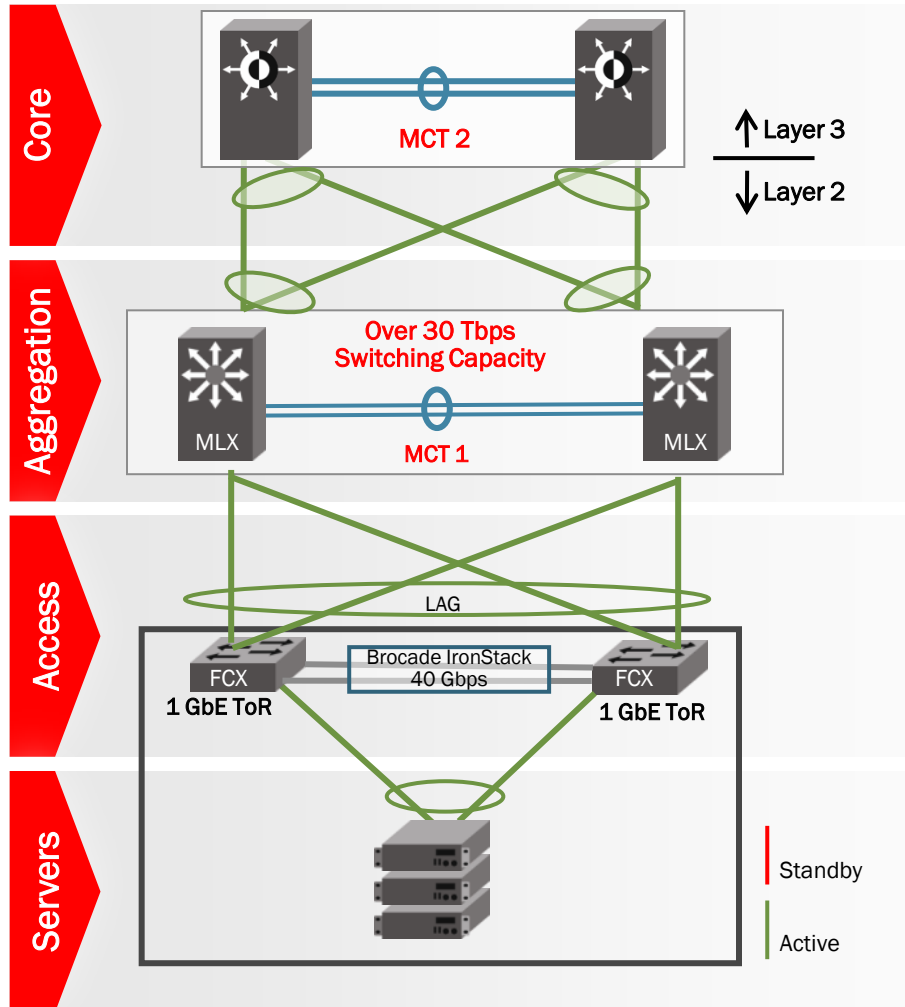
- Active-active L2/3 topology with VRRP-E and MCT from server to core
- Maximum VM mobility in IP hashing mode utilizing all links

- **Investment protection**

- 1/10/100GbE scalable wire-speed connectivity options
- 2 million MAC entries for large VM environments

Brocade MLXe Series

Extensible 1G FCX stack

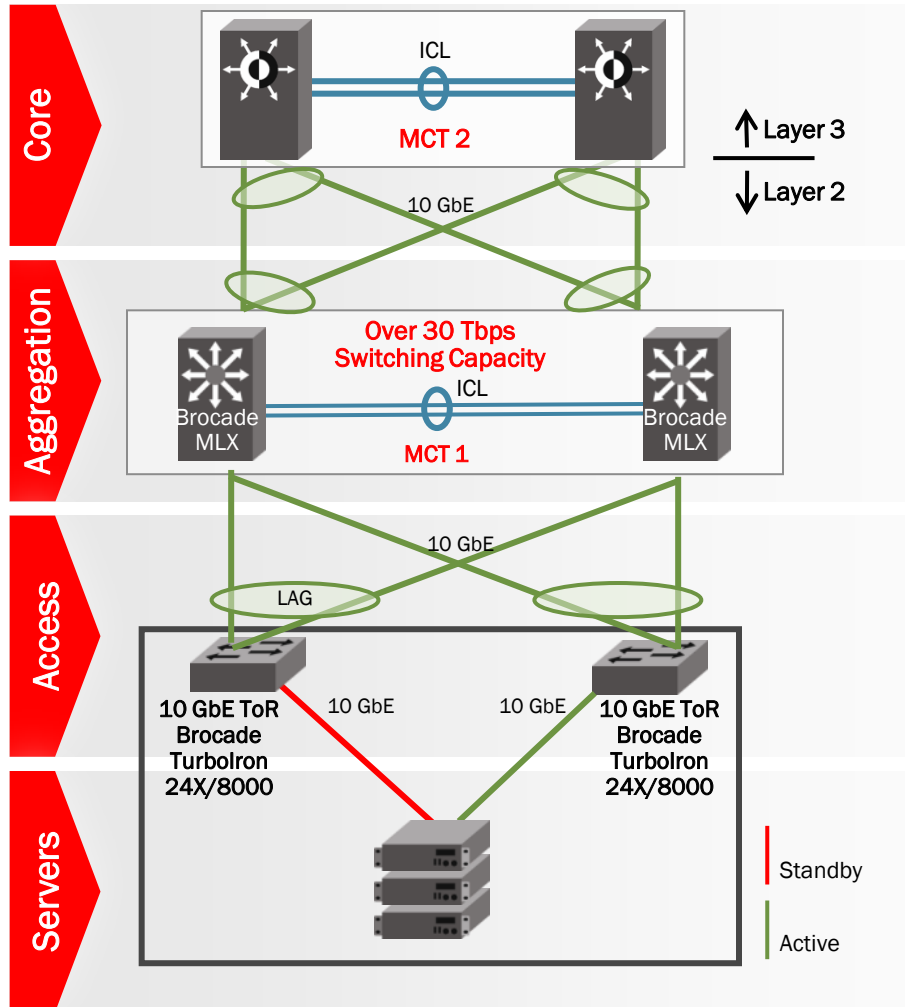


- **Easy to configure and manage**
 - No STP from server to core
 - Single management of Brocade FCX
- **Five times improvement**
 - Deterministic sub-200 ms link and node failover from server to core
 - Hitless failover within the FCX stack
- **Two times improvement**
 - All links forward traffic from server to core
 - Over 30 Tbps switching capacity
- **Investment protection**
 - 1/10/100 GbE scalable wire-speed connectivity options
 - Compatible with third-party devices
 - Pay-as-you-grow model for server access



Brocade MLXe Series

Resilient 10G architecture with TI24/8000



- **Easy to configure and manage**
 - No STP from access to core
 - Single pane NMS
- **Five times improvement**
 - Deterministic sub-200 ms link and node failover from access to core
- **Two times improvement**
 - All links forward traffic from access to core
 - Over 30 Tbps switching capacity
 - Non-blocking wire-speed 10 GbE
- **Investment protection**
 - 1/10/100 GbE scalable wire-speed connectivity options
 - Dual-speed 1/10 GbE with Brocade Turbolron
 - Server I/O consolidation FC and 10 GbE DCB/FCoE with Brocade 8000



Brocade MLXe Series

The Clear Choice

- **Unparalleled Scalability**

- Uniquely differentiated offerings to fit multiple business needs
- Flexible configuration options for more profitable services
- Affordable Larger Scale Rollouts with reduced “interconnect tax”

- **Industry Leading Performance**

- Highest density of line rate 1G & 100G interfaces
- Line rate routing performance for both IPv4 and IPv6

- **Investment Protection**

- Highly Programmable architecture that scales as you grow
- Flexible migration paths to higher density configurations without forklifts
- Simplified management. Virtualization ready.

- **Green**

- Lowest energy cost per Gbps
- Energy Efficient Front-to-Back airflow on highest density models

- **Lowest TCO**

- Low CapEx and OpEx
- Broad range of chassis sizes for deployment simplicity
- Proven Reliability track record lowers risk

Winner



Brocade MLXe

NetIron MLX 48-T-A Overview

Product overview

- Dense 1 GbE interface module
- Used in data center aggregation, HPC clusters
- Ideal for end-of-row and middle-of-row data center aggregation architectures
- Aggregates thousands of servers and reduces multiple layers in the data center

Key differentiators

- Industry-leading 1 GbE capacity in a single router
- Lowest power draw per GbE port
- All the advanced capabilities of the NetIron MLX modules
- Advanced data center virtualization with multi-VRFs, VLANs, and MPLS/VPLS
- Advanced load balancing and resiliency at the data center access layer



Highlights

- 48-port 1 GbE interface module
 - Eight MRJ21 connectors
 - Six 1 GbE links per connector
- Part Number: NI-MLX-48-T-A
- Requires high-speed fan upgrade on NetIron MLX-16 (part # NIBI-16-FAN-EXH-A)



8x10GE-M Module



| | |
|-----------------|---|
| Market Segments | All |
| Benefits | High density 10G module for large scale network build-out; reduced CAPEX and OPEX |
| Platforms | MLX |

Product Positioning

- High density wire-speed 10GE modules
- Leapfrog Cisco and Juniper to regain industry 10G leadership
 - 256 wire-speed 10GbE ports in a single MLX-32 router
- Significant investment protection – no forklift upgrades: Works on all MLX models

Hardware

- 8x10G-M
- New hSFM required

Competitive Differentiation:

- Industry leading 10G wire speed port density (256)
- Industry-leading power efficiency: < 31W per 10G port
- Converged Enhanced Ethernet (CEE) and H-QoS capable
- Extremely power efficient modules consuming ~42% less power consumption per 10GE than existing 4x10G
- Enhanced statistics for advanced services
- Deep buffers with Virtual Output Queues to accommodate large bursts (1GB per module)

Brocade FastIron CX Series Switches

Product highlights



Brocade Assurance
Limited Lifetime Warranty



Campus Access

- Highest PoE+ port density in a stackable design
 - Up to 48 PoE or 26 PoE+ ports in 1U
 - Up to 384 PoE or 208 PoE+ ports per stack
- Optimum flexibility and high availability
 - PoE+, non-PoE+ and 24 fiber ports models
 - All models support four 100/1000 MbE fiber uplinks
 - All models are field-upgradable to 2-port 10 GbE
 - Dedicated 64 Gbps stacking bandwidth (FDX)
 - Redundant, removable, load-sharing power supplies
 - Hot unit insertion/removal from the stack
 - Cross-unit link aggregation
 - Single IP management
- Advanced VoIP software features
 - Dynamic voice VLAN and QoS configuration
 - Dynamic power management through LLDP-MED
 - Robust security suite
 - sFlow for granular network traffic accounting
 - Support for variety of 802.3af and 802.3at devices
 - Support for full Layer 3 features and routing



Brocade FCX 624 and 648

Product highlights




Data Center Access

- Purpose-built data center design
 - 24 or 48 10/100/1000 Mbps ports in a 1U form factor
 - Optional 4 ports of 10 GbE for high-speed uplinks
 - Optional 4 100/1000 Mbps fiber uplinks
 - Efficient SFP+ 10 GbE optics
- Optimal flexibility and reliability
 - Horizontal stacking to create one logical switch per row
 - Hot unit insertion/removal from the stack
 - Cross-unit link aggregation from the stack
 - Single-IP stack management
 - Dual removable and load-sharing power supplies
 - Reversible front-to-back airflow
 - Field replaceable fan unit with 4 redundant blowers
- Advanced software features
 - Robust and secure Brocade IronWare OS
 - sFlow for detailed network traffic accounting
 - Enterprise-class Layer 3 routing
 - Optional BGP routing
 - Network-wide management with Brocade IronView Network Manager

Brocade Turbolron 24X Switch

Product highlights



- Dual 10/1 GbE ToR and Aggregation
- 10 GbE HPCC and iSCSI SAN/NAS
- Blade Server Aggregation
- Collapsed Midmarket Aggregation and Core

Data center-class performance

- Data center modularization with 10 GbE Top-of-Rack (ToR) solution
- Up to 488 Gbps line rate, non-blocking performance
- Up to 24 10 GbE SFP+ ports in 1RU
- Cut-through architecture with ultra-low latency
- Front-to-back cooling

Greatest flexibility

- 24 1 GbE/10 GbE ports for seamless migration
- Four 10/100/1000 Mbps copper ports
- Redundant, removable, and load-sharing AC power supplies
- Hot-swappable triple-fan assembly

Advanced software features

- Advanced Layer 2/3 features
- Hardware-based security features and sub-second failover
- sFlow for detailed network traffic accounting

Operational efficiency

- Best-in-class power efficiency
- Rack space saving
- Brocade Assurance Limited Lifetime Warranty



Brocade ServerIron ADX Application Delivery Controllers (ADCs)

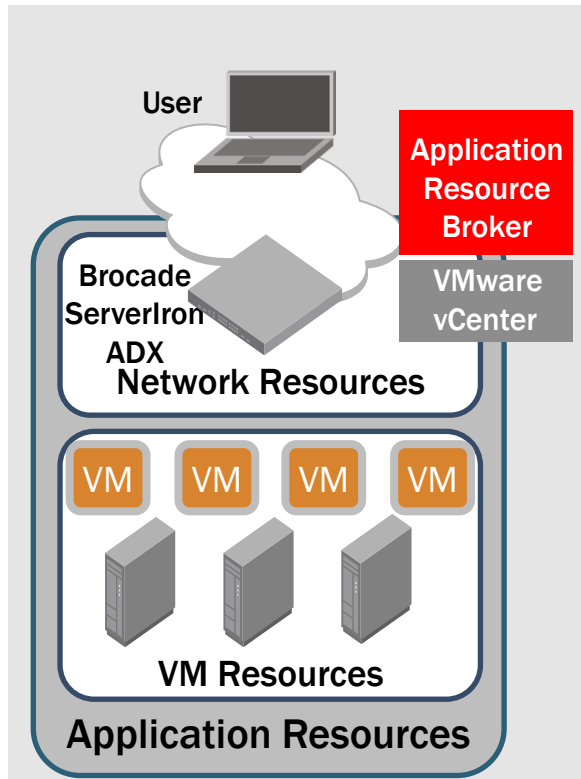
How Did Brocade ADCs Become the Key to Virtual Server Provisioning?

- Leading performance and scalability
 - 70 Gbps Layer 4-7 throughput
 - 16 million Layer 4 transactions/sec
 - 120 million SYN attacks/sec DoS protection
 - 14 million DNS queries/sec
 - 320 Gbps switching
- Optimum flexibility
 - Fixed and chassis configurations with interchangeable modules
- Advanced functionality
 - Content switching and rewrite
 - Transparent cache switching
 - Hardware SSL acceleration
 - TCP/HTTP multiplexing
 - Multi-site redundancy using GSLB, FWLB
 - VMware application provisioning

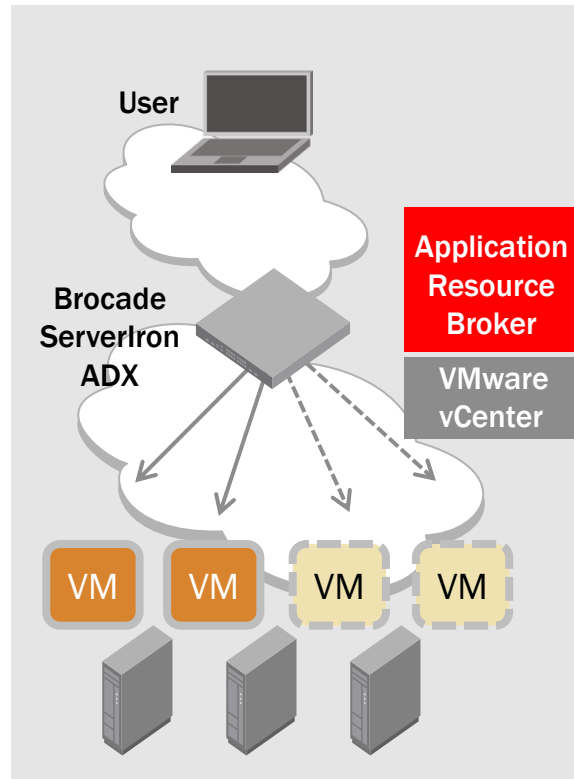


Brocade Application Resource Broker

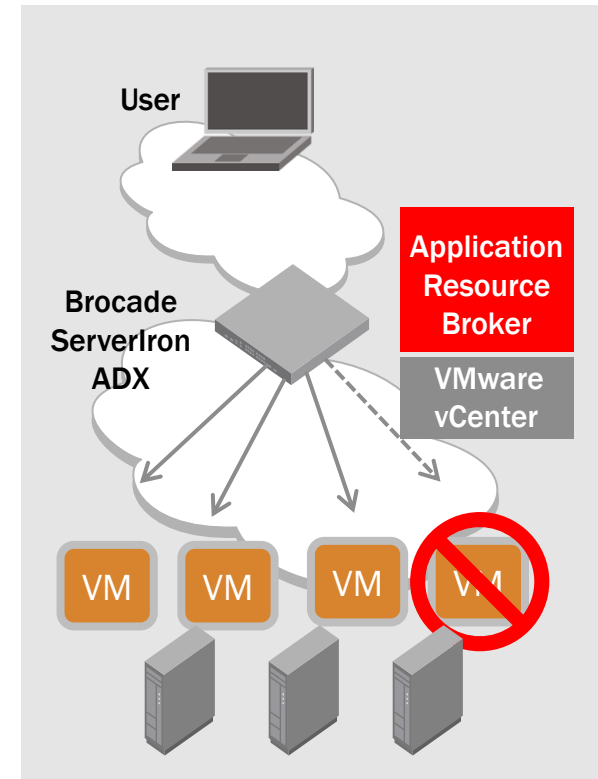
Matching Data Center Capacity with End-User Demand in Real Time



User Experience and
VM Resource
Monitoring



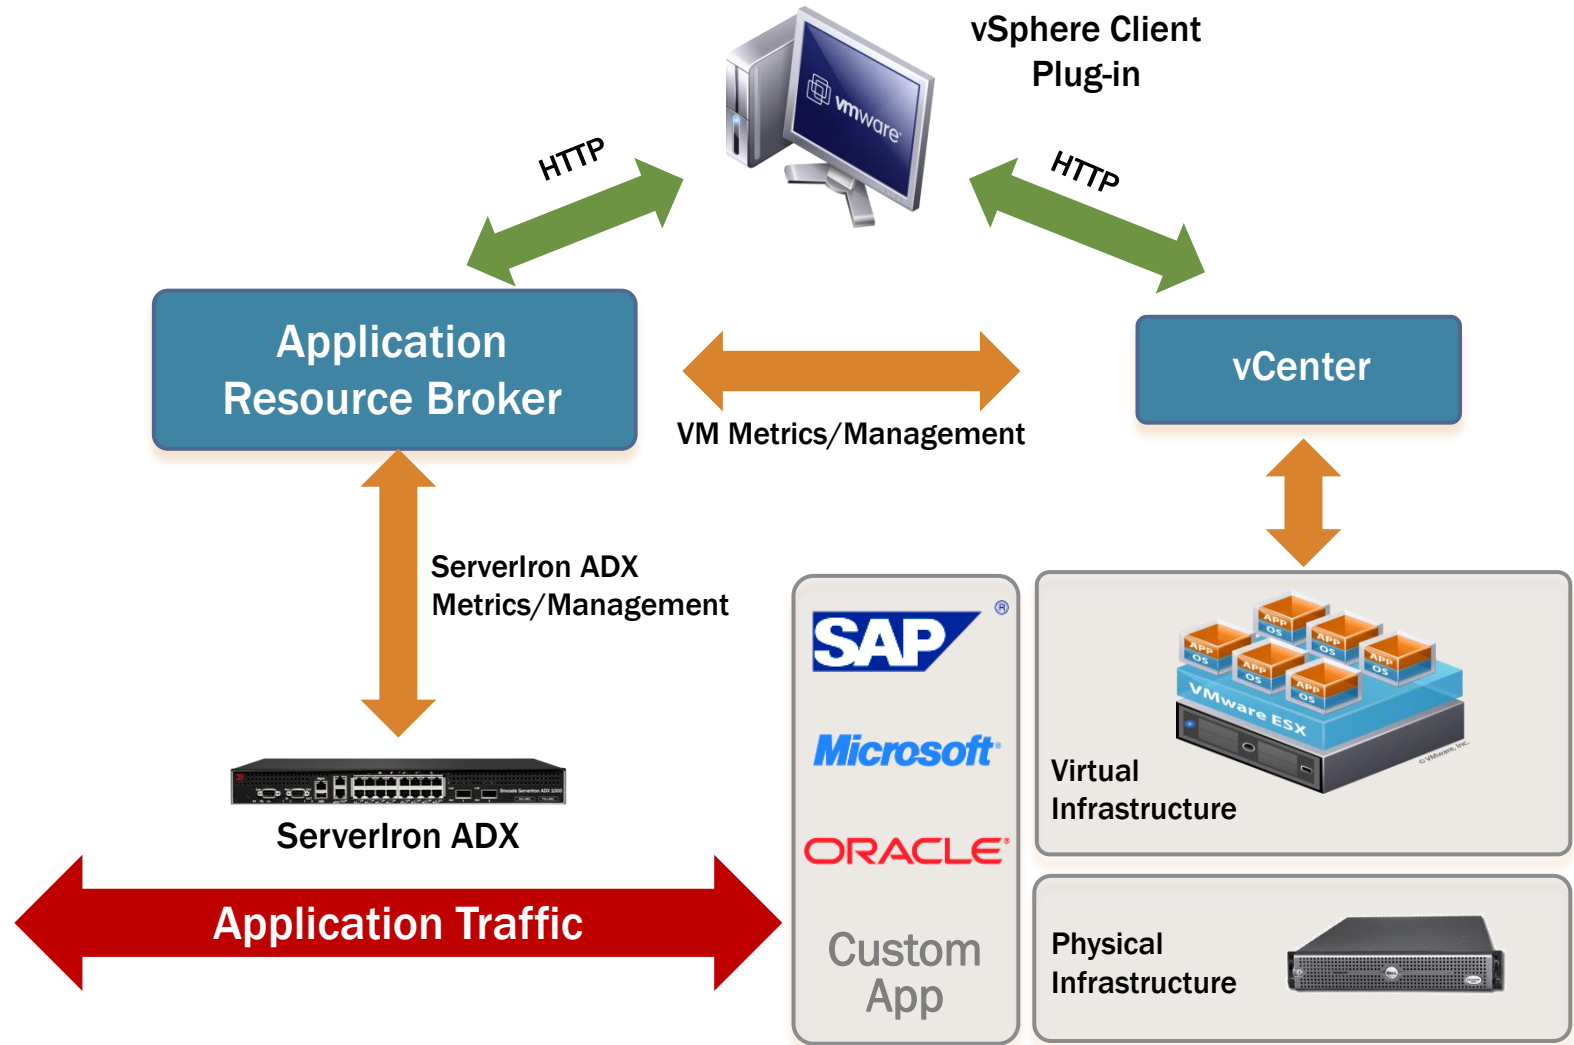
Resource
Commissioning
When Load Increases



Resource
Decommissioning
When Load Decreases

Brocade Application Resource Broker

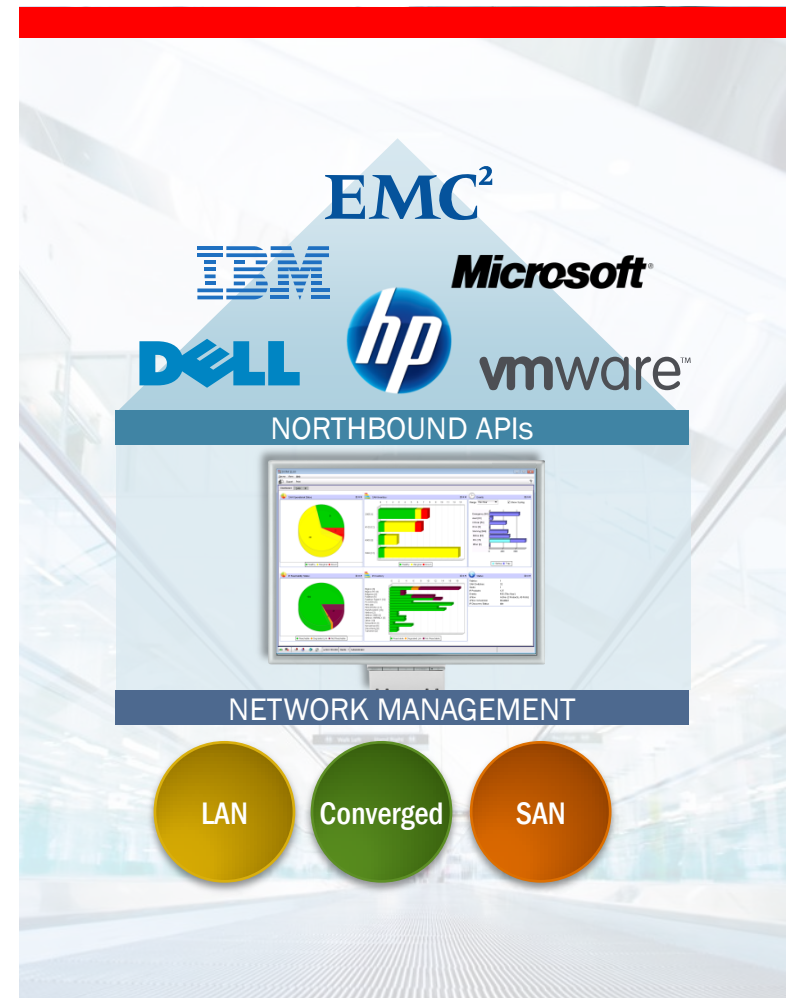
Solution Architecture



Brocade Network Advisor

End-to-End Service Orchestration with Leading Partner Products

- Open architecture with industry-standard APIs (SMI-S, Web Services, NETCONF, SNMP)
- Seamless integration with
 - Orchestration Frameworks
 - Security Partners
 - Service Delivery platforms
- VMware and Microsoft hypervisor plug-ins



Brocade Network Advisor

Simplified Management for Data Center, Enterprise Campus and SP Networks

- Unified Network Management product for SAN, IP, Application Delivery, and Converged Networks
 - One management GUI across FC, IP/MPLS, FCoE protocols
 - Custom views based on Operator specialization
 - Flexible user management with Role Based Access Control
- Standards-based architecture
- Provides seamless integration with leading partner products

