Improving Your Ethernet Services with Virtual Cluster Switching (VCS)

Session 8260
Dr. Steve Guendert
Brocade Communications
sguender@brocade.com

Agenda

- Both sides now
- VCS Technology Details
  - Ethernet Fabric
  - Distributed Intelligence
  - Logical Chassis
  - Dynamic Services
- VCS Use Cases

Both sides now

- **Rows and flows of angel hair,**
  And ice cream castles in the air,
  And feather canyons everywhere,
  I've looked at clouds that way. But now they only block the Sun,
  They rain and snow on everyone.
  So many things I would have done,
  But clouds got in my way.
- **I've looked at clouds from both sides now,**
  From up and down, and still somehow,
  It's cloud illusions I recall,
  I really don't know clouds, at all.

Skyrocketing Data Growth and Network Complexity

>900 EB	Digital Data Created by Individuals in 2010[1]
> 35 ZB	Digital Data Forecasted by 2020[1]
10X	Growth in Virtual Machines by 2012[1]

1. Search Storage, May 2010

The Network Is Central to the Cloud

- Virtualization Brings New Requirements and Challenges

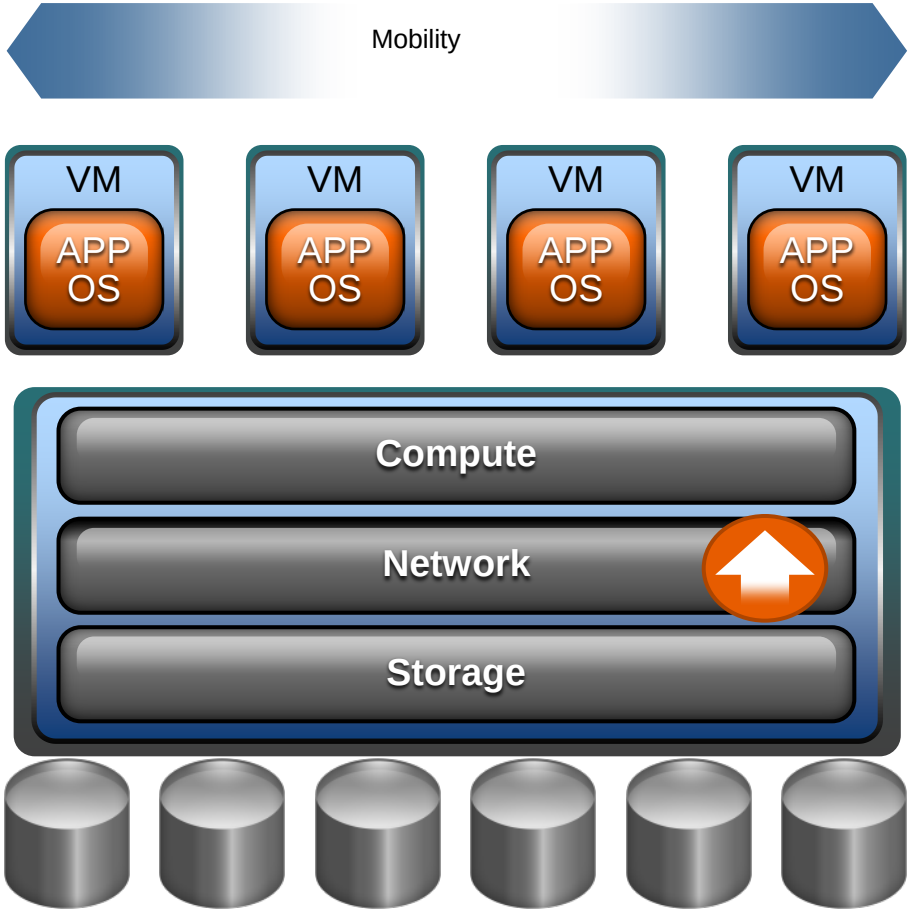

Challenges

- Network performance/scalability constraints
- Application resiliency and performance under load
- VM mobility limits
- Infrastructure complexity
- Management silos

The Cloud Enabled Data Center

- Building the Virtualized Data Center

SIMPLER, AGILE NETWORK

Fewer layers

Full performance

High resiliency

Flexible transport

Shared intelligence

Lower OpEx/CapEx

MAXIMUM COMPUTE AND STORAGE

VM scale and mobility freedom

Hypervisor offload

Predictive application provisioning

Fibre Channel, iSCSI, NAS, FCoE optimization

Ecosystem leverage

A SINGLE VIEW

Management "hand shake" across tools/orchestration

Comprehensive monitoring

Predictive event notification and response

Choice of VMs

LOWER COST OPTIMUM EFFICIENCY CLOUD-ENABLED

Cloud Implications

- The Dawn of the Virtual Enterprise

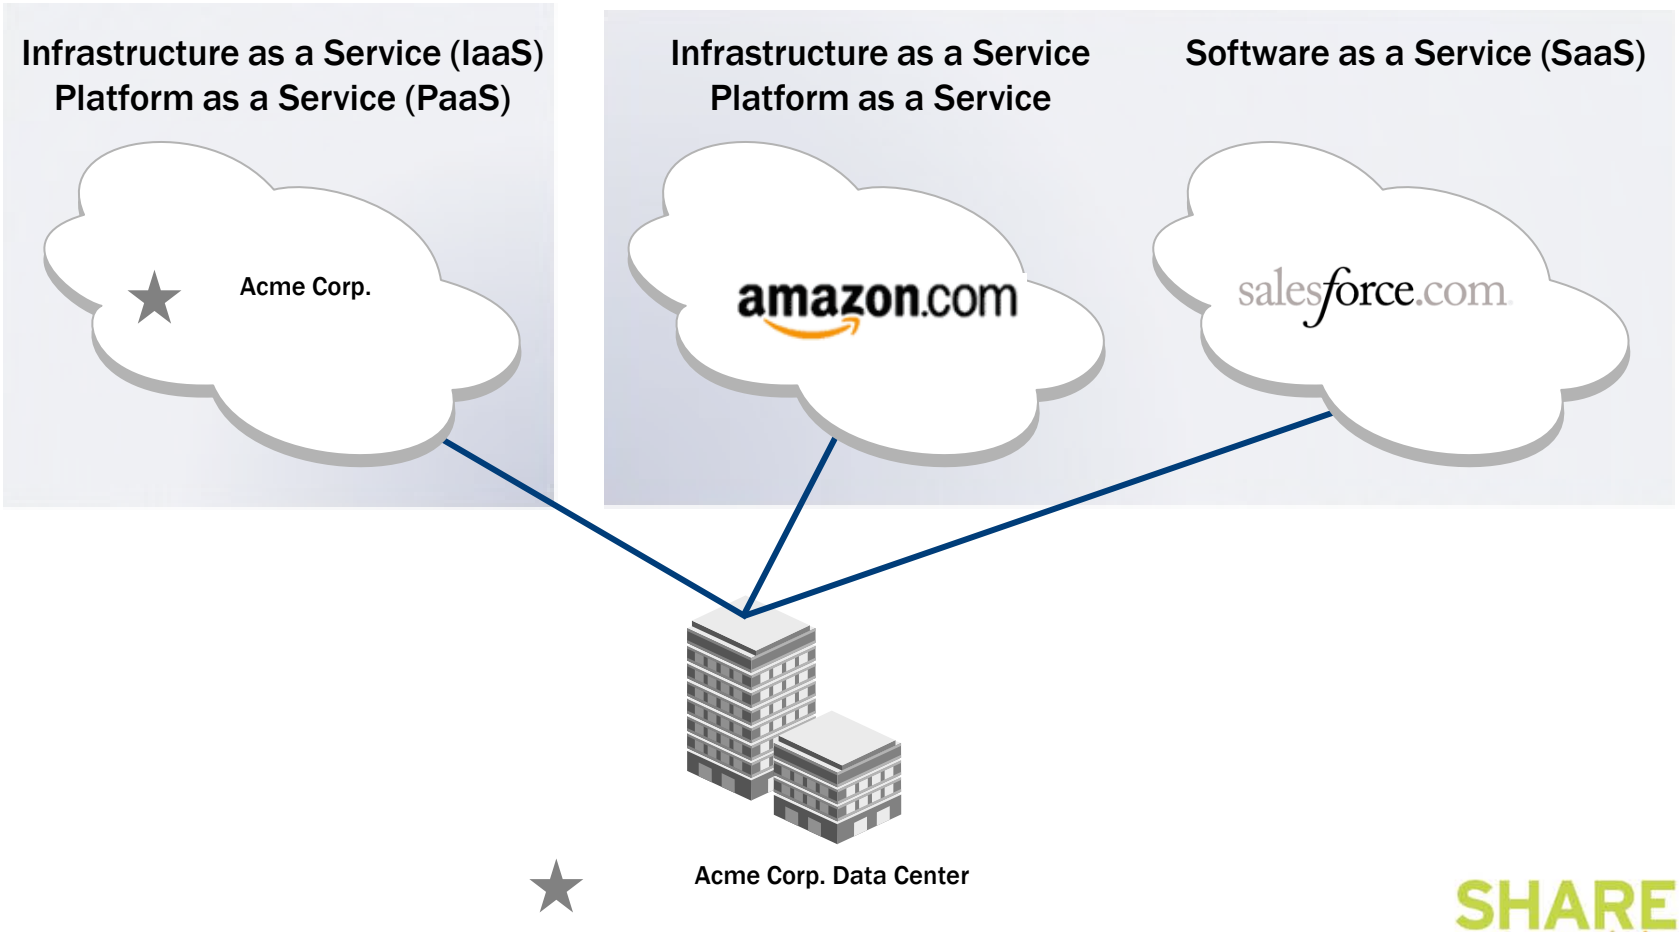

Virtual Data Centers Are Key

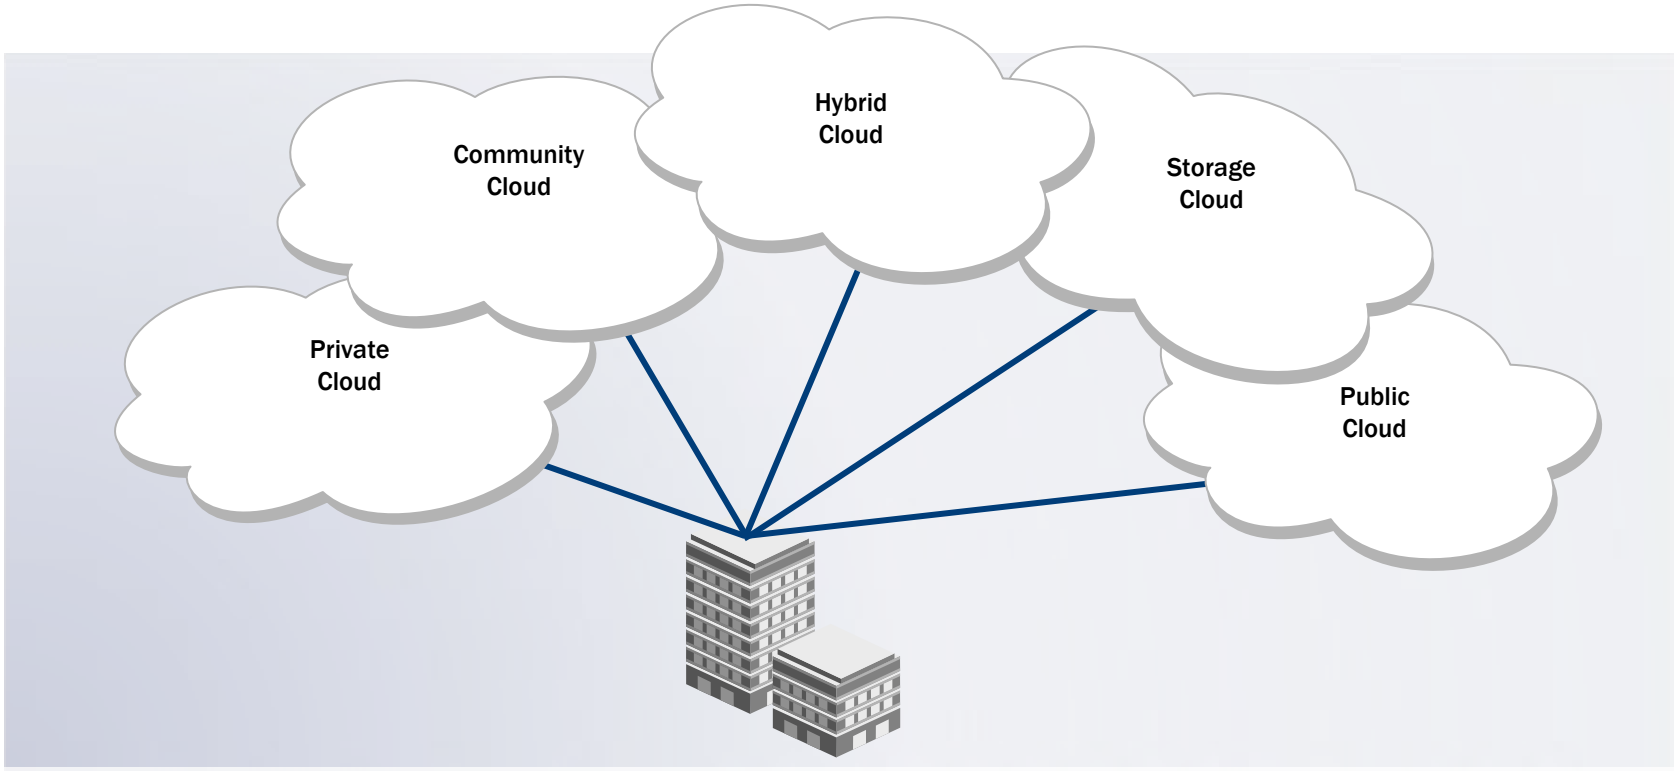

Physical

Data Center Construction

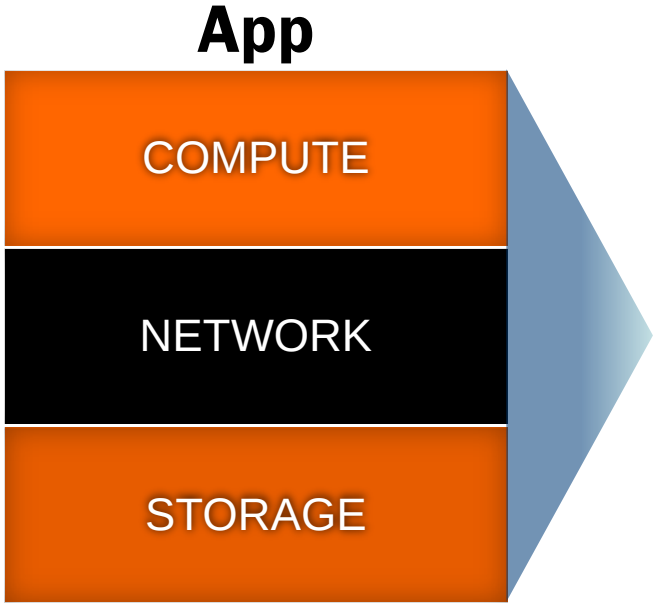

Traditional Construction

Manual

Replication

Rigid

Inflexible

15%
Utilization

70%
Operating Cost

Data Center Construction

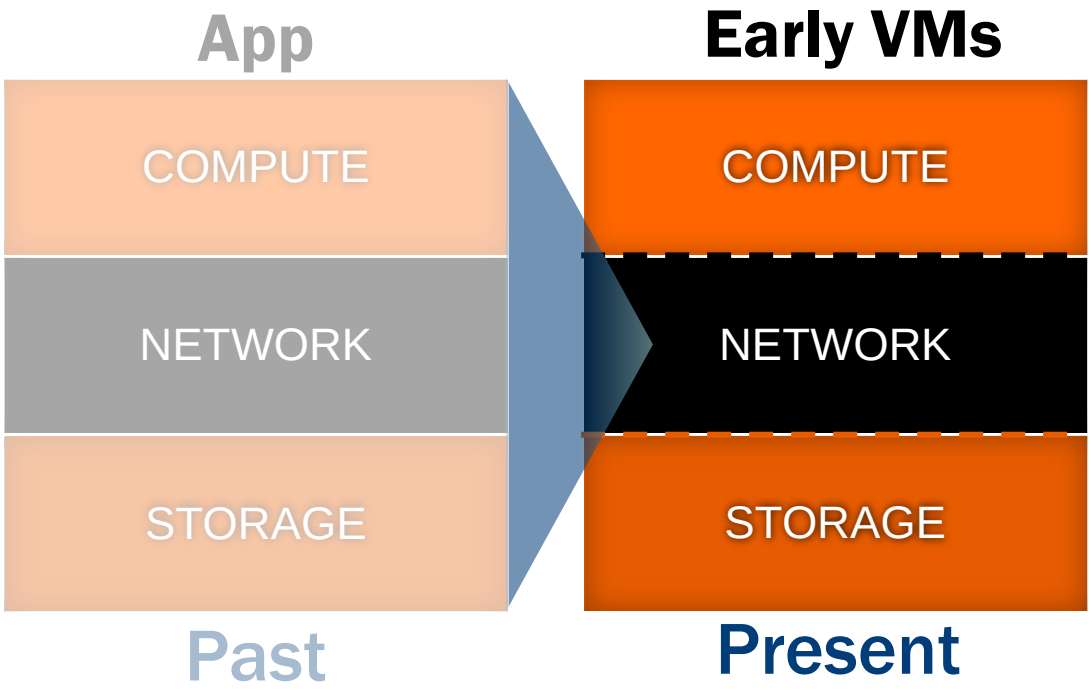

50%
Virtual Apps (2012)[1]

Escalating Complexity

[1] Source: Gartner, March 2010

Data Center Construction

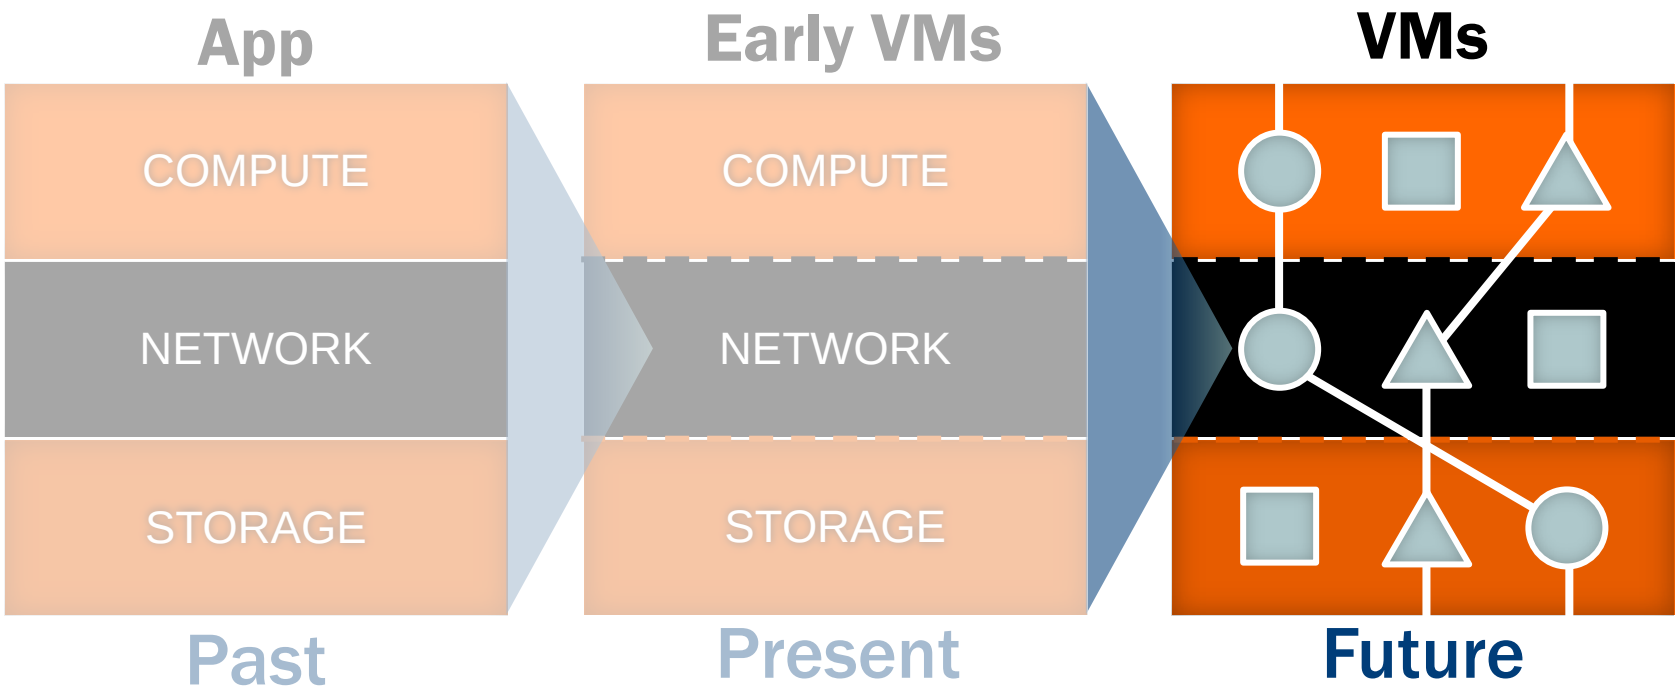

Challenges of Current Data Center Networks

- Must Be Solved for Virtualized Data Centers
  - Layer 2 performance, scalability, reliability
    - Limitations of Spanning Tree Protocol (STP)
  - Scaling virtual server environments
  - Virtual machine mobility
    - Must run in Layer 2; IP address stays the same
    - Enforce the same policies and permissions
  - Infrastructure complexity
    - Lots of switches to manage
    - Layer 3 protocols to the edge
  - Management overhead
    - High operational costs

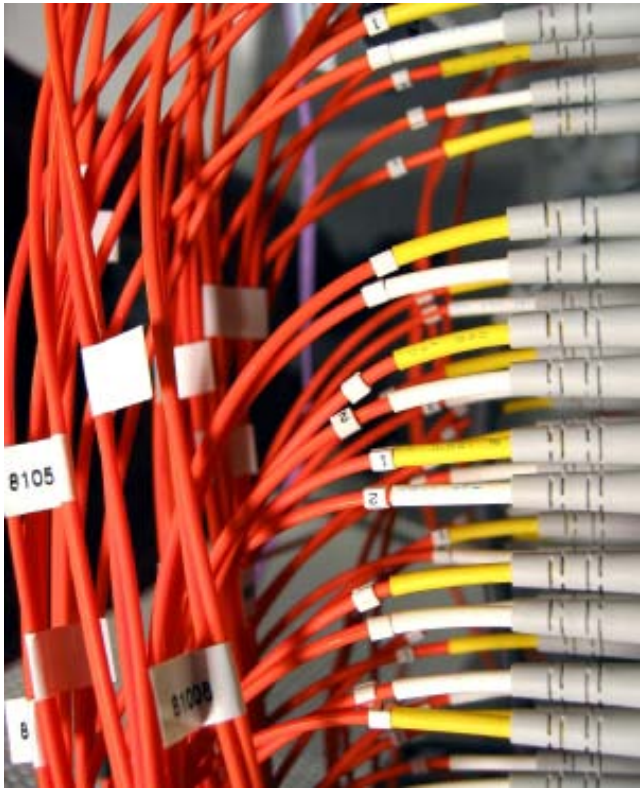

Brocade Virtual Cluster Switching (VCS)

VCS

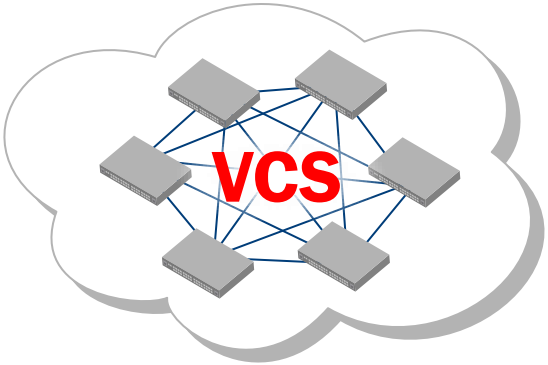

- First true data center Ethernet fabric
- Revolutionizes Layer 2 connectivity
- Increases scalability of virtual server environments and sphere of mobility
- Maximizes network performance—reduces network complexity

Brocades VCS

- VCS technology comprises three main innovations that extend the capabilities of a typical Layer 2 network:
  - Ethernet fabric
  - Distributed intelligence
  - Logical chassis

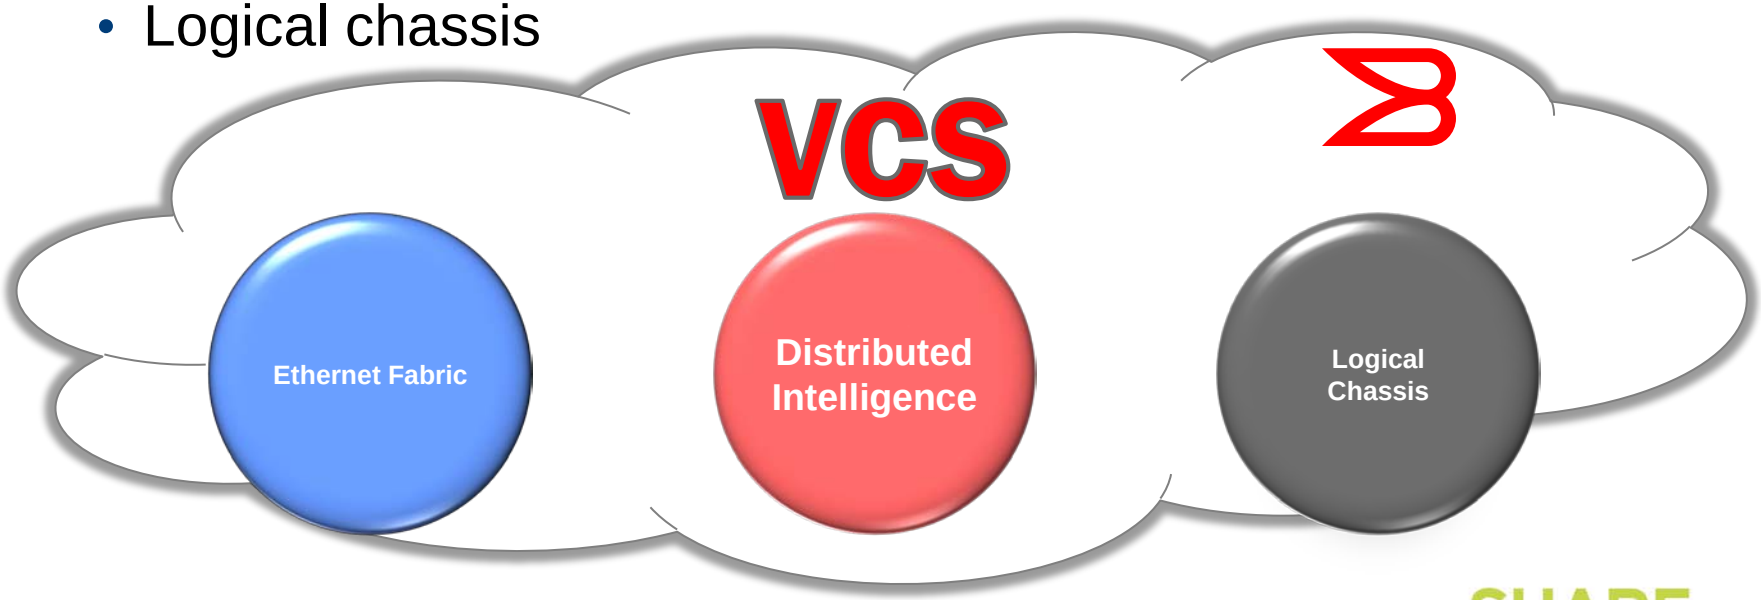

Scaling Virtual Server Environments

- Challenges Today

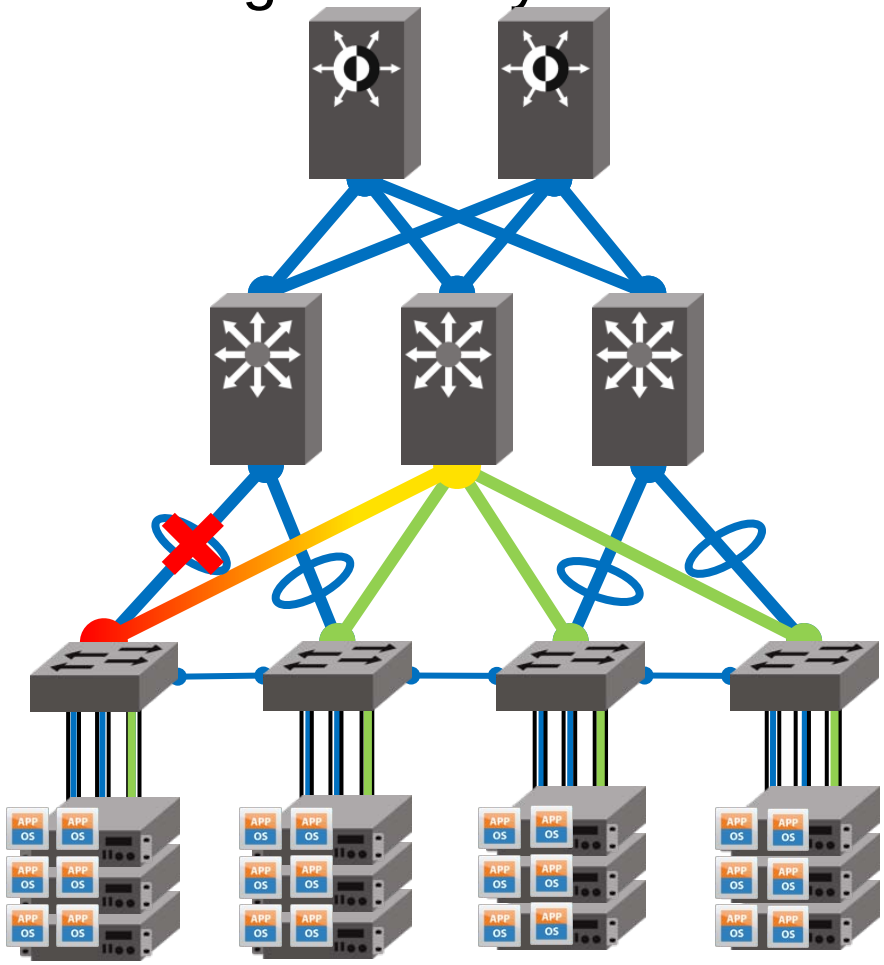

Layer 2: Only one active path

- STP disables other paths
- Not “virtualization-optimized”

Add virtual machines

- Add 1 GbE connections
- Move to 10 GbE for simplicity and higher 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

Ethernet Fabrics

A new network architecture

SHARE
Technology • Connections • Results

Classic Hierarchical Ethernet Architecture

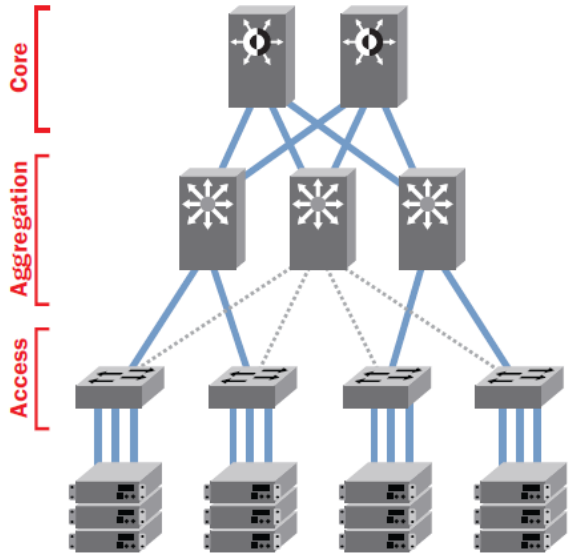

Servers with 10 Gbps Connections

- Classic architectures often require three tiers in the physical network
- STP disables links in the fabric to prevent loops, limiting network utilization
- Each switch has to be managed individually

Ethernet Fabric Architecture

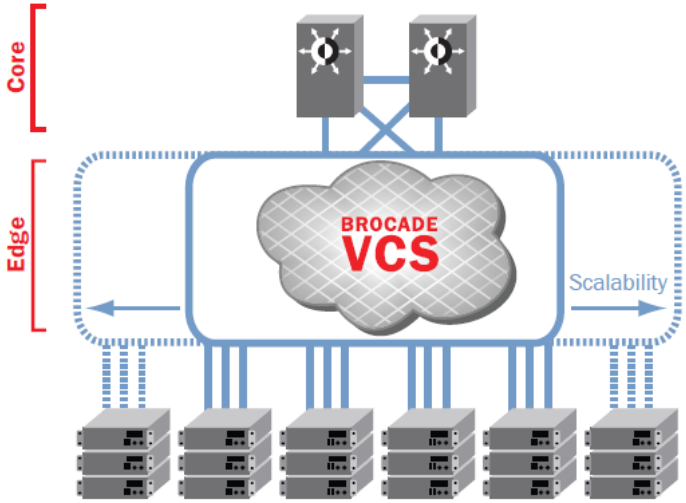

Servers with 10 Gbps Connections

- Fabric architectures flatten and seamlessly scale out the Layer 2 network at the edge
- All links in the Ethernet fabric are active when utilizing VCS technology

SHARE in Anaheim 2011

VCS Technology and Ethernet Fabrics

- In VCS technology, Ethernet fabrics are defined as a group switches exchanging information between each other to implement distributed intelligence
  - Interconnected using regular front-end ports
  - Presented as one unified and transparent Ethernet switching service to the external network
- Ready for FCoE and iSCSI traffic
- Extends existing Ethernet infrastructure
- Fabric auto-configures
  - Once VCS is enabled, only minor configuration is necessary

Ethernet Fabric Components

- Data Center Bridging (DCB)[1]
  - Data Center Bridging Exchange (DCBX)
  - Priority-based Flow Control (PFC)
  - Enhanced Transmission Selection (ETS)
- Transparent Interconnection of Lots of Links (TRILL)
  - Active multipath
  - Multihop routing
  - Highly available, rapid link recovery
- Fabric services
  - Link state routing provided by Fabric Shortest Path First (FSPF) routing protocol
  - Ethernet Name Server (eNS)

VCS: Distributed Intelligence

- Distributed Fabric Services
  - Self-forming fabric (with minimum configuration)
  - Information shared across all fabric members
  - Fabric aware of all connected devices
- Masterless control
  - Switch or link failure does not require full fabric reconvergence
- Shared port profile information
  - Automatic Migration of Port Profiles (AMPP)
    - Enables seamless virtual server migration

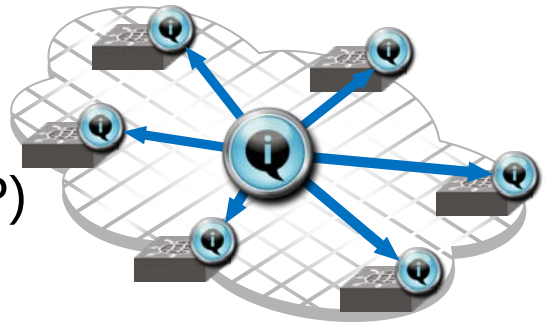

VCS: Logical Chassis

- Fabric behaves as a single logical chassis
  - Devices outside the cloud sees one switch
- Logically flattens and collapses network layers
  - Fabric is self-aggregating
  - Flexible fabric topologies

Ethernet Fabric Details

Enabling End-to-End Convergence

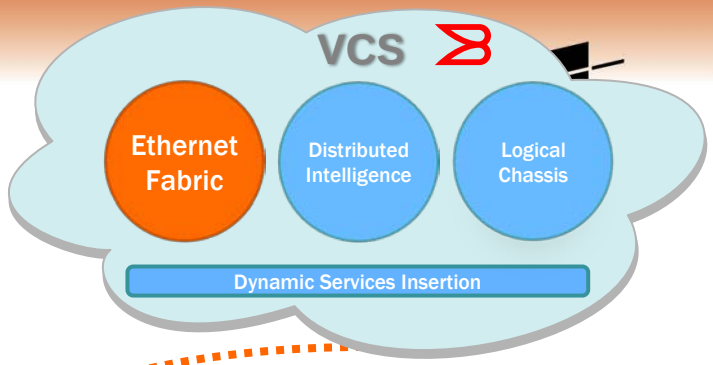

- Massive scale for VMs
- Eliminates STP
- Provides multi-pathing, reliability, and increased utilization of links
- Enables end-to-end network convergence
- Simplifies configurations and diagnostics
- Storage-enabled

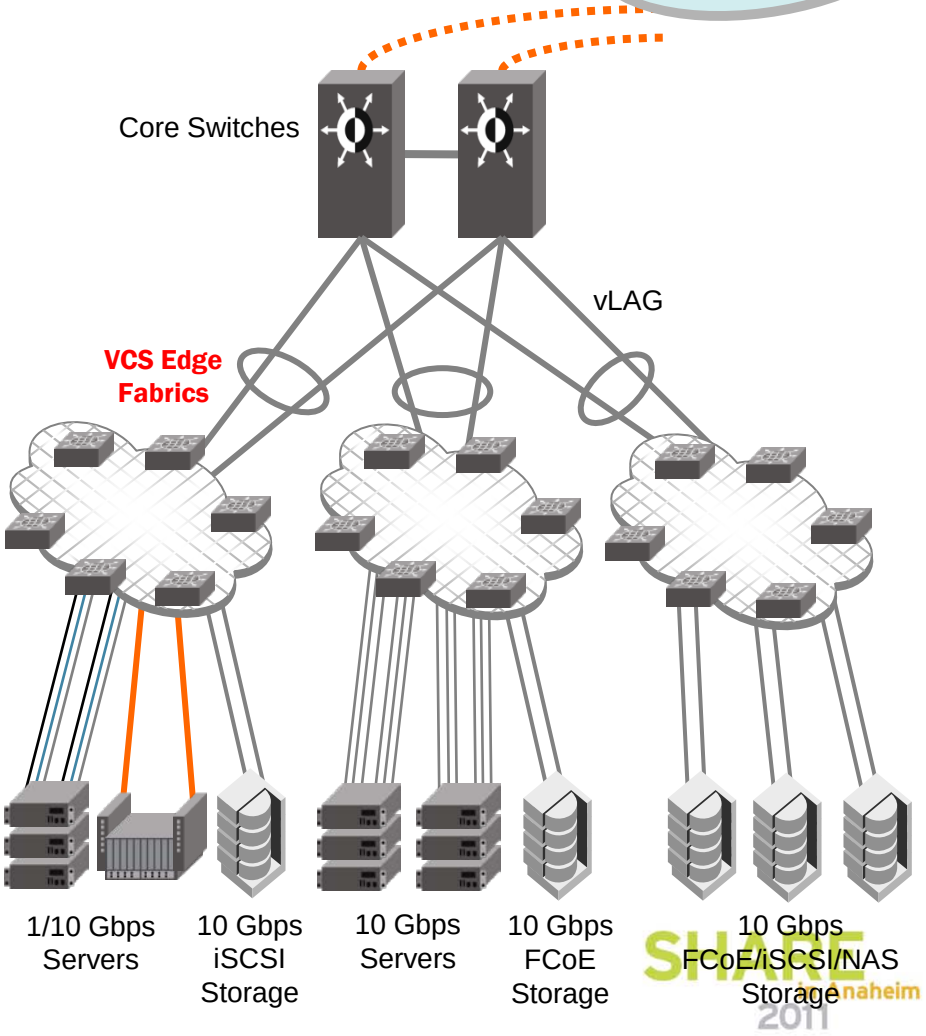

Ethernet Fabric Details

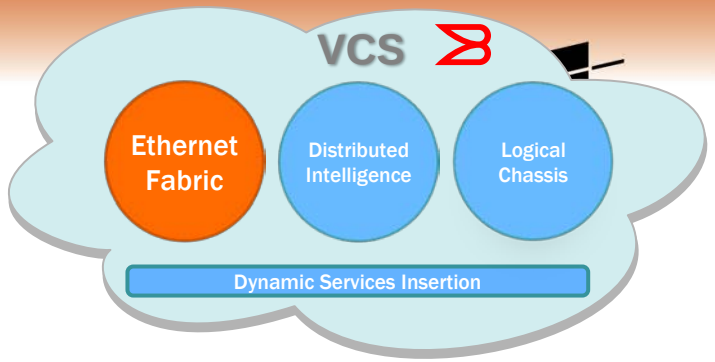

- First true Ethernet fabric
  - Layer 2 intelligent, lossless network
- Link-speed agnostic
- Data Center Bridging (DCB)
  - Lossless, deterministic
  - Priority-based Flow Control (PFC)
  - Enhanced Transmission Selection (ETS)
  - Data Center Bridging Exchange (DCBX)
- Transparent Interconnection of Lots of Links (TRILL)
  - Active multi-path
  - Multi-hop routing
  - Highly available, sub-250 ms link recovery
- LAN/SAN convergence-ready
  - FCoE, iSCSI, and NAS traffic
- Standards-based
  - Integrates into existing Ethernet infrastructure outside of fabric

Ethernet Fabric Details

What is a Fabric?

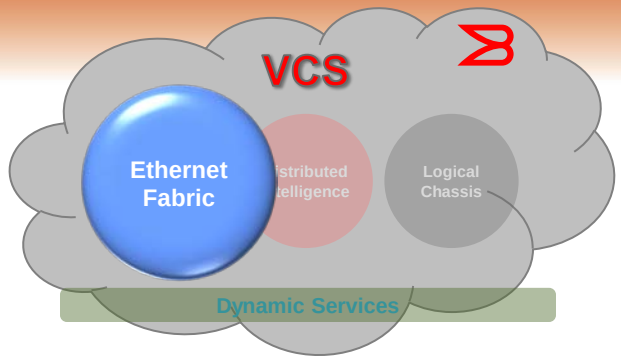

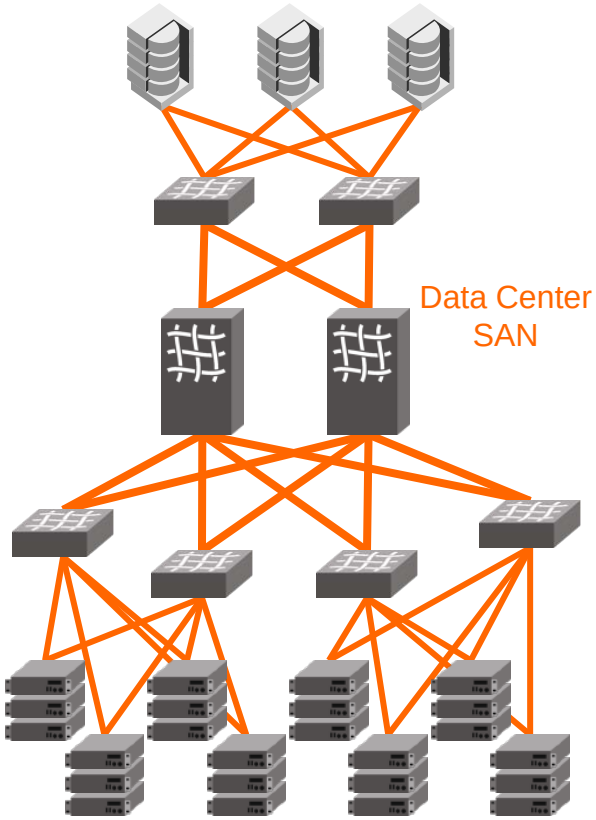

- Common fabric attributes
  - Switched network
  - Fabric members and devices connected always know about each other
  - All paths are available for high performance and high reliability
  - Traffic travels across the shortest path
  - Traffic can be routed from fabric to fabric

SHARE in Anaheim 2011

Ethernet Fabric Details

Data Center Bridging (DCB)

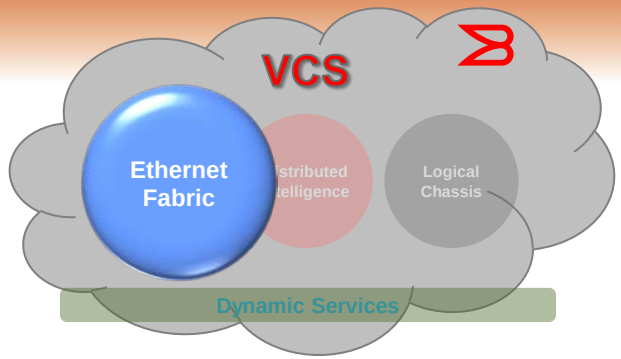

- Making Ethernet Lossless
- 802.1Qbb – Priority-Based Flow Control
  - PFC: Allows Identification and prioritization of traffic
- 802.1Qaz – Enhanced Transmission Selection/Data Center Bridging Exchange
  - ETS: Allows grouping of different priorities and allocation of bandwidth to PFC groups
  - DCBX: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

SHARE in Anaheim 2011

Ethernet Fabric Details

Transparent Interconnection of Lots of Links (TRILL)

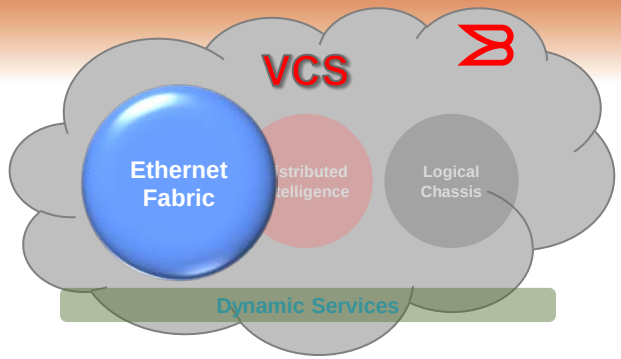

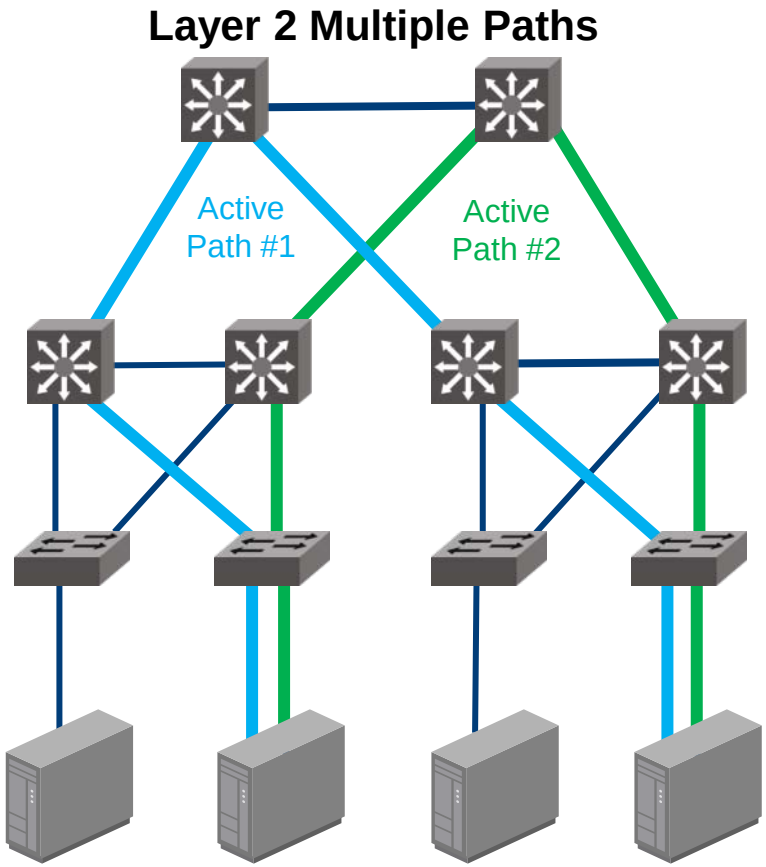

- Multi-path Layer 2 switching
  - All paths are active and traffic is distributed across all paths
  - Fully utilize all fabric bandwidth
- Establishes shortest paths through the Layer 2 fabric
- Uninterrupted response to link failures
- Backward-compatible and connects into existing infrastructures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

SHARE in Anaheim 2011

Distributed Intelligence

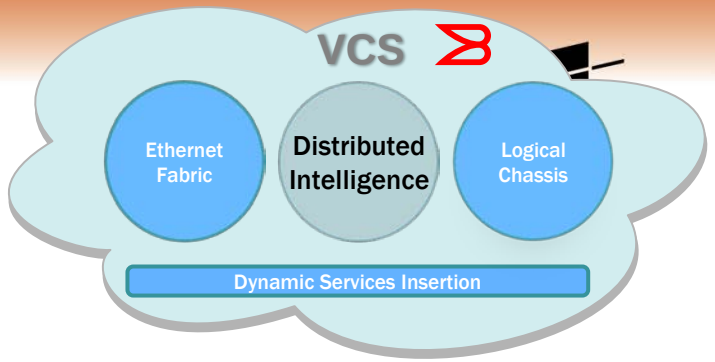

- Distributed Fabric Services
  - Fabric is self-forming
  - Information is shared across all fabric members
  - Fabric is aware of all devices connected
- Masterless control
  - Switch or link failure does not require full fabric reconvergence
- Shared port profiles information
  - Automatic Migration of Port Profiles (AMPP)
  - Enables seamless VM migration without compromise
- Optimized Virtual Access Layer
  - VEPA; frees host resources from switching and policy enforcement

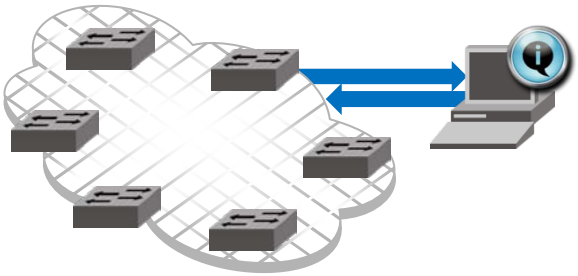

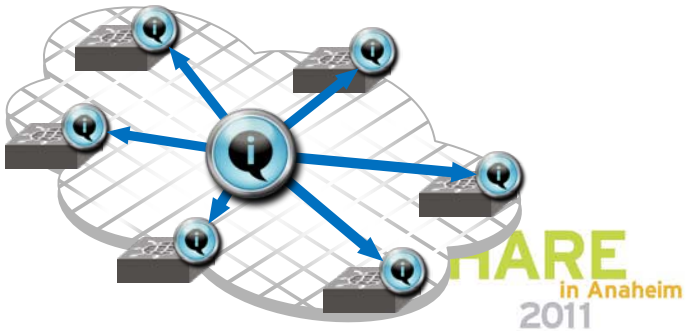

Distributed Intelligence

Automatic Migration of Port Profiles (AMPP)

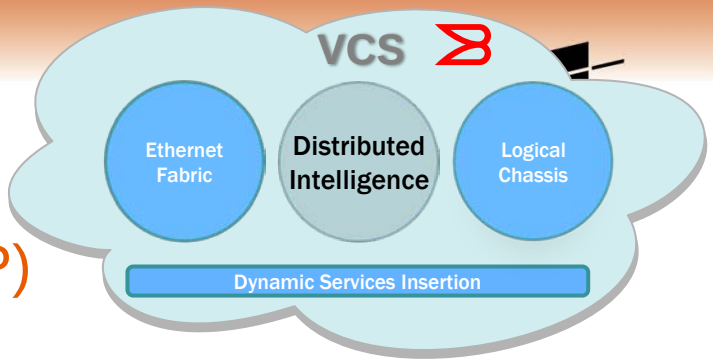

Allows VM to move, with automatic configuration in the network

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

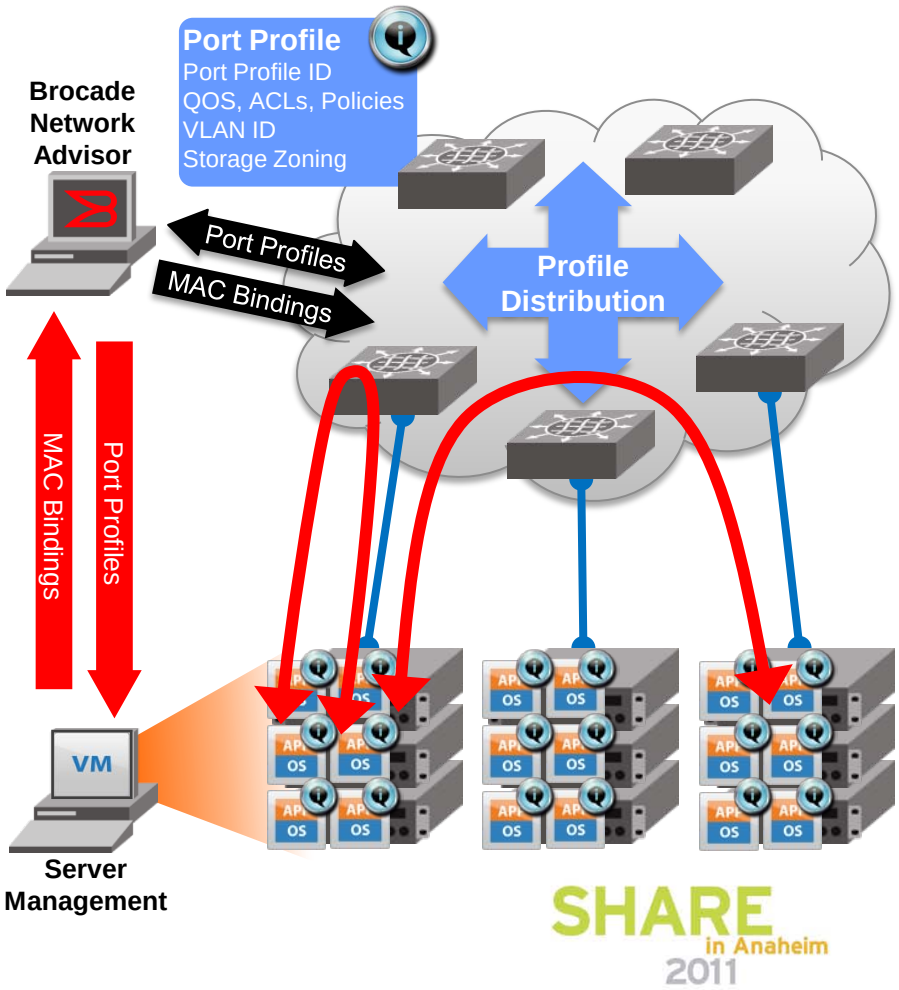

Distributed Intelligence Details

Optimized Virtual Access Layer

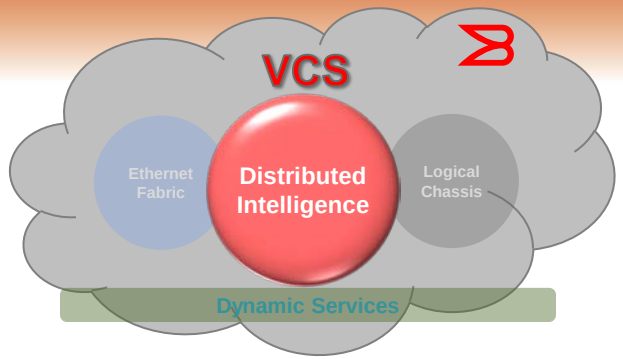

- Today, access to the network lives in the virtual hypervisor
  - Consumes valuable host resources
- Virtual switch is offloaded to the physical switch
  - Eliminates the software switch; the advantages of a distributed virtual switch plus Distributed Intelligence
  - Leverages Virtual Ethernet Port 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

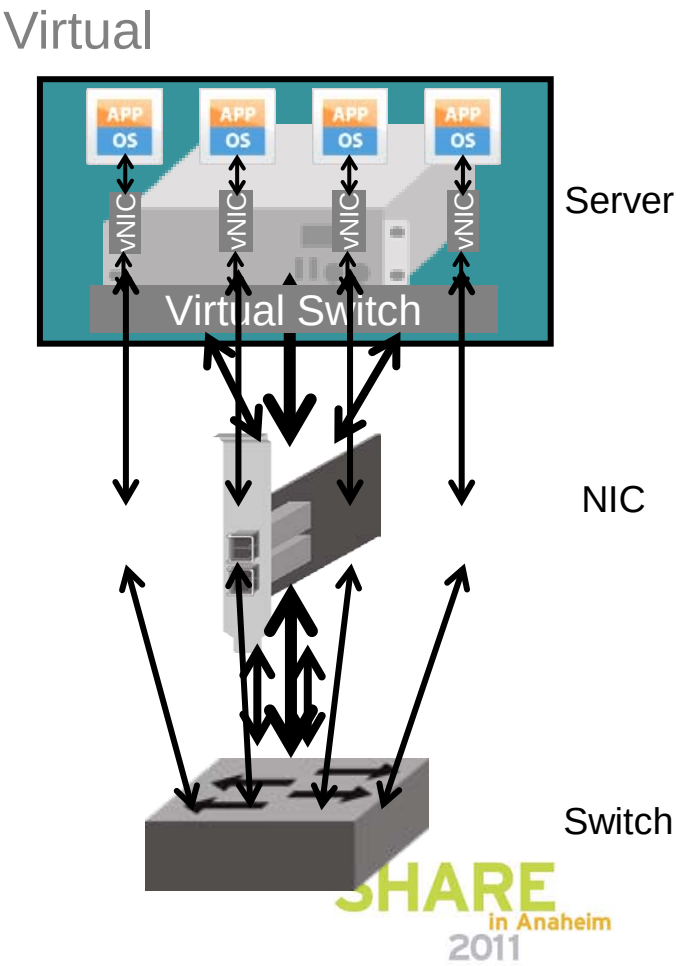

Logical Chassis Details

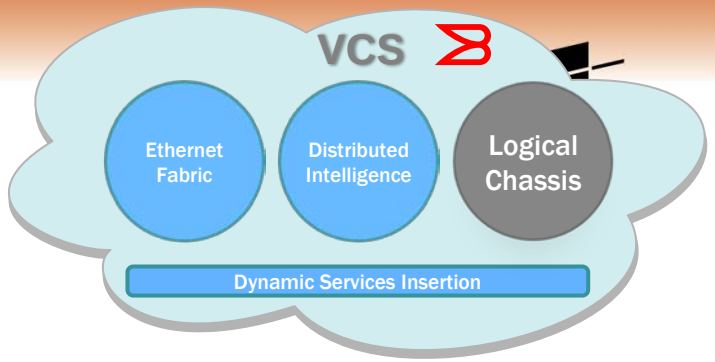

- Fabric auto-configures
  - Once VCS is enabled, no configuration is necessary
- Fabric behaves/is managed as a single logical chassis
  - Aggregation (or core) layer sees one standard Ethernet switch
  - Fabric members act like a blade in a chassis
- Logically flattens and collapses network layers
  - Fabric is self-aggregating
  - Flexible fabric topologies
- Scales to 1000s of ports without added management
  - Acts as a standard Ethernet switch outside of the fabric

Logical Chassis Details

Auto-Configuration

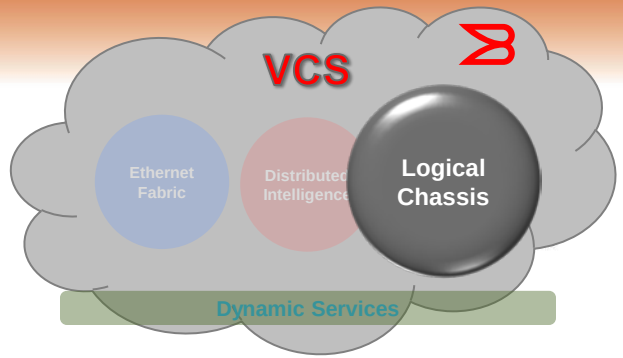

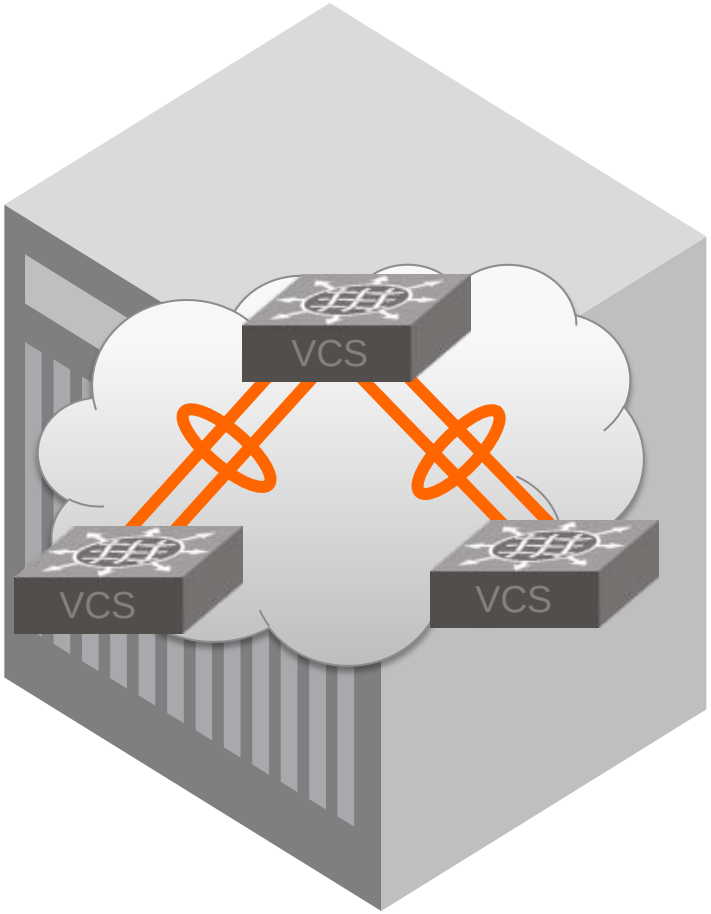

- VCS simplified deployment, scalability, and management of the network
- Enable VCS on each switch
- Connect the switches
- Fabric automatically forms
  - Common configuration across all switches
  - vLAGs auto-configure
- Managed as a single logical chassis

SHARE in Anaheim 2011

Logical Chassis Details

Single Logical Switch Behavior

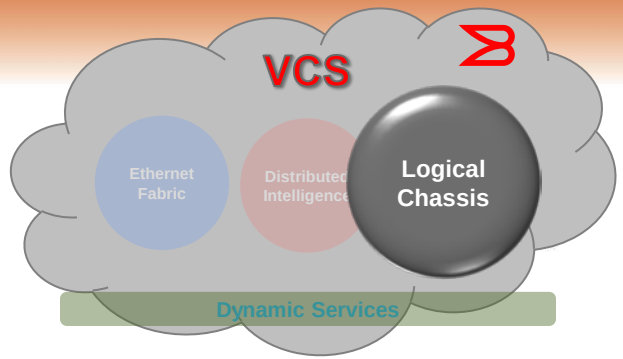

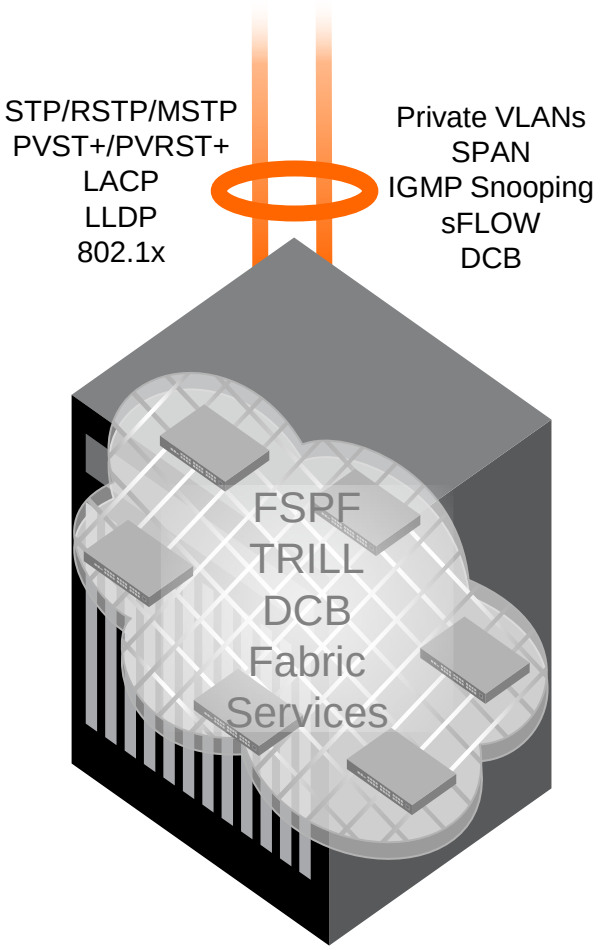

- VCS behaves like a single industry-standard Ethernet switch
  - Fabric members are like blades in a modular chassis
- Standards-based and closed protocols used within the fabric
  - FSPF, TRILL, Fabric Services, etc.
- Industry-standard protocols used to communicate outside the fabric
  - RSTP, LACP,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

SHARE in Anaheim 2011

Intelligent Services

Challenges Today

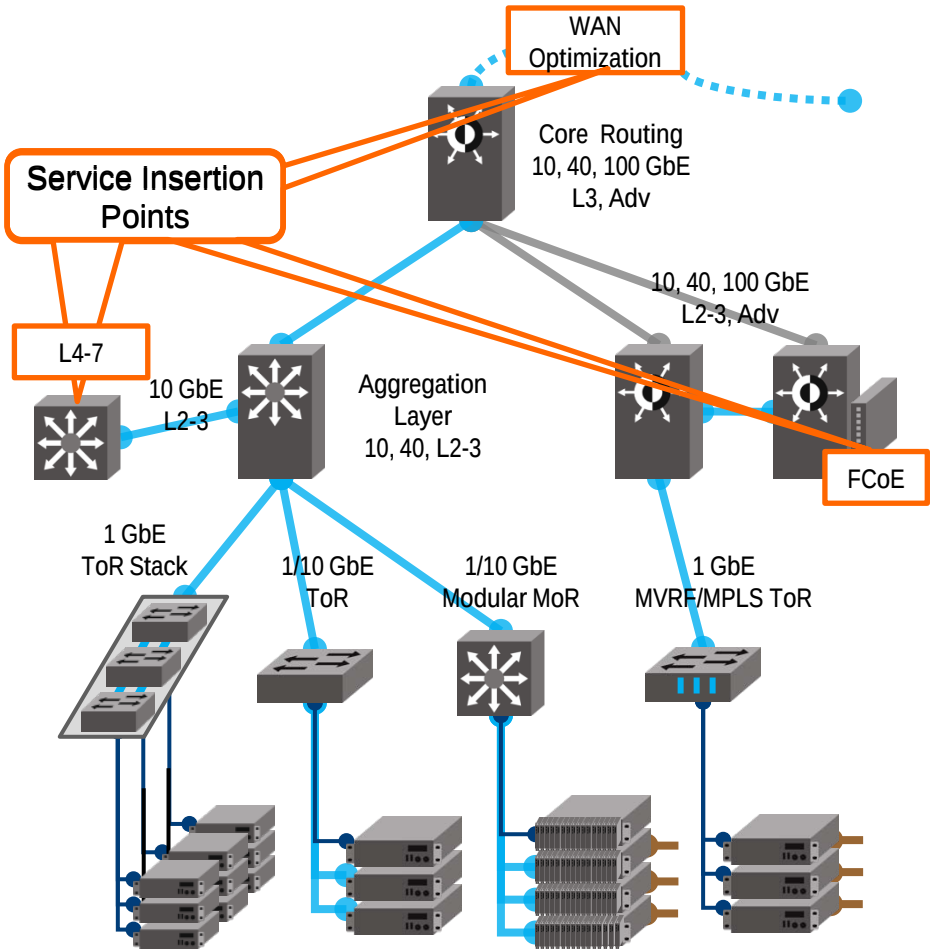

- Complex to engineer, deploy, and manage
- Static and rigid
- Leads to reconfiguration and connectivity interruption

Dynamic Services Details

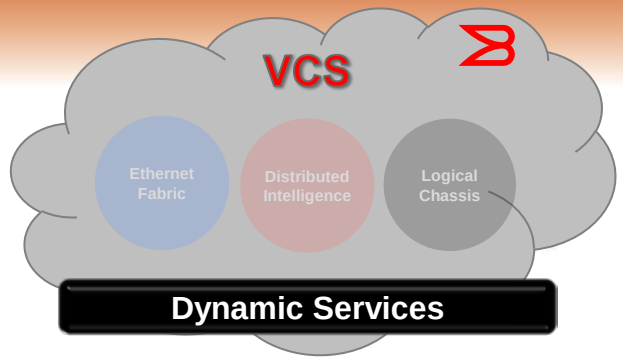

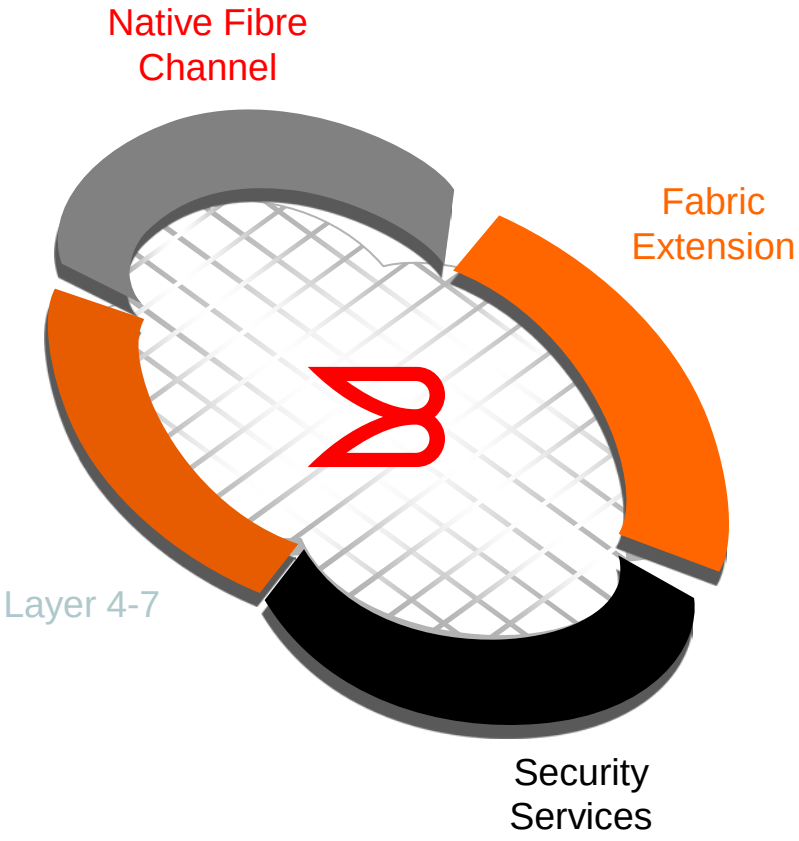

- Add services into the Ethernet fabric
  - Extends the capabilities of VCS
  - Fabric extension, native Fibre Channel, security services, layer 4-7, etc.
- Purpose-designed hardware
  - Switches with unique functionality can be added to the fabric
  - Like service modules in a chassis
  - Functionality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

SHARE in Anaheim 2011

Dynamic Services Details

Data Center to Data Center Connectivity

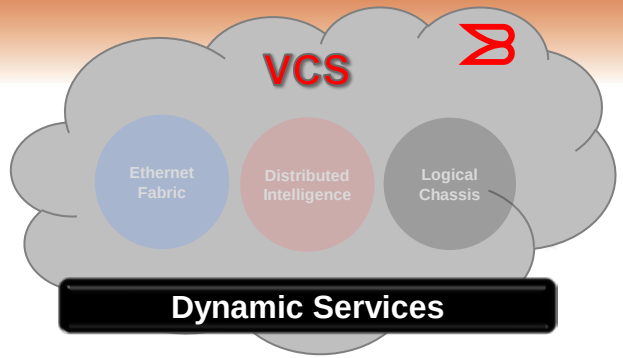

- Dynamic Service to connect Data Centers
  - Extend the layer 2 domain over distance
  - Maintains fabric separation while extending VCS services to secondary site (e.g. discovery, distributed configuration, AMPP)

- 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

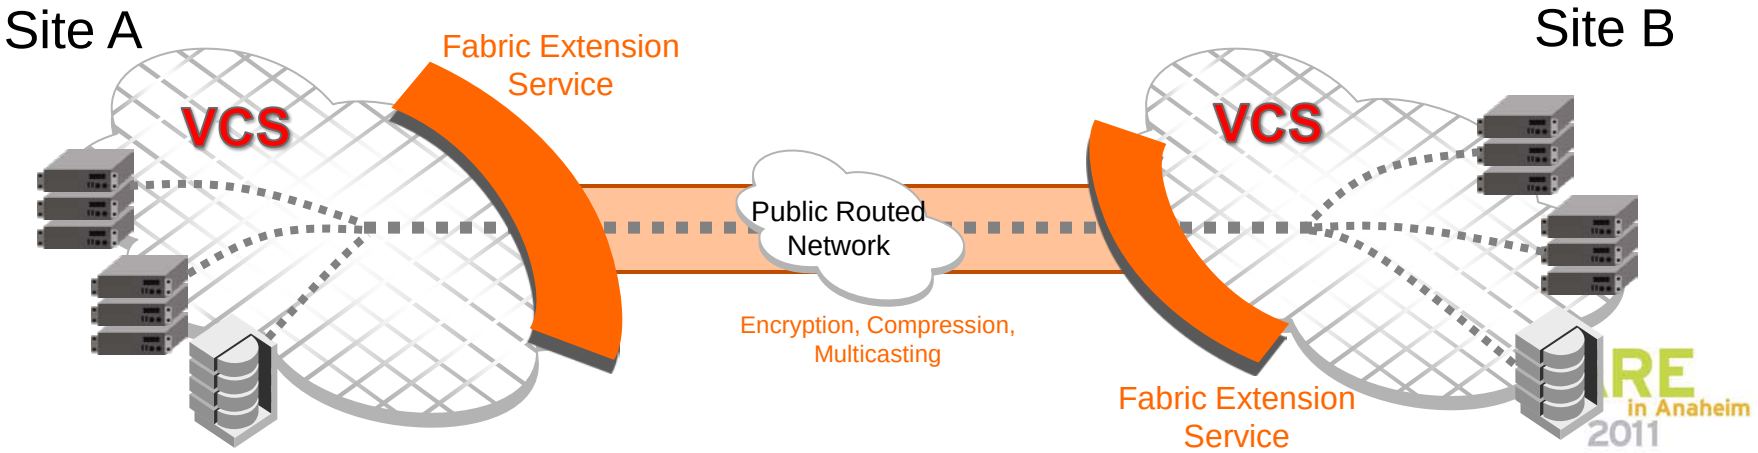

Dynamic Services Details

Native Fibre Channel Connectivity

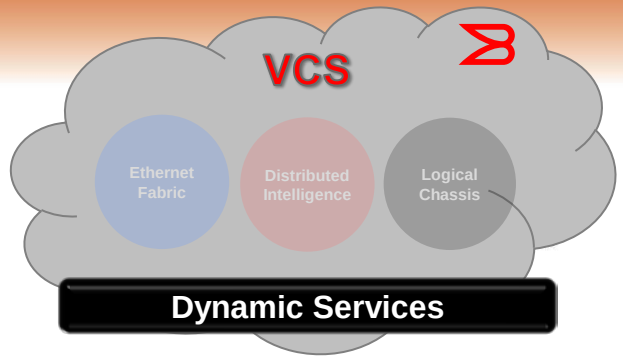

- Provide VCS Ethernet Fabric with native connectivity to FC storage
  - Connect FC storage locally
  - Leverage new or existing Fibre Channel SAN resources

- 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.

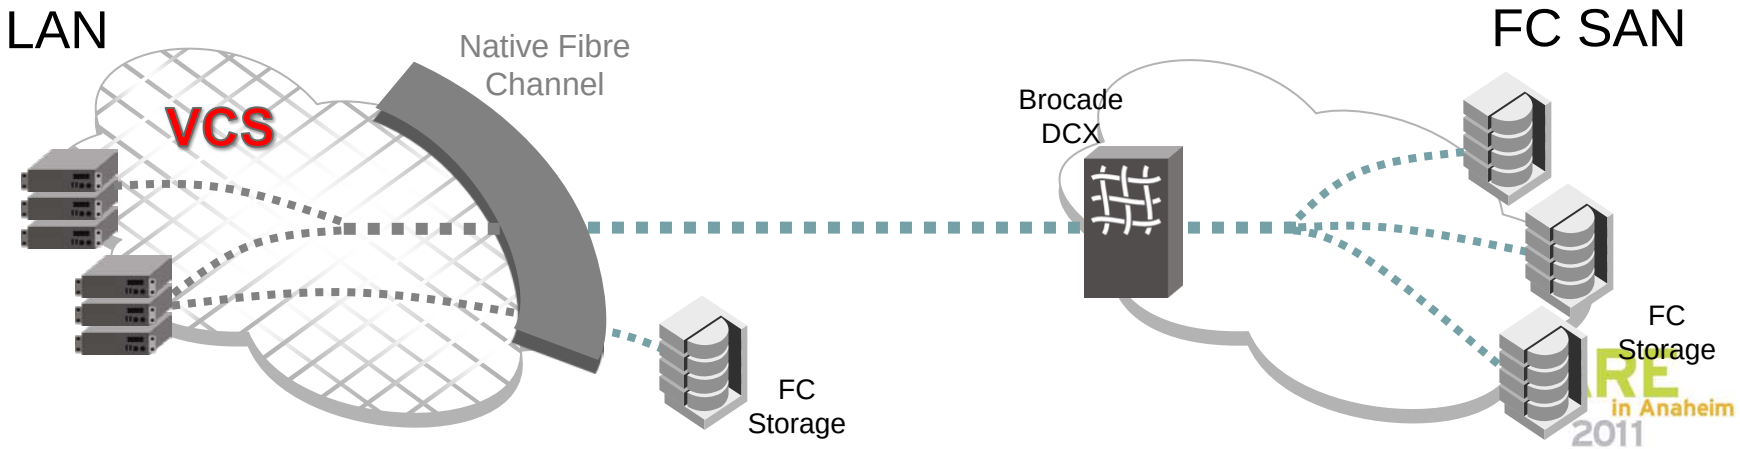

Dynamic Services for Network and Storage Convergence

Multi-hop FCoE + Bridging to Fibre Channel SAN

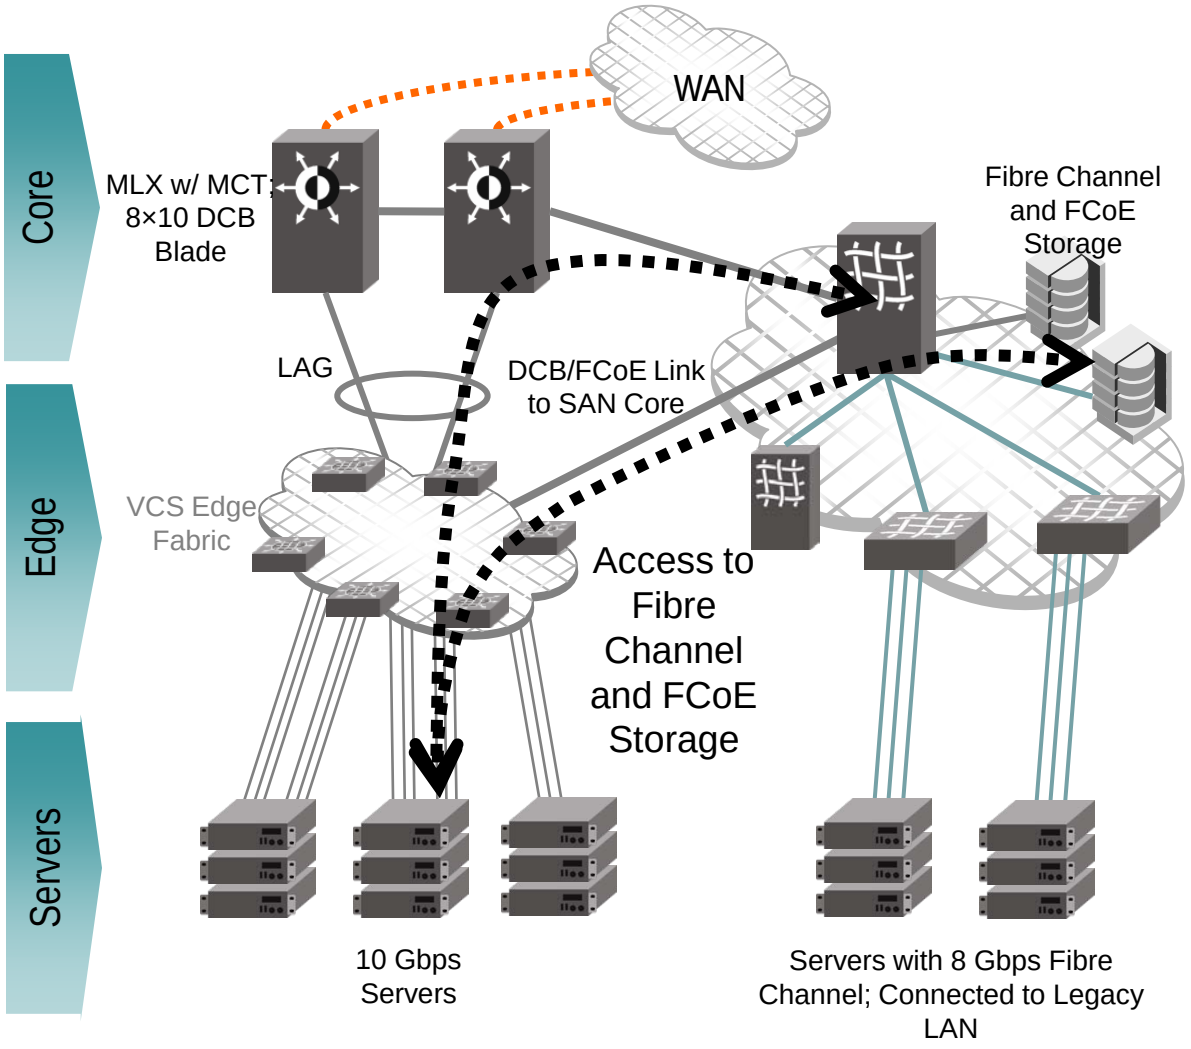

Leverage existing resources

- Connect VCS fabrics into Fibre Channel SAN; new servers can access existing storage

Maximum storage flexibility

- Fibre Channel, FCoE, iSCSI, NAS
- Deploy the right storage technology without isolating it

Optimal performance, availability

- No single point of failure
- Hardware-based trunking

VCS Use Cases

1. 1/10 Gbps Top-of-Rack Access
2. 10 Gbps Top-of-Rack Access for Blade Servers
3. 10 Gbps Aggregation; 1 Gbps Top-of-Rack Access
4. 1/10 Gbps Access; Collapsed Network Layers
   a. Top-of-Rack Mesh Topology
   b. Clos Fabric Topology
5. 1/10 Gbps Access; Network Convergence
6. 1/10 Gbps Access; Convergence + FC SAN

In the interest of time, I will cover 1,3,5,6

VCS Use Case #1

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

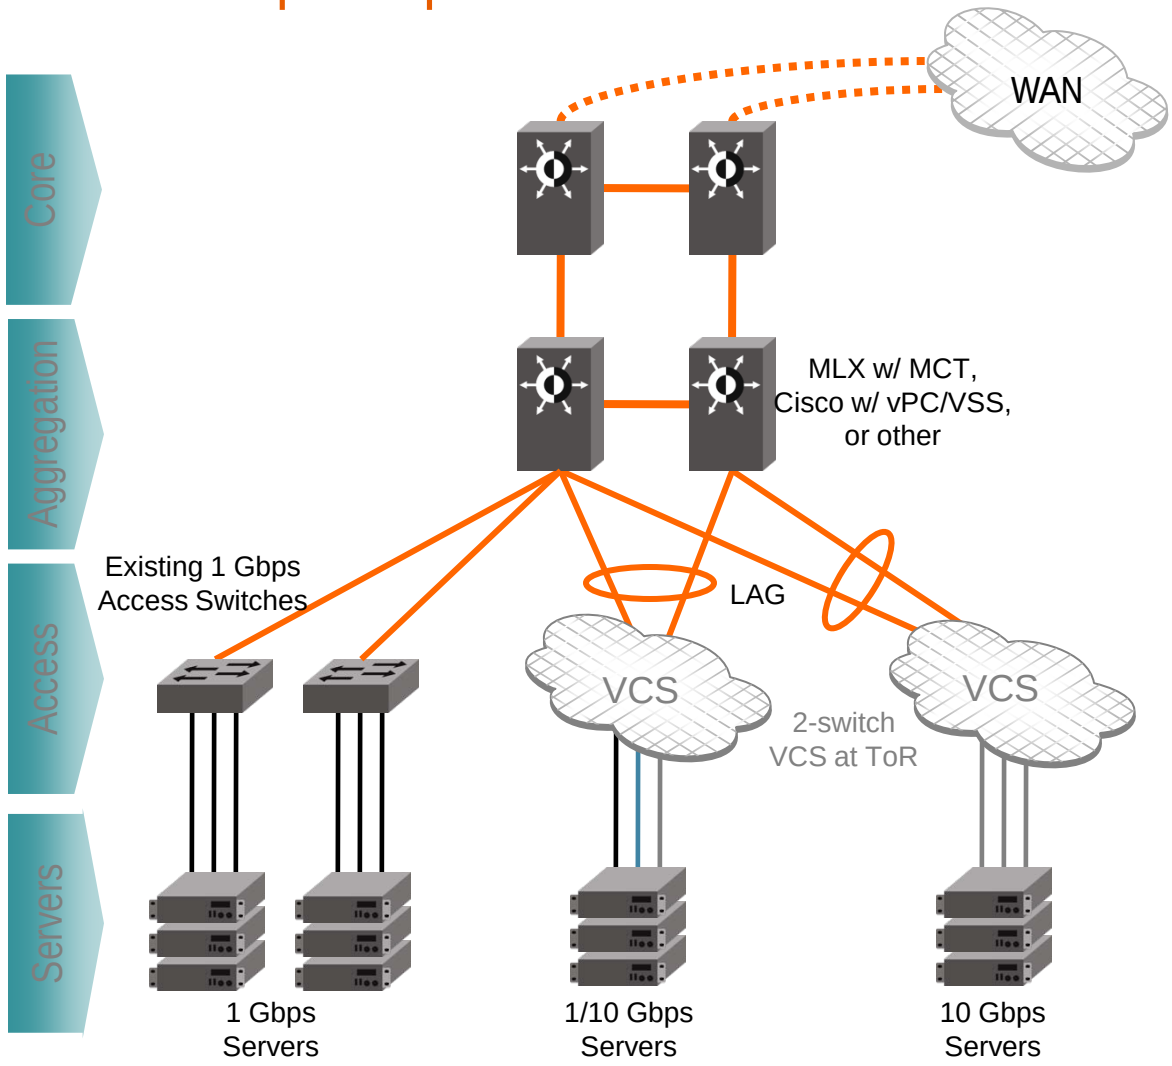

- Preserves existing architecture
  - Leverages existing core/agg
  - Co-exists with existing ToR switches
- Supports 1 and 10 Gbps server connectivity
- Active-active network
  - Load splits across connections
- No single point failure
  - Self healing
- Fast link reconvergence
  - < 250 milliseconds
- 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 – Topology

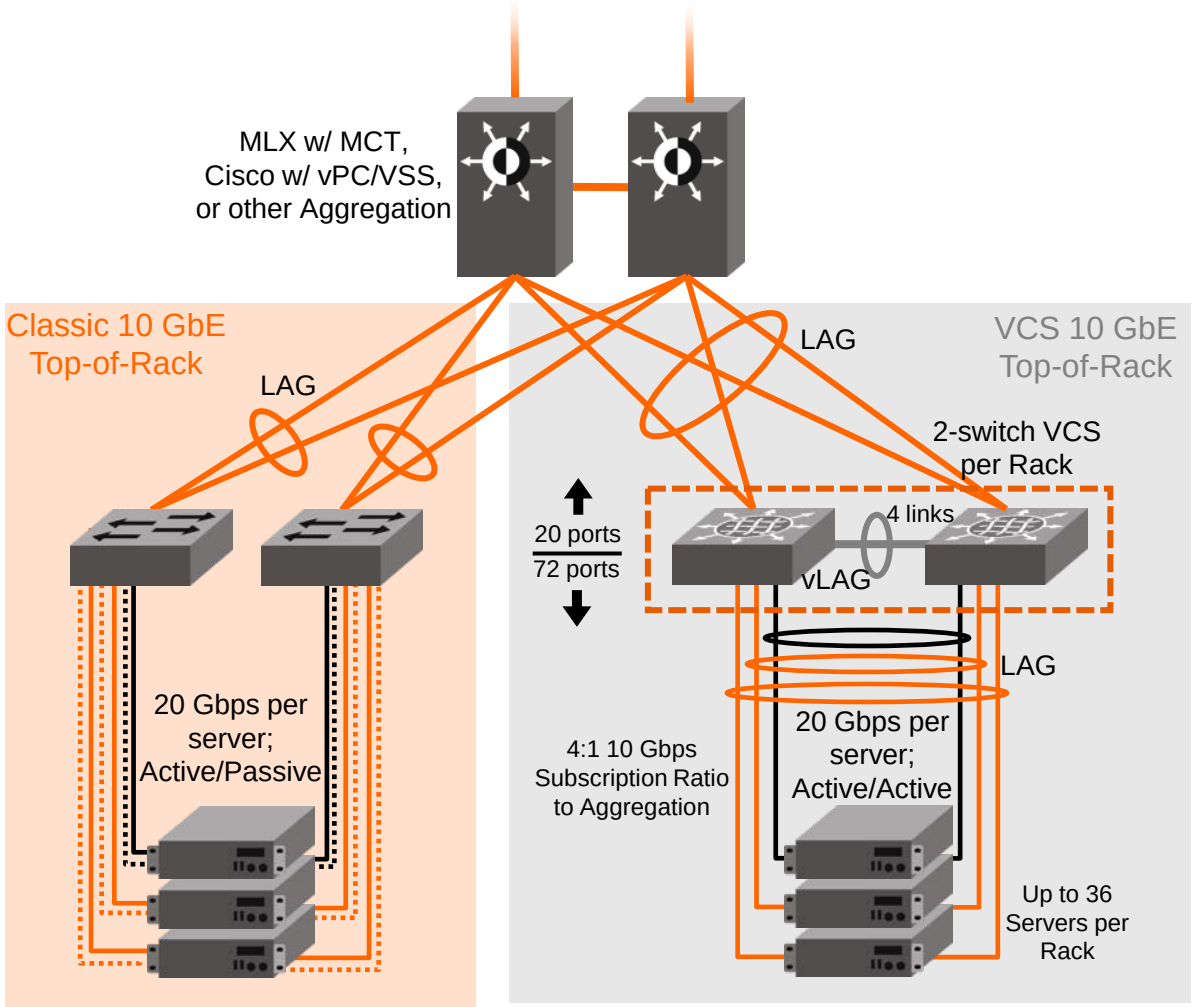

- Active/Active server connections
  - Servers only see one ToR switch
  - Half the server connections
- Reduced switch management
  - Half the number of logical switches to manage
- Unified uplinks
  - One LAG per VCS

	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

VCS Use Case #1

1/10 Gbps Top-of-Rack Access – Layout

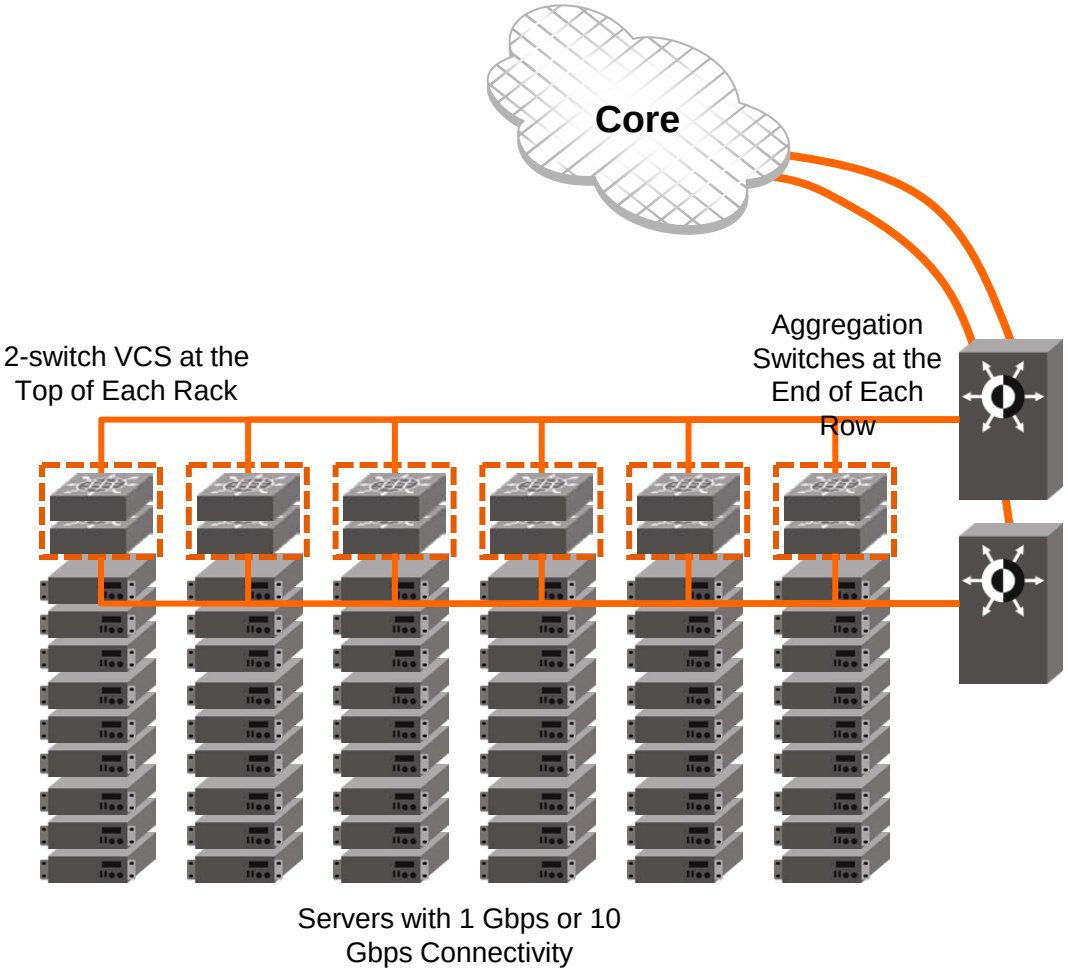

- Preserves existing network architecture
  - Leverage VCS technology in stages
- 2-switch VCS in each server rack
  - Managed as a single switch
  - 1 Gbps and 10 Gbps connectivity
  - Highly available; active/active
- High performance connectivity to End-of-Row Aggregation
  - One LAG to core for simplified management and rapid failover

VCS Use Case #3

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

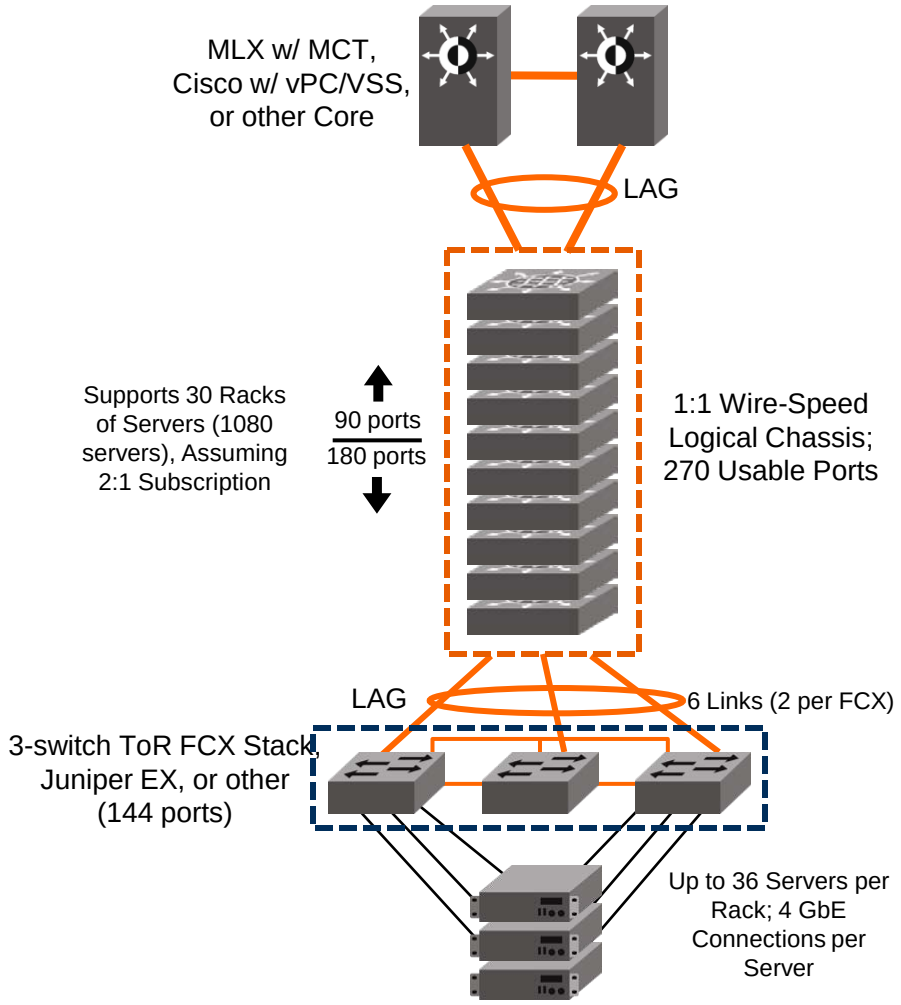

1 GbE | 10 GbE DCB | 10 GbE | Logical Chassis

Scalable VCS Aggregation

- Cost effective building blocks
- 270 usable ports with 1:1 subscription through VCS
- User-determined port count and subscription ratio

Aggregates 1 GbE Access

- 3-switch stack in each server rack
- LAG across stack members to VCS
- Reduced management; no single point of failure

VCS Use Case #3

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

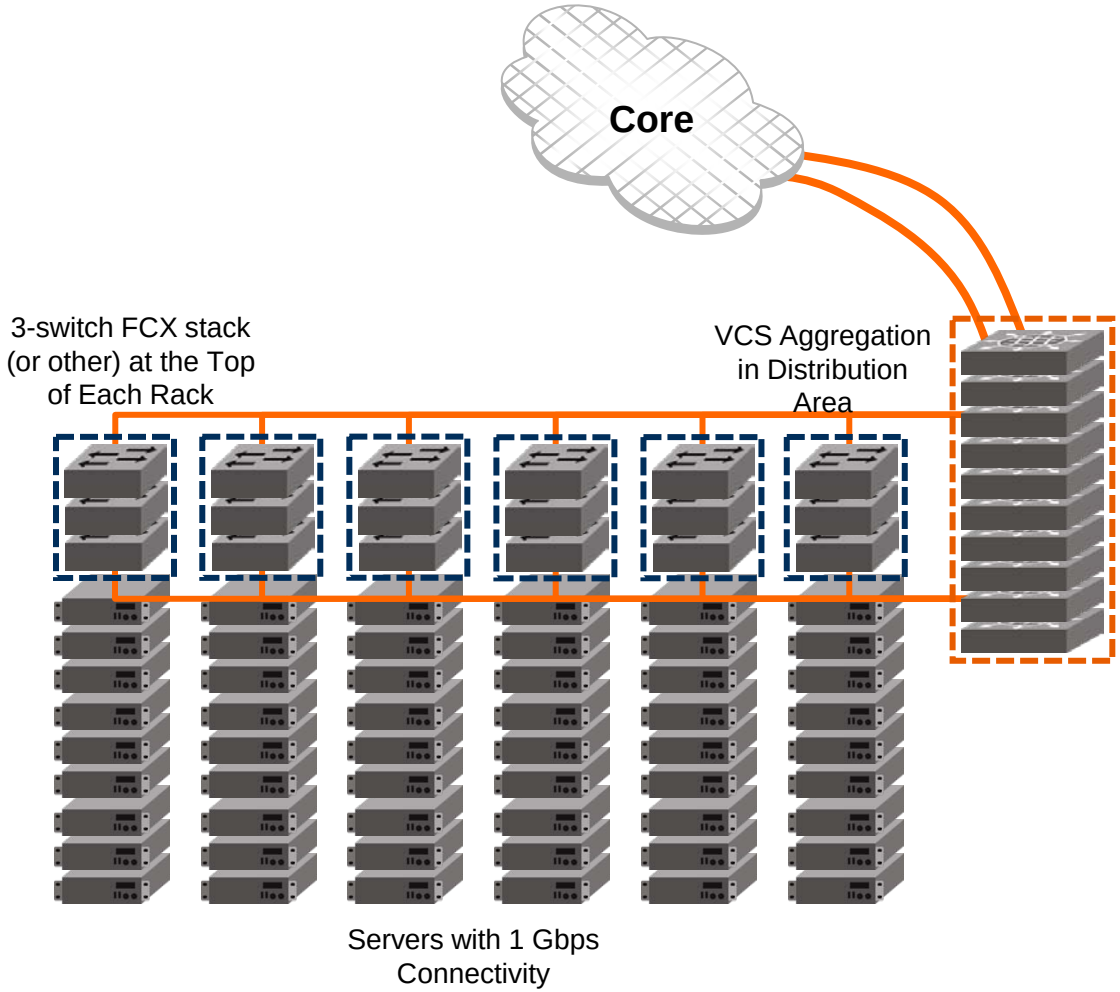

- 3-switch stack in each rack
  - Managed as a single switch
  - Redundancy throughout network, without STP
  - High density 10 Gbps LAG to VCS aggregation
- Logical Chassis Aggregation Router in Distribution Area
  - Build out aggregation as needed
  - Supports 30 racks of servers
- High performance, resilient connection to Core
  - One LAG for simplified management and rapid failover

VCS Use Case #5

1/10 Gbps Access; Network Convergence – Architecture

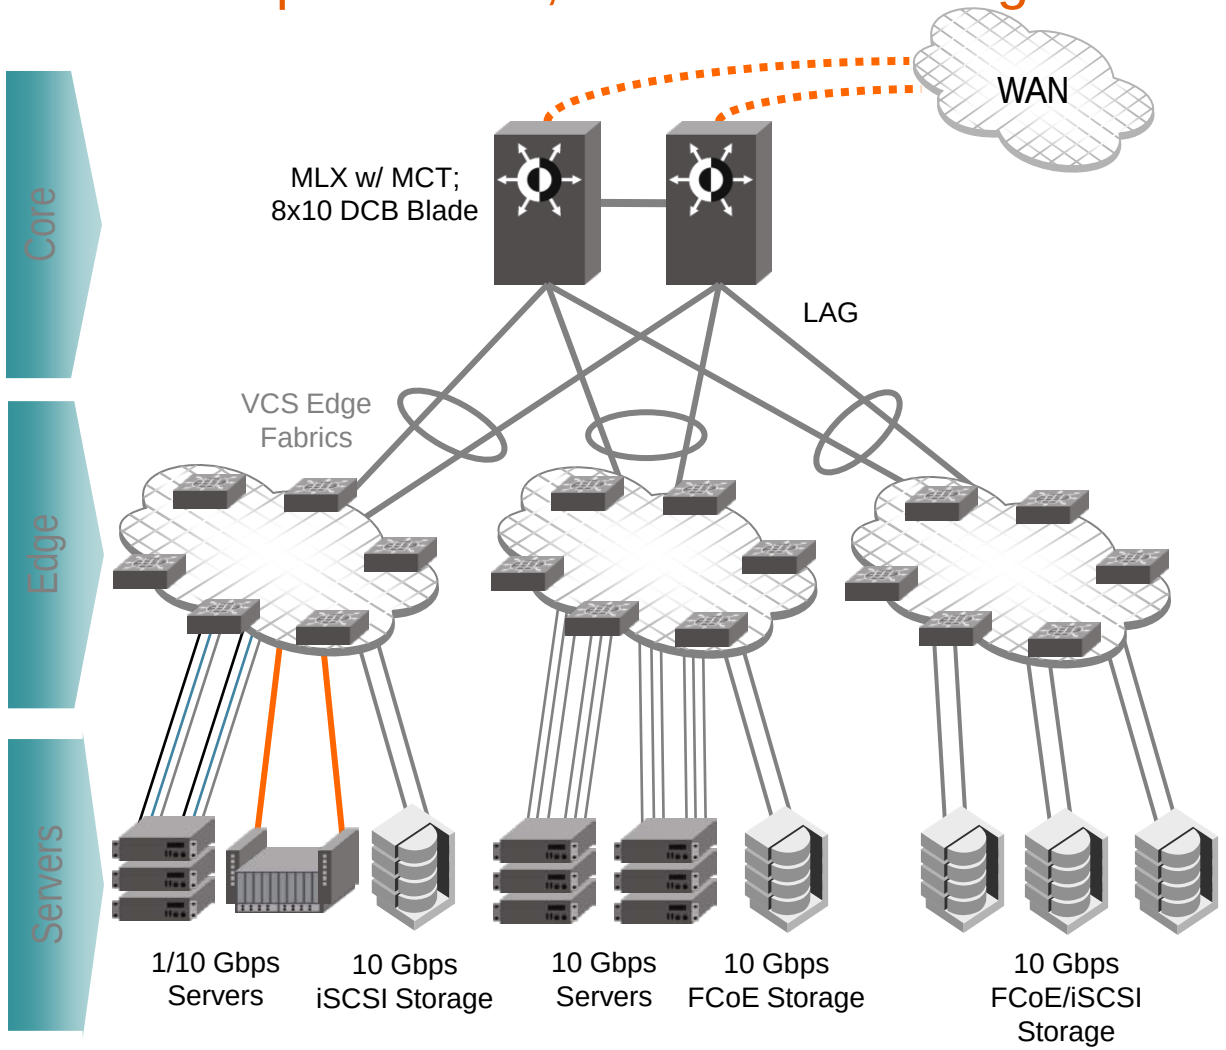

- Flatter, simpler network design
  - Logical two-tier architecture
  - VCS fabrics at the edge
- Greater layer 2 scalability/flexibility
  - Increased sphere of VM mobility
  - Seamless network expansion
- Optimized multi-path network
  - All paths are active
  - No single point failure
  - STP not necessary
- Convergence ready
  - End-to-end enhanced Ethernet (DCB)
  - Multi-hop FCoE support
  - Lossless iSCSI

VCS Use Case #5

1/10 Gbps Access; Network Convergence – Topology

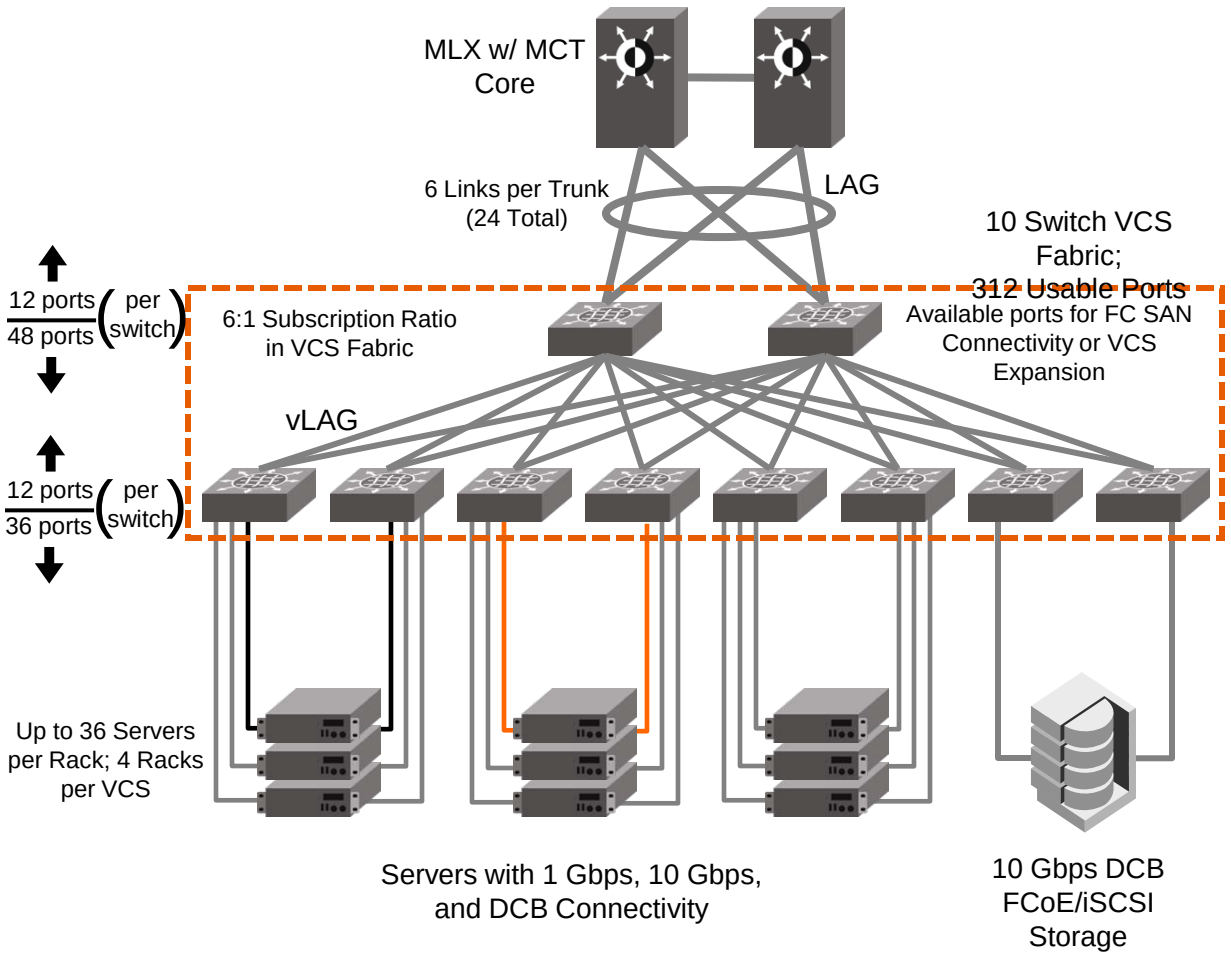

- Scale-out VCS edge fabric
  - Self aggregating, flattens the network
  - Clos Fabric topology for flexible subscription ratios
  - 312 usable ports per 10-switch VCS
  - Supports 144 servers in 4 racks, all with 10 Gbps connections
- Drastic reduction in management
  - Each VCS managed as a single logical chassis
- Enables network convergence
  - DCB and TRILL capabilities for multi-hop FCoE and enhanced iSCSI

VCS Use Case #5

1/10 Gbps Access; Network Convergence – Layout

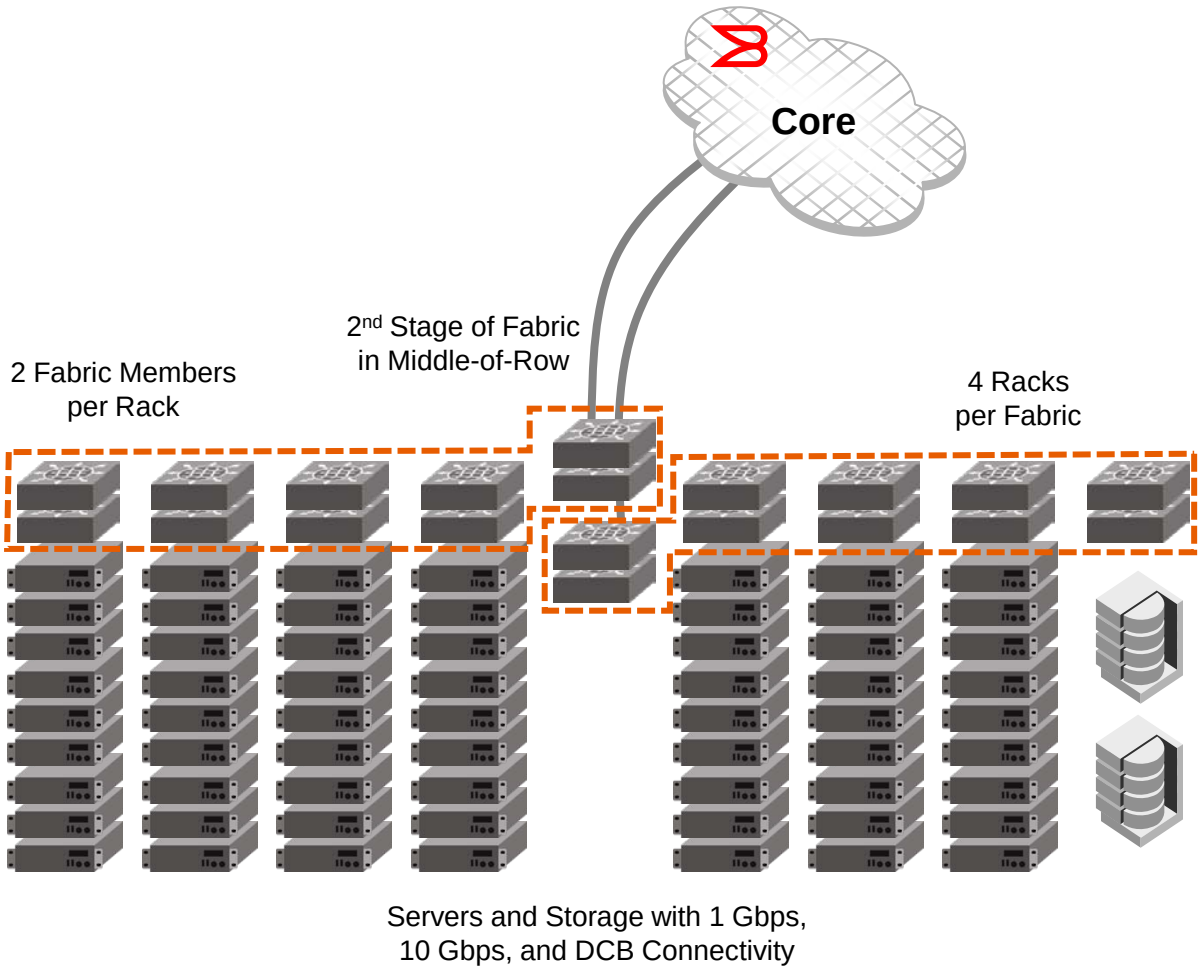

- 2 fabric members in each rack
  - Dual connectivity into fabric for each server/storage array
  - Low cost Twinax cabling in rack
- 2nd stage fabric members in a middle-of-row rack
  - Low cost Laserwire cabling from top-of-rack switches
  - 1 VCS fabric per 4 racks of servers (assuming 36 servers per rack)
- Fiber optic cabling only used for connectivity from edge VCS to core
  - Single LAG per fabric
  - Reduced management and maximum resiliency

VCS Use Case #6

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

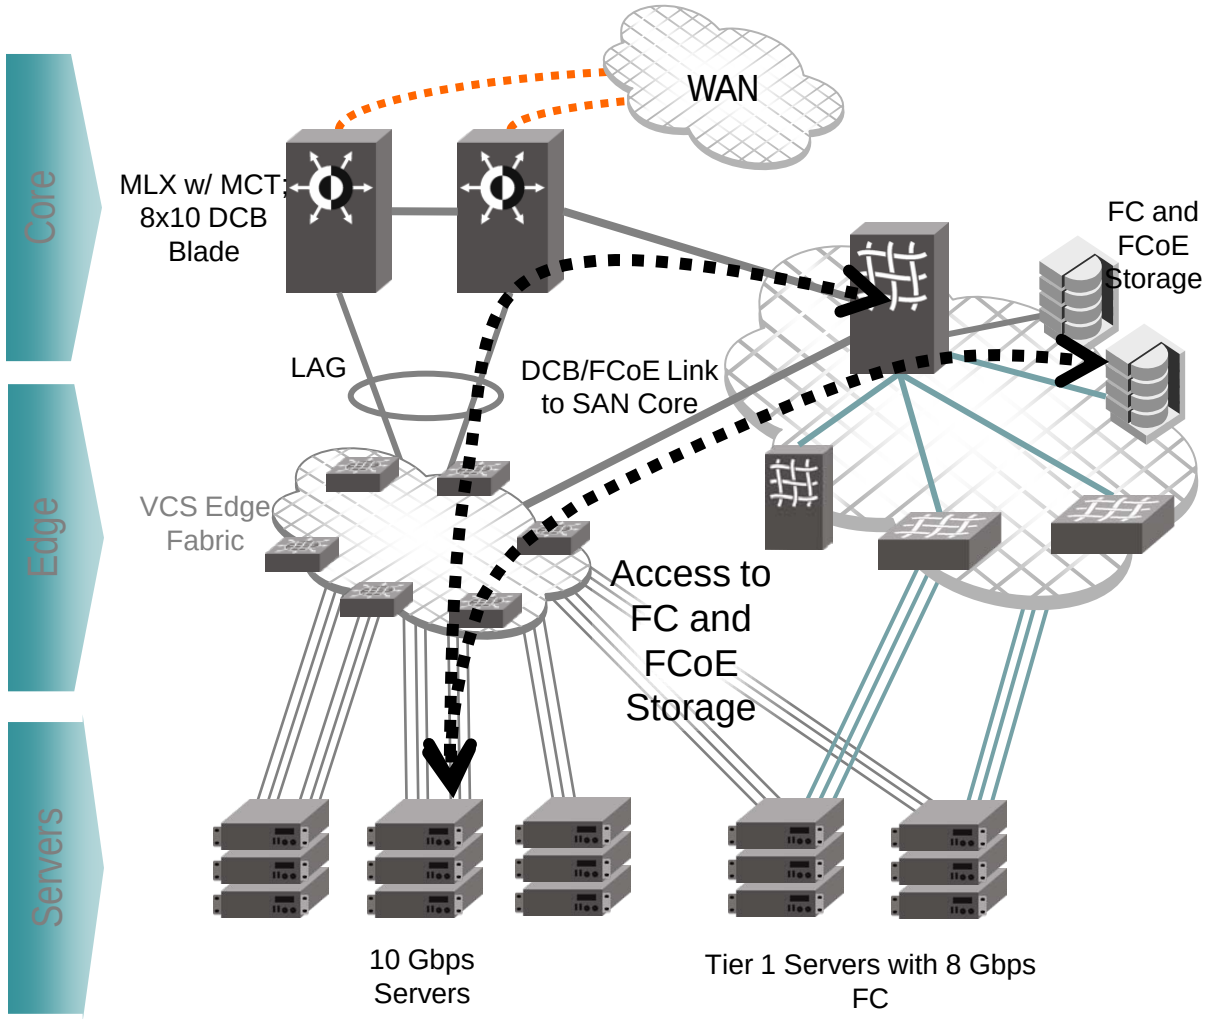

- Leverage existing resources
  - Connect Ethernet fabrics into Fibre Channel SAN – new servers have access to existing storage
- Maximum storage flexibility
  - Fibre Channel, FCoE, iSCSI, NAS
  - Deploy the right storage technology without isolating it
- Optimal performance, availability
  - No single point failure
  - Frame-level, hardware-based trunking between nodes

VCS Use Case #6

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

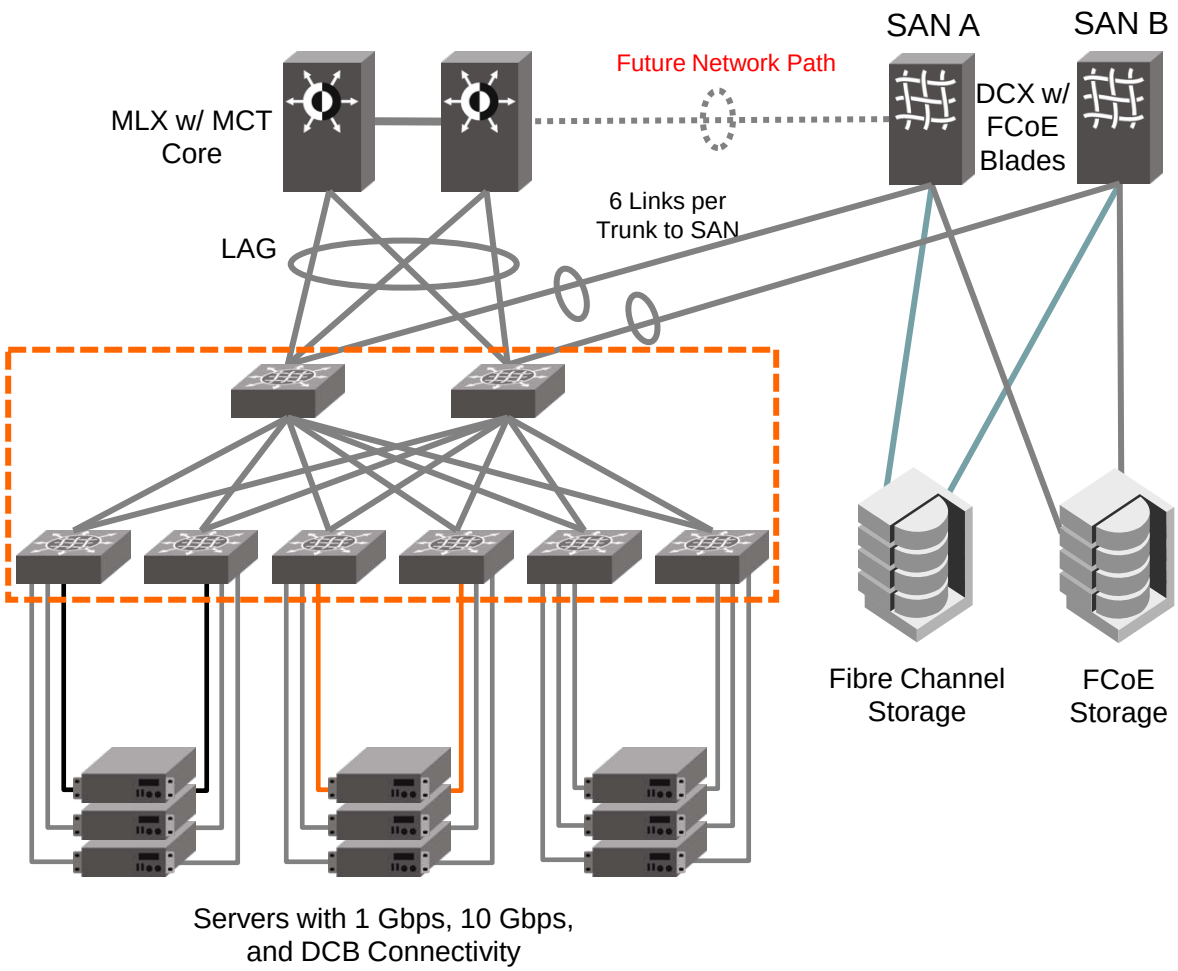

- VCS fabric connectivity into Fibre Channel SAN
  - High performance Ethernet trunks from VCS to DCX core
- Allows shared storage resources to exist in SAN
  - Fibre Channel and FCoE storage
  - Can be accessed by servers with Converged Network Adapters
- Future connectivity from converged LAN aggregation to SAN core
  - MLX with DCB connects to DCX with FCoE blade

What about VCS with?

- System z?
- zBX?
- zEnterprise?

Summary: Virtual Cluster Switching (VCS)

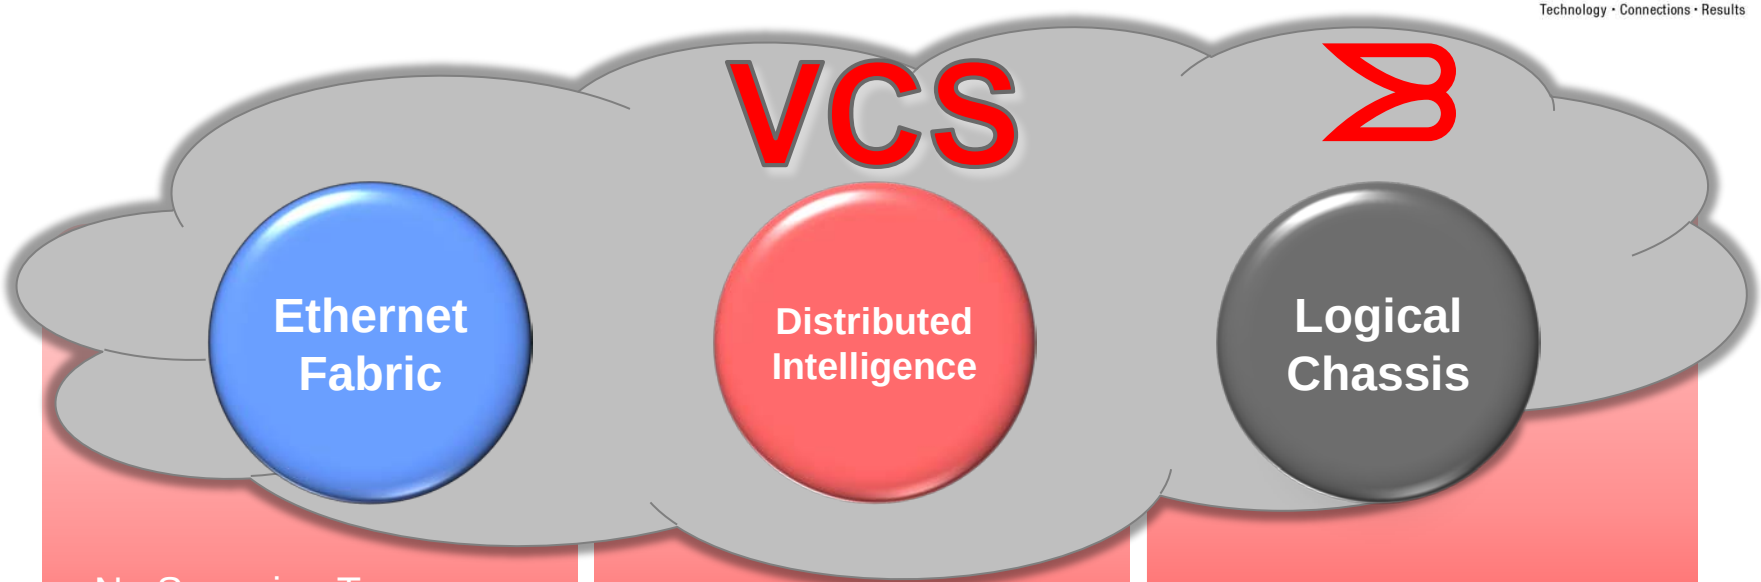

- No Spanning Tree Protocol
- Multi-path, deterministic
- Auto-healing, non-disruptive
- Lossless, low latency
- Convergence-ready

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

Both sides now

- **Rows and flows of angel hair,**
  And ice cream castles in the air,
  And feather canyons everywhere,
  I've looked at clouds that way. But now they only block the Sun,
  They rain and snow on everyone.
  So many things I would have done,
  But clouds got in my way.
- **I've looked at clouds from both sides now,**
  From up and down, and still somehow,
  It's cloud illusions I recall,
  I really don't know clouds, at all.
- Hopefully now, after the week at SHARE, you know clouds a little better!

THANK YOU