BROCADE

THE VIRTUALIZED SP DATA CENTER

Joerg Ammon (jammon@brocade.com)

Systems Engineer Service Provider

Legal Disclaimer

- All or some of the products detailed in this presentation may still be under development and certain specifications, including but not limited to, release dates, prices, and product features, may change. The products may not function as intended and a production version of the products may never be released. Even if a production version is released, it may be materially different from the pre-release version discussed in this presentation.
- NOTHING IN THIS PRESENTATION SHALL BE DEEMED TO CREATE A WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE, INCLUDING BUT NOT LIMITED TO, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT OF THIRD-PARTY RIGHTS WITH RESPECT TO ANY PRODUCTS AND SERVICES REFERENCED HEREIN.
- Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, IronView, NetIron, SAN Health, ServerIron, and TurboIron are registered trademarks, and Brocade Assurance, DCFM, Extraordinary Networks, and Brocade NET Health are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Agenda

- Trends in the SP Data Center
- Virtualizing the server
- Virtualizing the network using MPLS
- Virtualizing the network using TRILL
- Virtualizing the IO
- Summary

Trends in the Data Center

Capitalize—Stay Ahead

Data Center

- Virtualization (→ Cloud)
- Consolidation
- Convergence ready
- Environmentals
- Gigabit ramp

Data Center Evolution

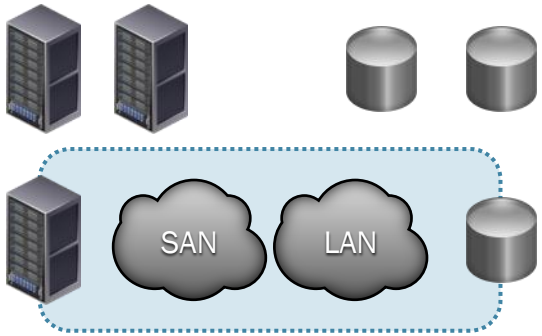

ENTERPRISE

HISTORICAL

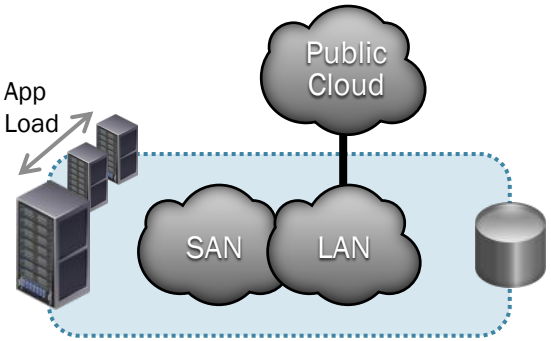

DYNAMIC NETWORKS

TODAY

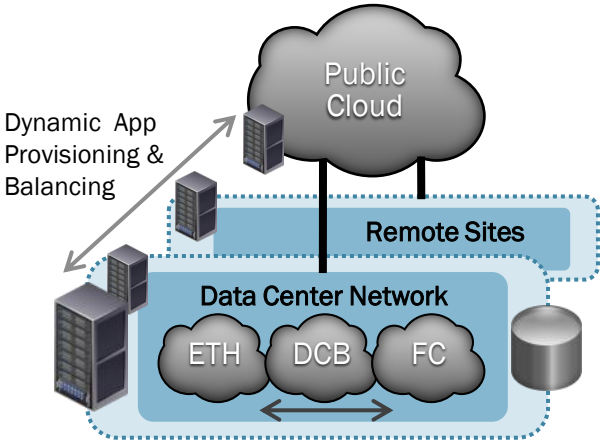

PRIVATE CLOUD

FUTURE

The Network is Central to the Cloud

Virtualization Brings New Demands

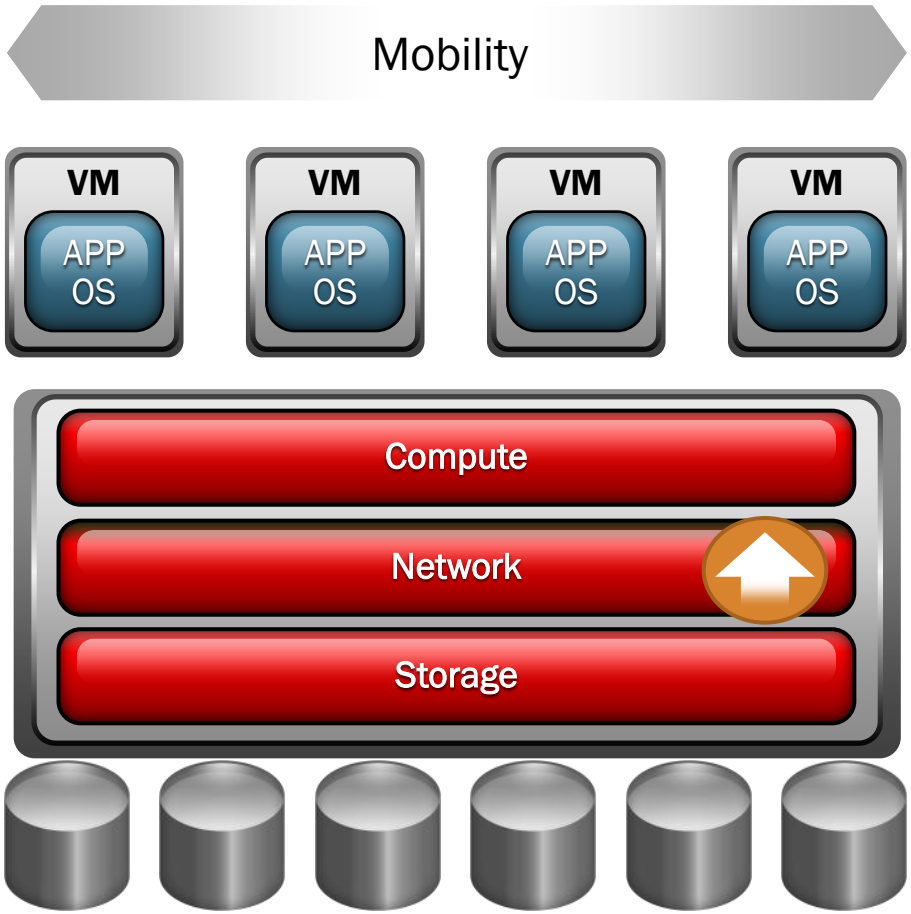

Challenges

Network performance/scalability constraints

Application resiliency and performance under load

VM mobility limits

Infrastructure complexity

Management silos

Building the Virtualized Data Center

Brocade Advantage

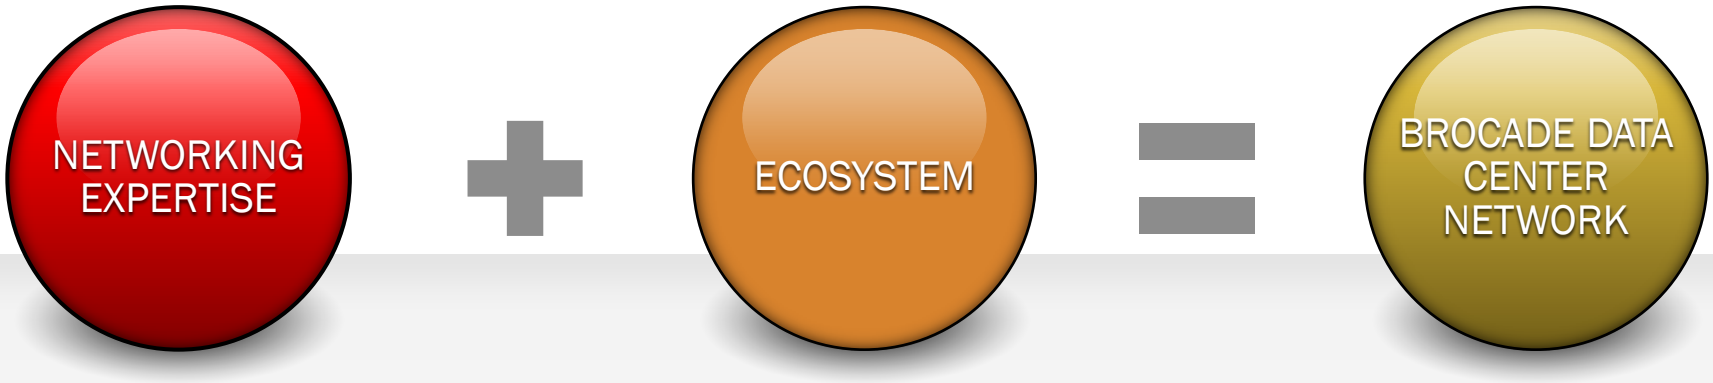

LAN, SAN and DCB; App awareness

Industry leading performance, scalability, reliability and energy efficiency

Professional Services

Strong server and storage industry partnerships

Broad interoperability

Standards-based

Best of Breed solutions

Choice and Agility

Leverage investments

Maximum value

Virtualizing the server

Components of Virtual Networking

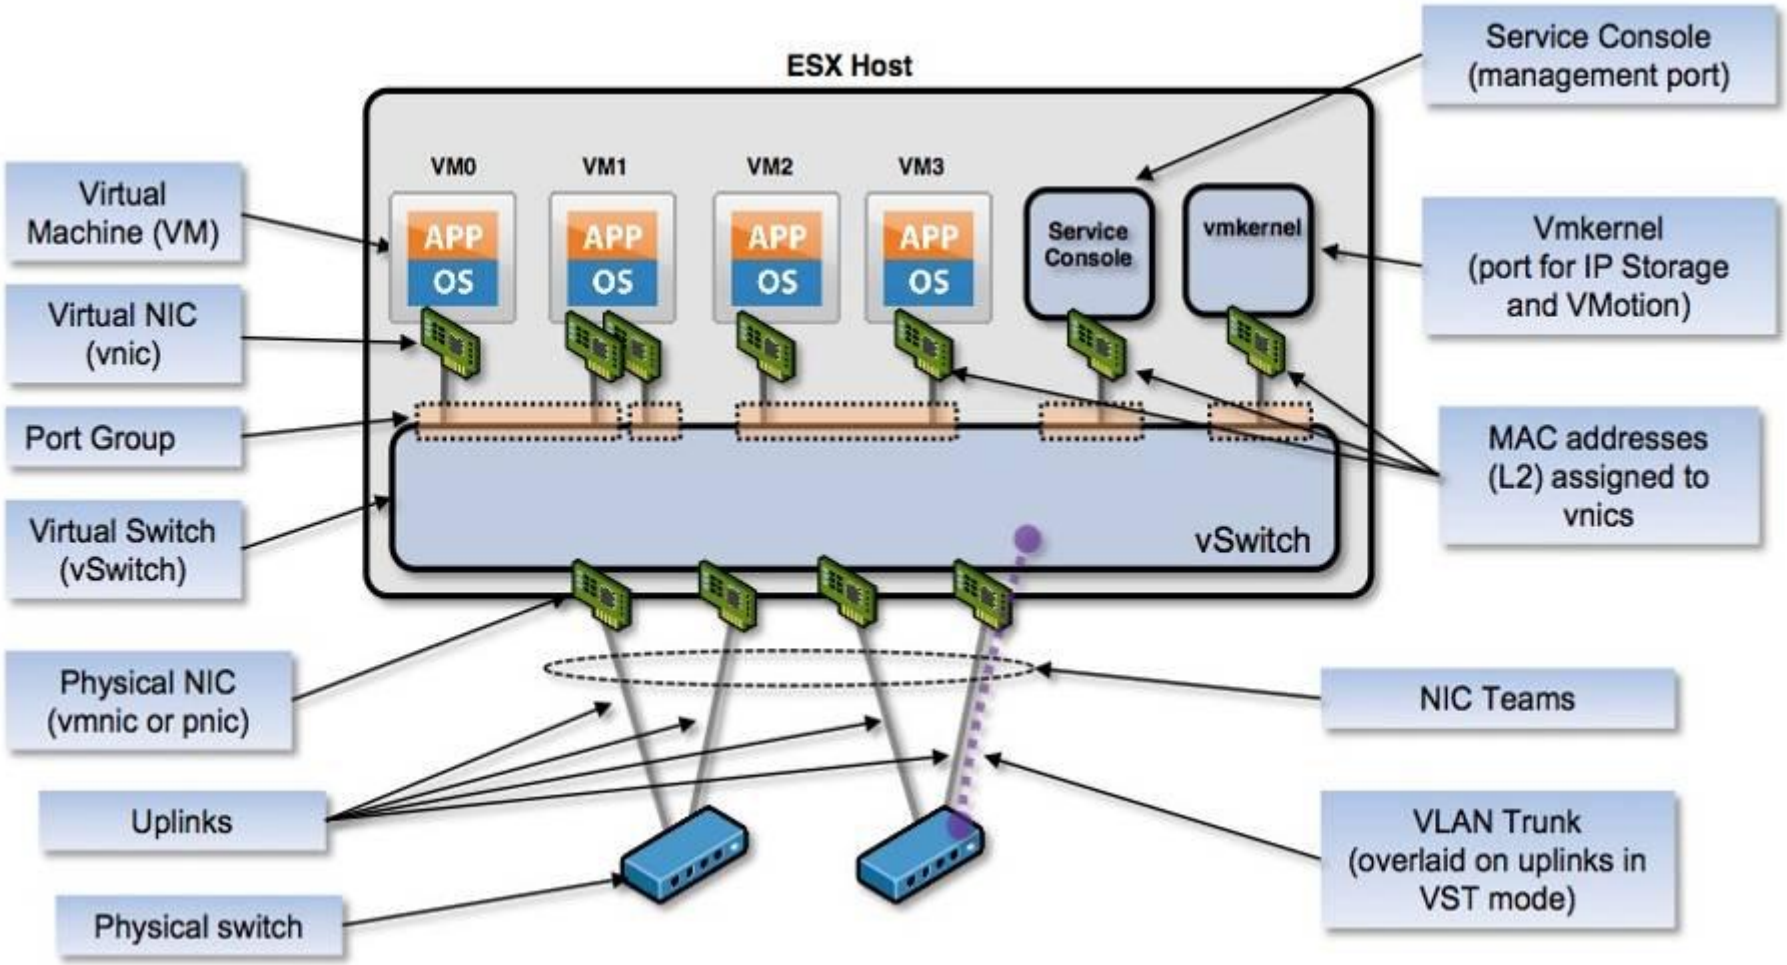

Components of Virtual Networking

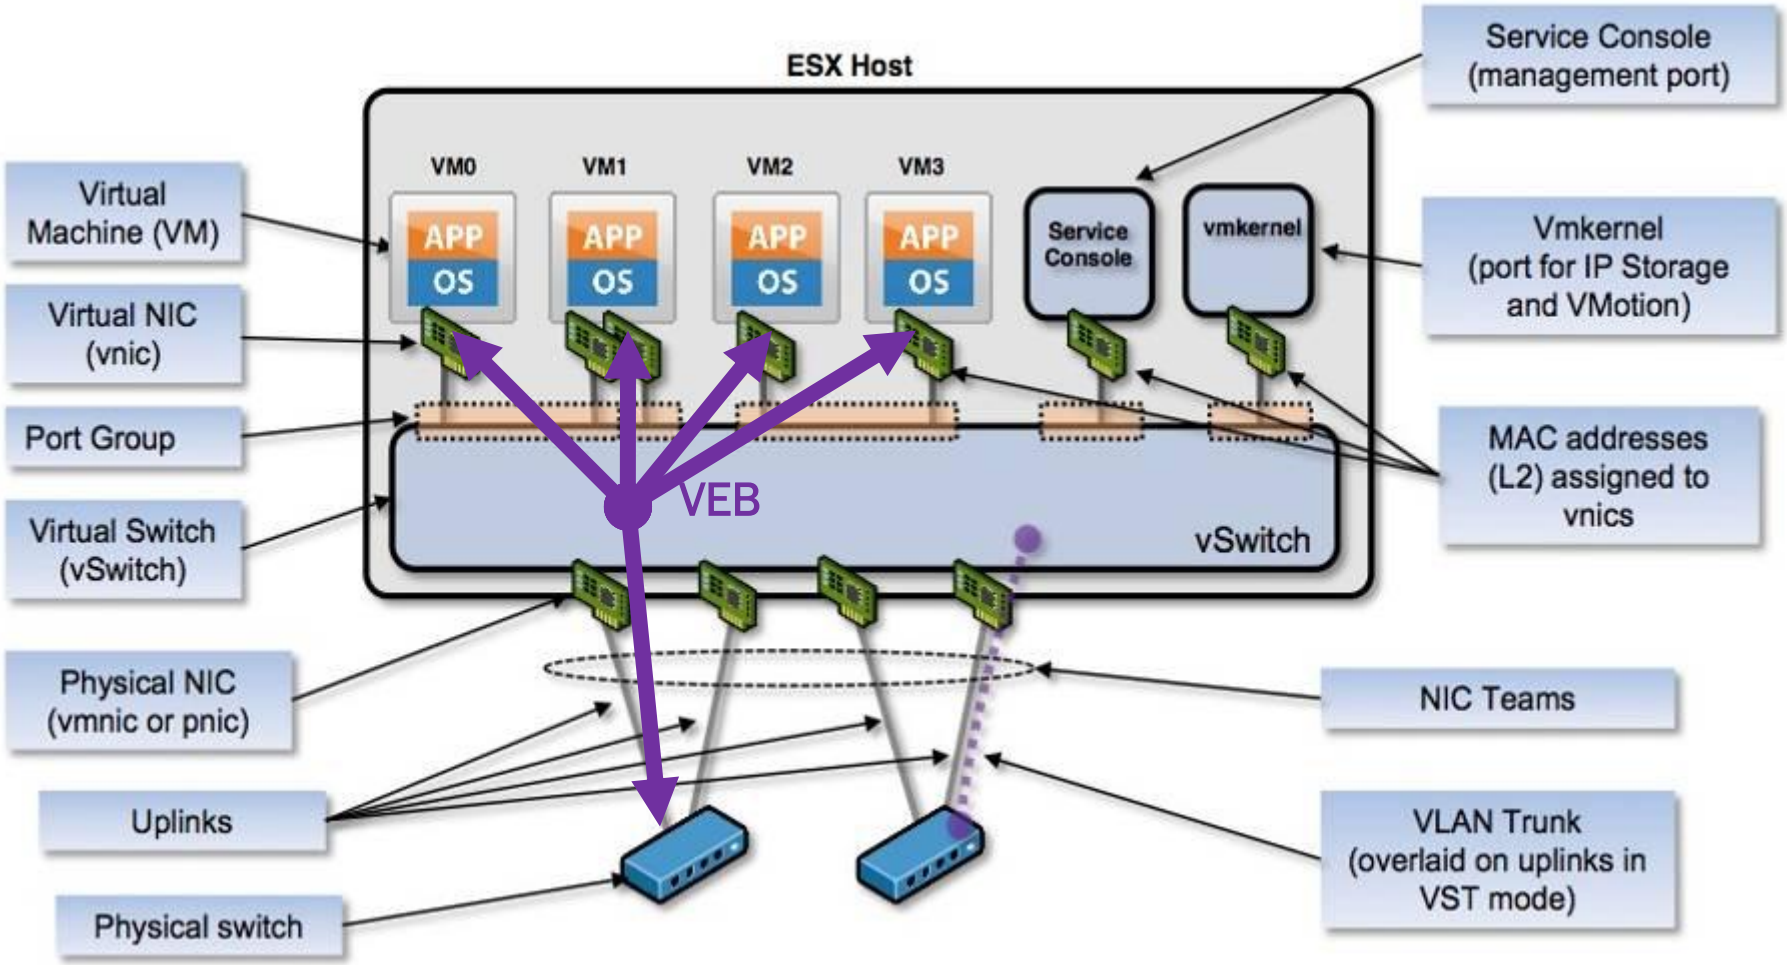

Up to 20% resource overhead if implemented as virtual switch

Components of Virtual Networking

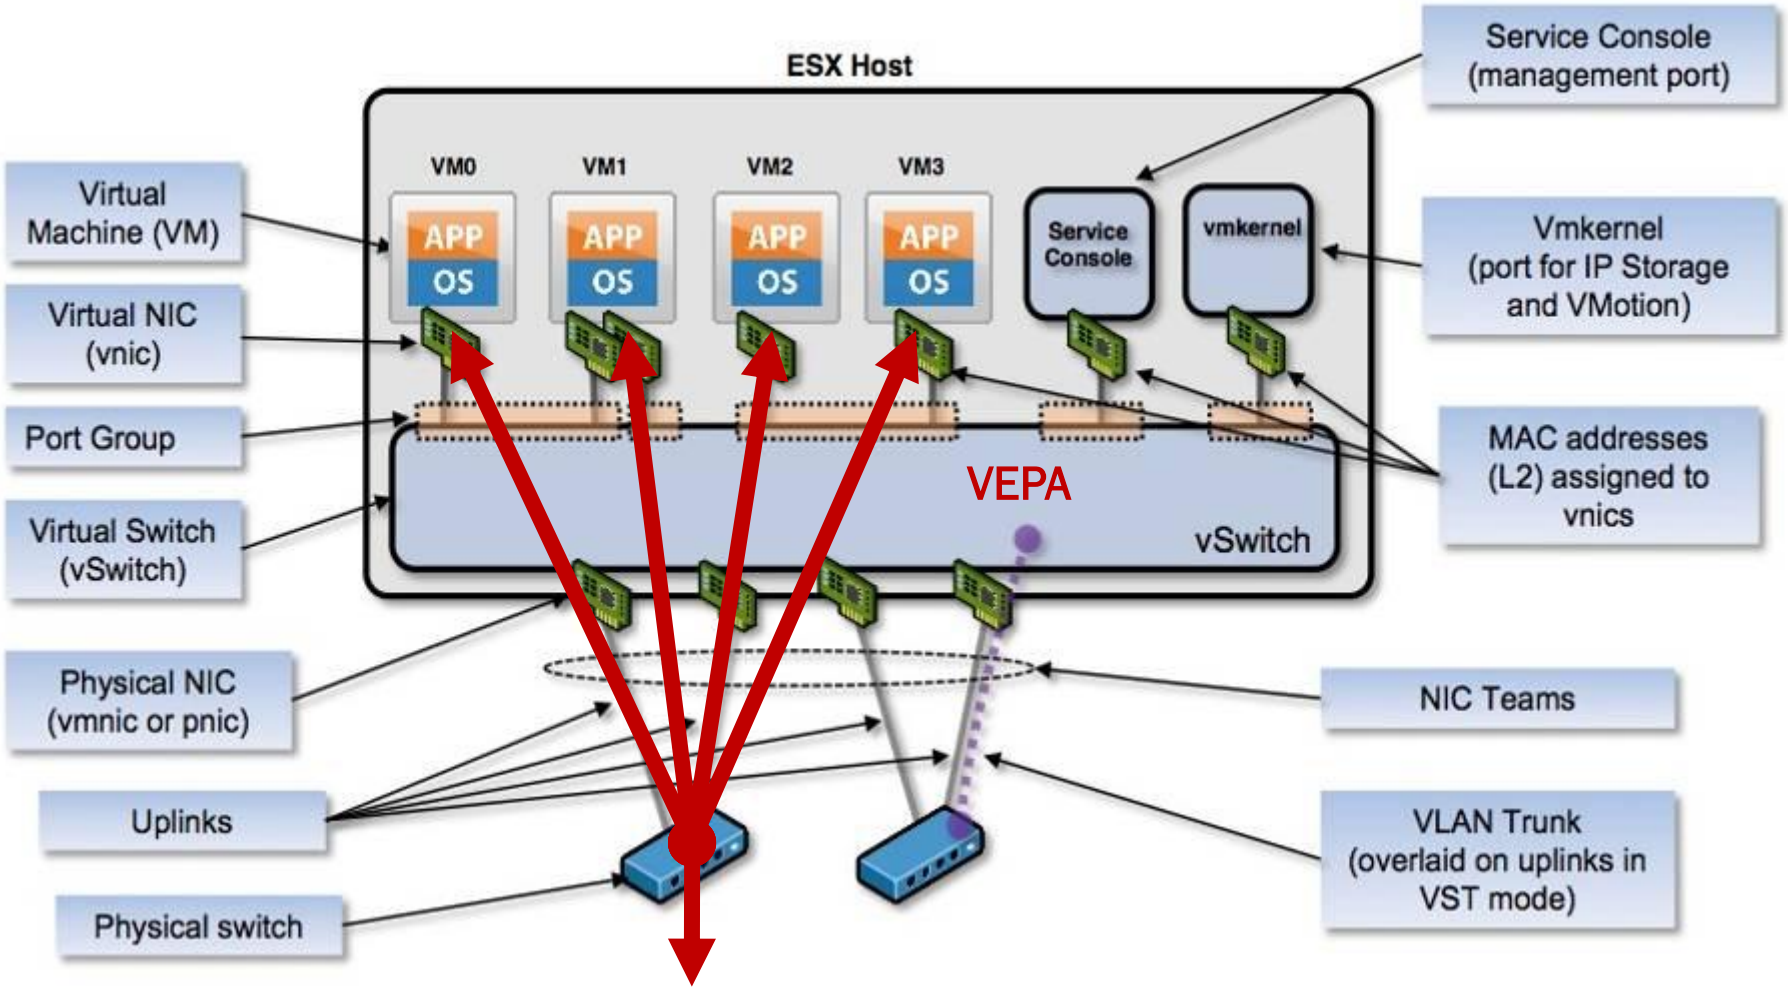

Hairpin-switching requires modification to „standard Ethernet“

Requirements for a virtual server move

- Similar environment: CPU, no direct attached devices
- Server needs access to his storage device
- Server needs the same network environment after move
- Large flat L2-network

LANs Get L2 Multipathing

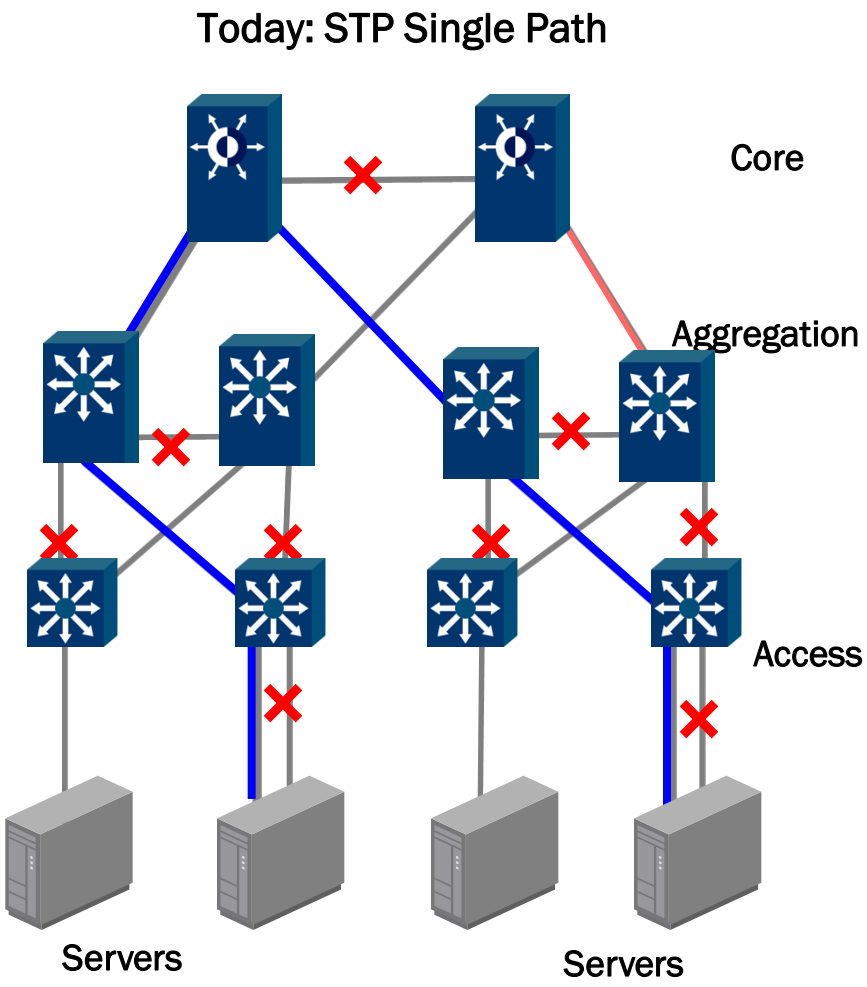

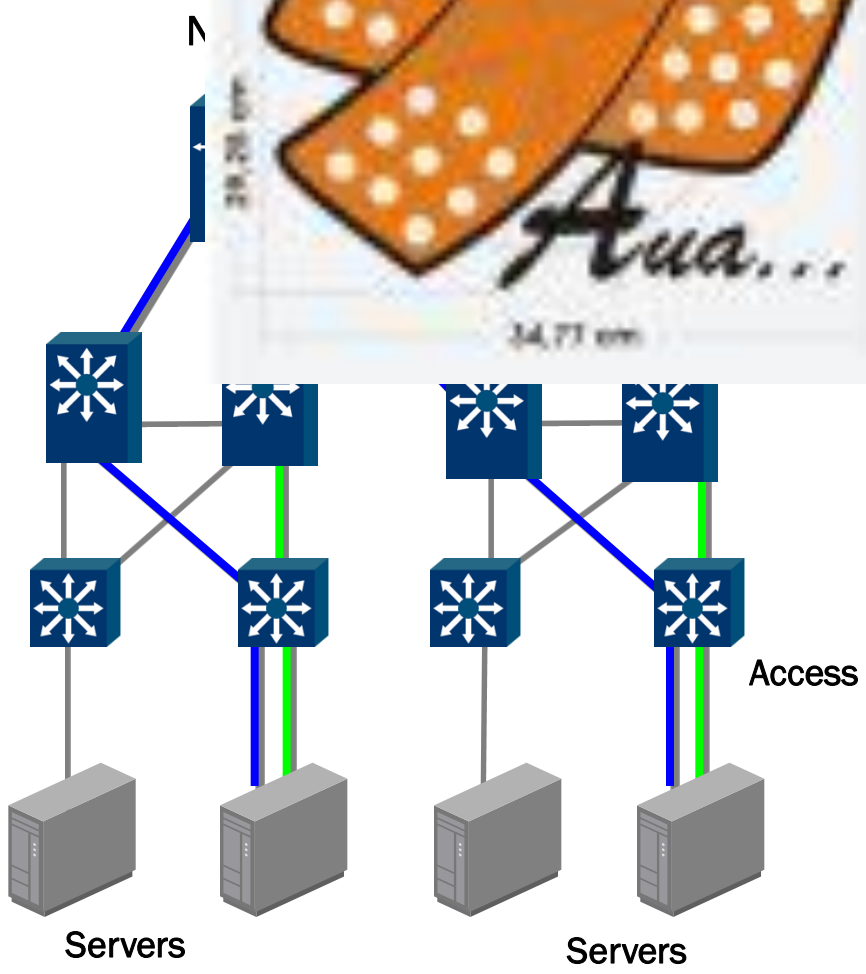

Virtualizing the network (and the service) using MPLS

Real world Data Center Challenge

Security-As-A-Service Company

Requirements:

- Data Center consolidation
- Green Design – reduce space, power and cooling
- Investment Protection with any new hardware
- Technology to support geographically dispersed data centers
- Need for Virtualization and Isolation
- Support growth of existing and additional systems
- Mission critical High Speed transport and access
- High Availability – redundancy at every level and path
- Security - appropriate for customer facing applications
- Operational Efficiency + “Lights-out” remote management

Example Data Center Design

MPLS-Enabled Data Center

GbE
10 GbE

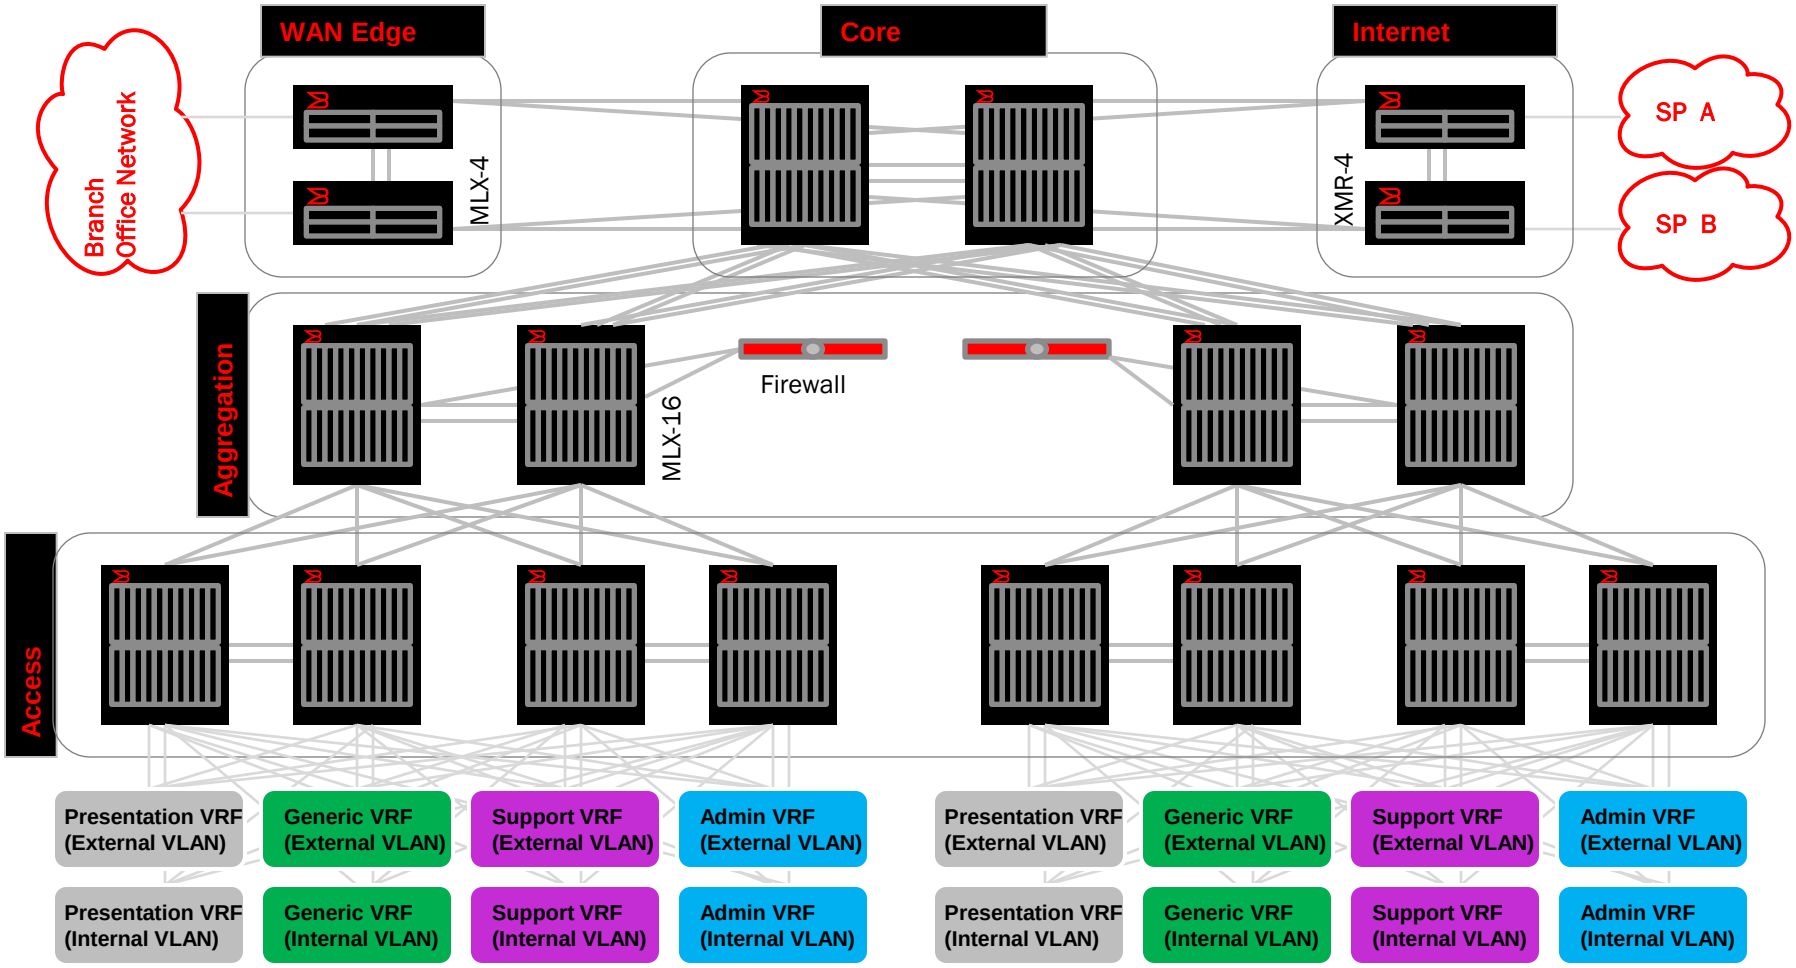

Example Data Center Design

MPLS-Enabled Data Center

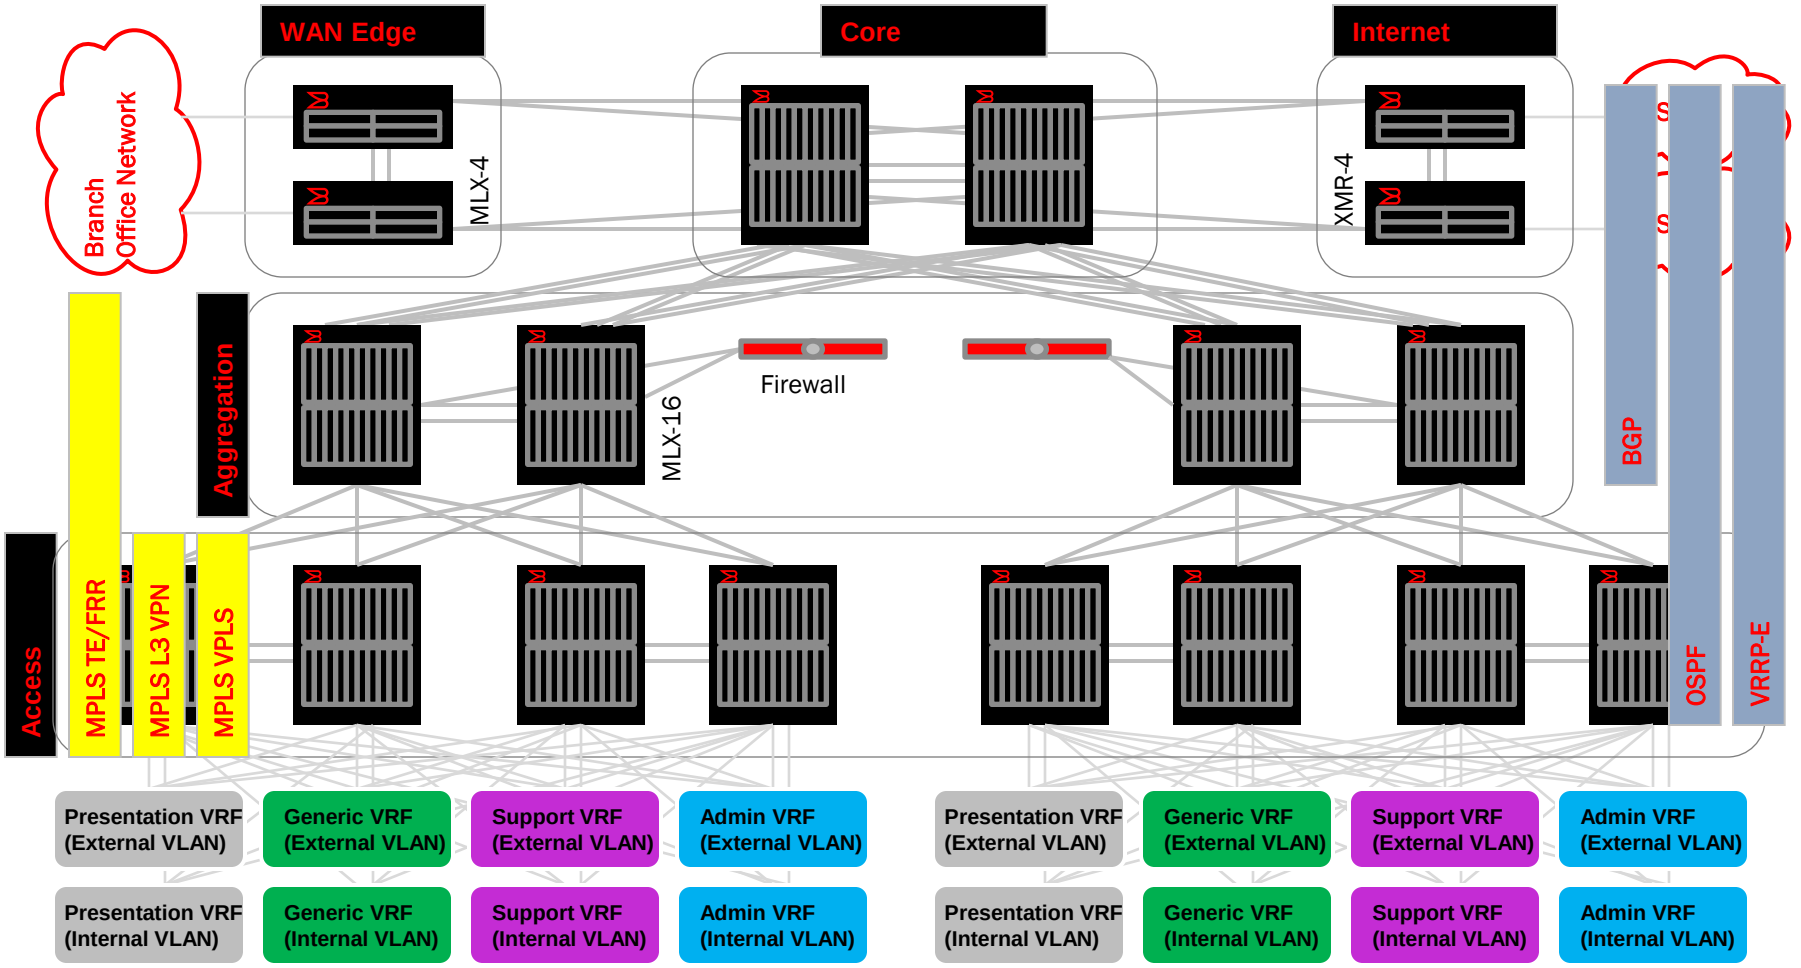

Example Data Center Design

MPLS-Enabled Data Center

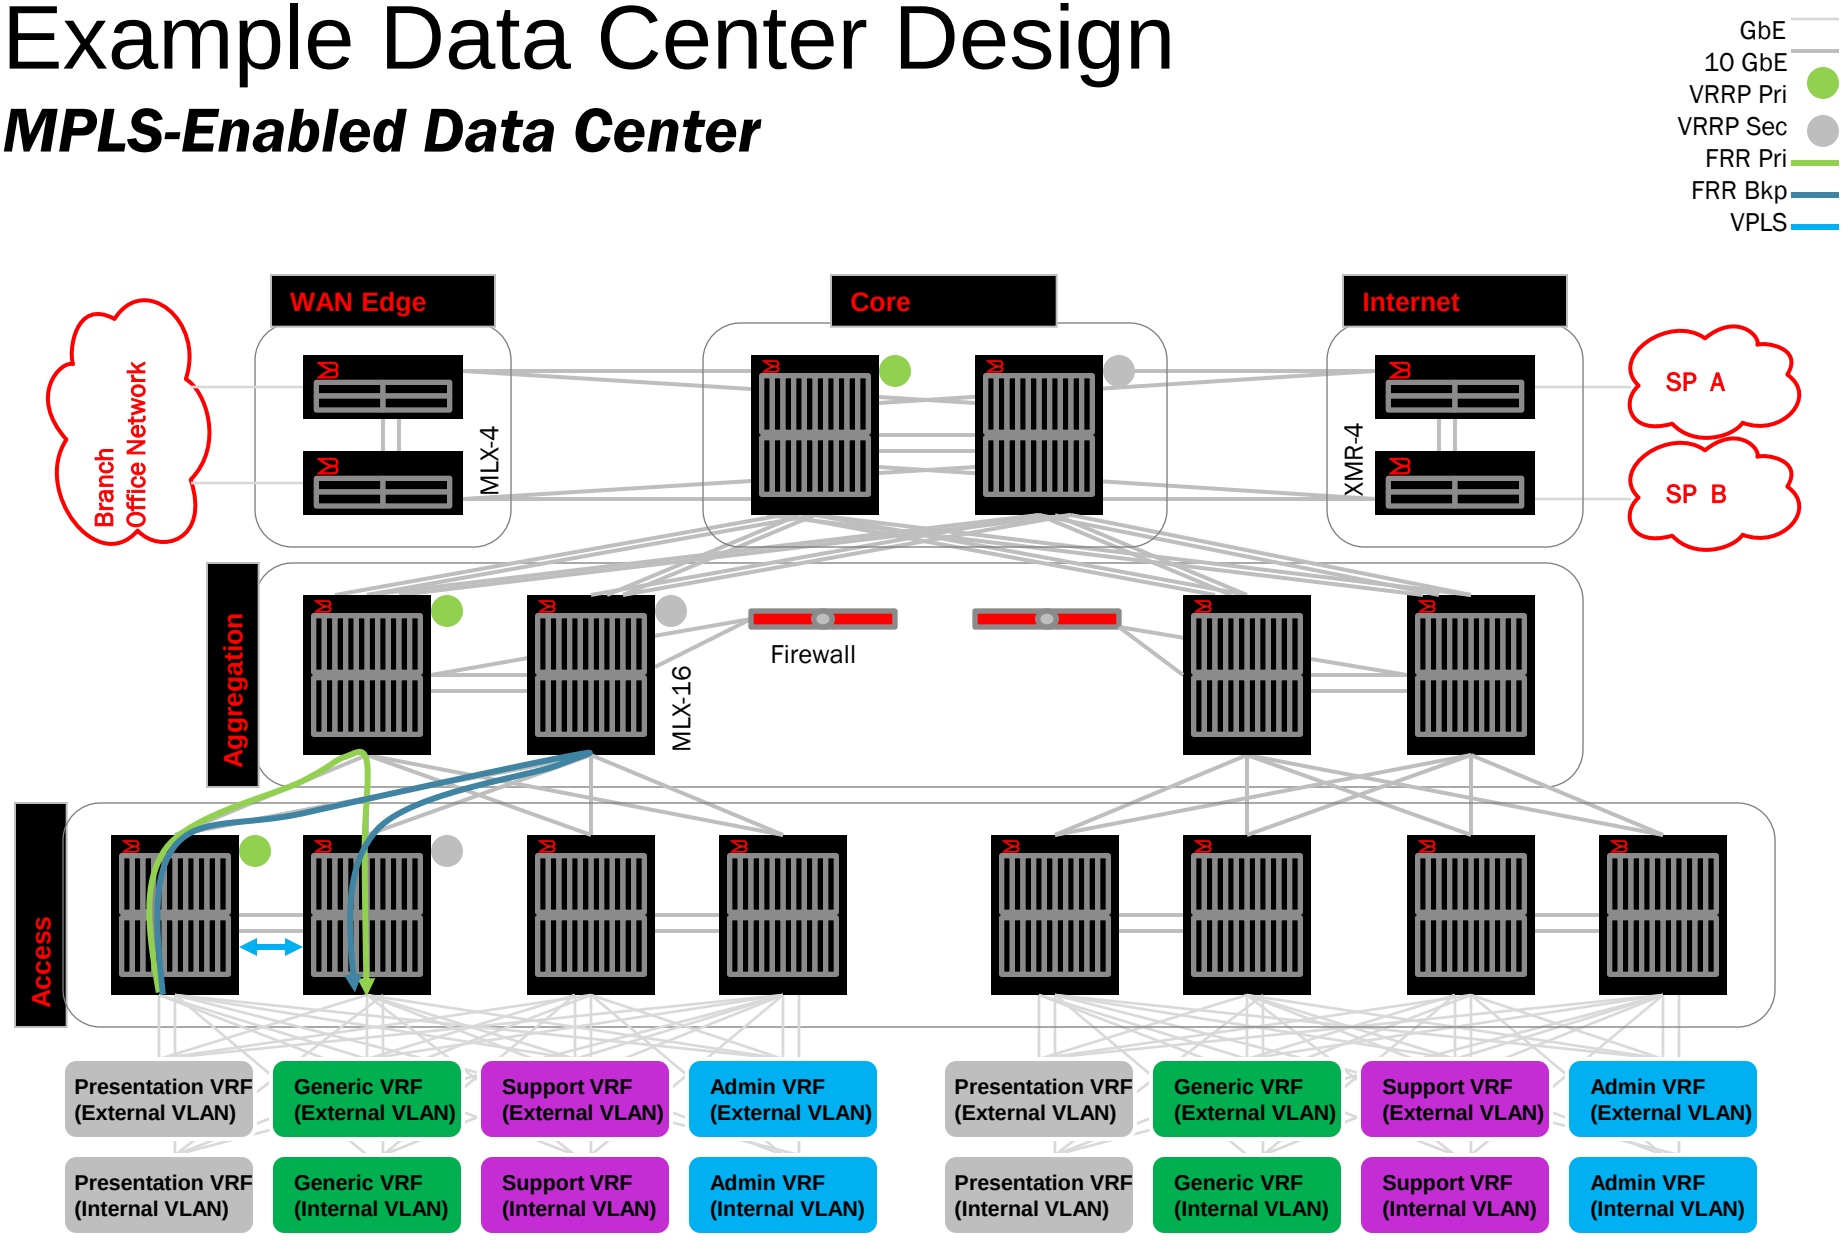

Alternatives and their key components

- End-/Middle-of-Row: High density, non-blocking line card with intelligent cabling solution

CapEx savings: 35%; Reduction in power and cooling: 58%

- Top-of-Rack: Full MPLS-featured, high performance, affordable device

Load Sharing on MPLS systems

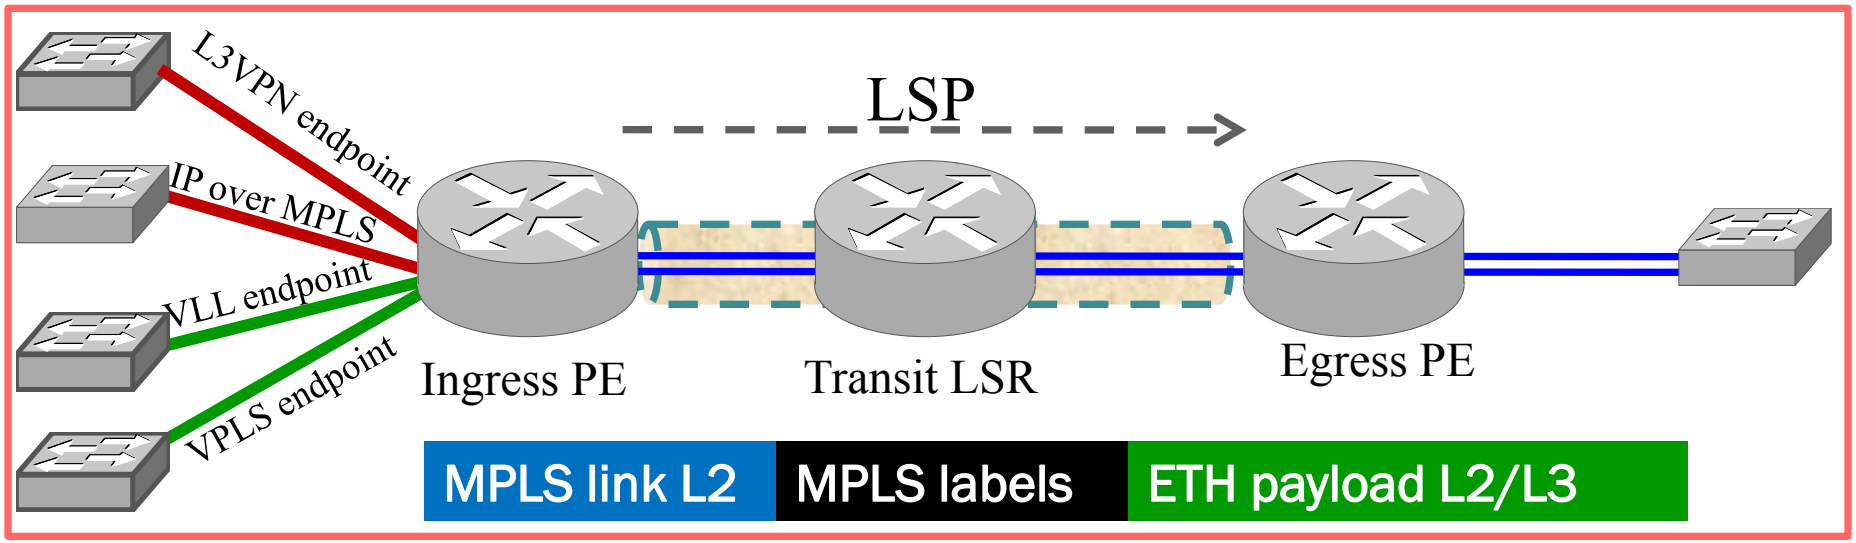

- At Ingress PE (packets entering a MPLS LSP):
  - Load shares IP packets (IP/MPLS, L3VPN, IPv4/v6 in VPLS/VLL) using "L2/L3/L4" headers
- At transit LSR (Checks first nibble after bottommost label):
  - If 4/6, load shares using "MPLS link L2/LSP Label/VC label/**Payload(L3/L4)**" headers
  - Else, load shares using "MPLS link L2/LSP Label/VC label/**Payload(L2/L3)**" headers
- At Egress PE (packets exiting a MPLS LSP):
  - Load shares IP packets with no Ethernet header (IP/MPLS, L3VPN) using "MPLS link L2/Label1/Label2/**Payload(L3/L4)**" headers
  - Load shares IP packets with Ethernet header (IPv4/v6 in VPLS/VLL) using "MPLS link L2/Label1/Label2/**Payload(L2/L3)**" headers

Virtualizing the network (and the service) using TRILL

RBridges

Overview

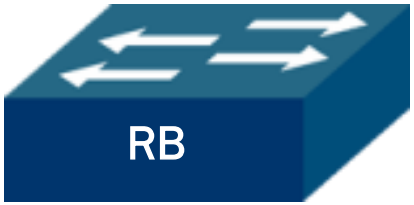

- Implement TRILL protocol
- Use link state routing
- Perform L2 forwarding
- Provide point-to-point forwarding with zero configuration
- Can auto configure themselves
- RBridges forwarding tables scale with the number of RBridges
- RBridges know what options other RBridges support

- Support multipathing for unicast and multicast traffic
- Compatible with classic bridges and can be deployed in bridged LANs
- Ingress RBridge adds *TRILL & outer MAC headers* to frames
- Outer MAC header is modified hop-by-hop as with routing
- Egress RBridge decapsulates frames and learns the association of the “Inner MAC SA” with the Source RBridge nickname

How RBridges Work?

Designated RBridge - DRB

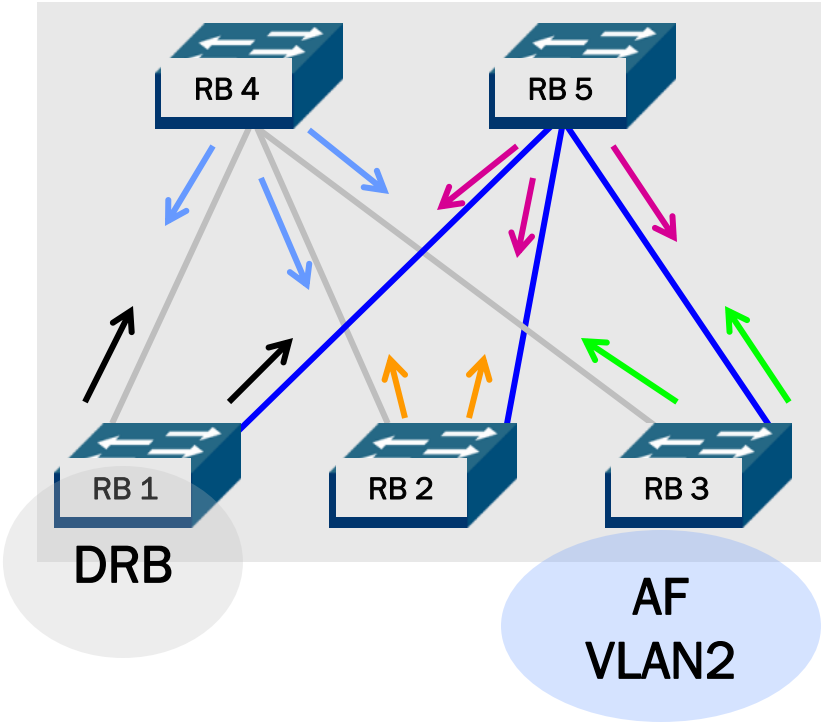

- RBridges discover each other by exchanging TRILL IS-IS (or FSPF) Hello frames
  - TRILL Hellos are sent to the All-IS-IS-RBridges multicast address
- Using link state protocol (IS-IS or FSPF), a single Designated RBridge (DRB) is elected from among all RBridges on the LAN
  - The DRB specifies the Appointed Forwarder (AF) for each VLAN
  - The DRB also specifies the Designated VLAN for inter-RBridge communication

How RBridges Work?

Appointed Forwarder - AF

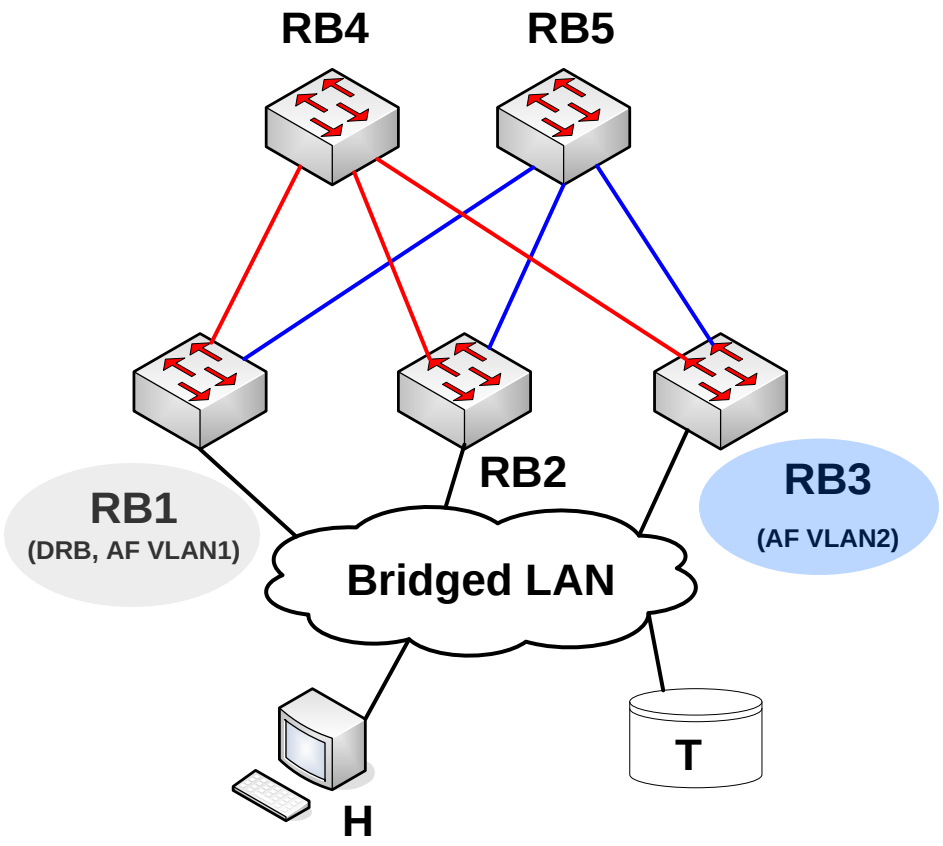

- The DRB specifies the Appointed Forwarder (AF) for each VLAN
  - DRB can also be the AF
- Only *ONE* AF can be appointed per VLAN; One VLAN - One AF
- The AF is in charge of handling all native frames in the VLAN
  - Ingress RBridge function: Encapsulates TRILL data frame
  - Egress RBridge function: Decapsulates TRILL data frames

TRILL Encapsulation

Unicast data path

Known unicast destinations are forwarded RBridge hop by RBridge hop toward the egress RBridge

2. RB1 adds TRILL header and outer MAC header

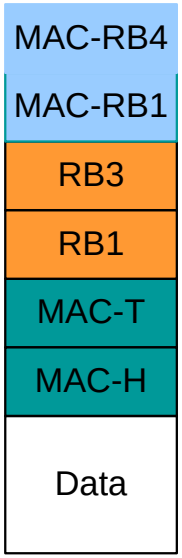

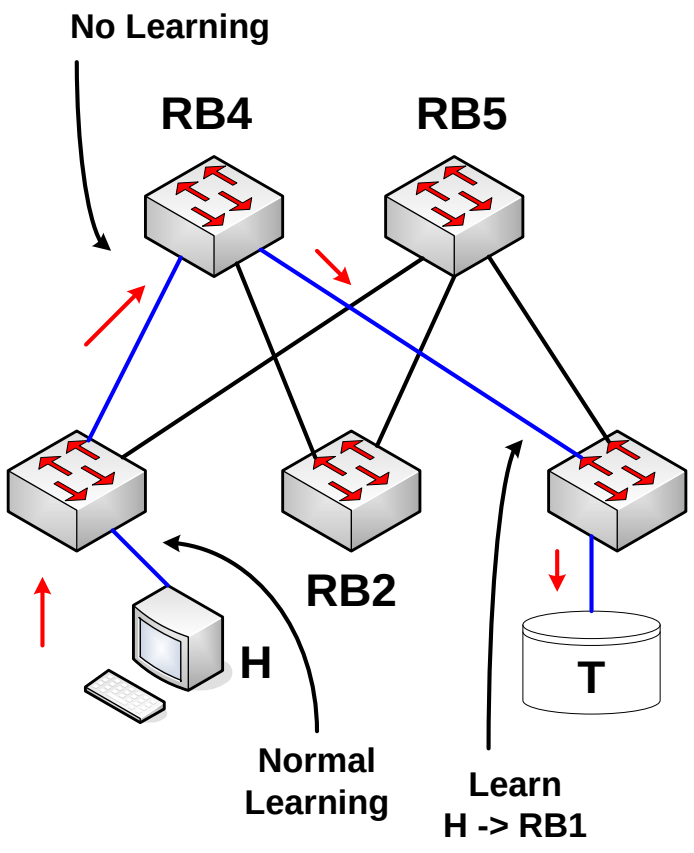

3. RB3 removes TRILL header and associates MAC-H with RB1

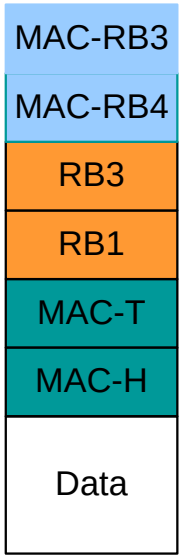

1. Host sends frames

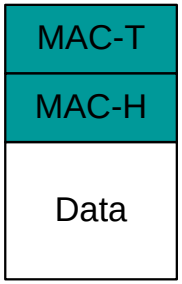

4. Target receives frames

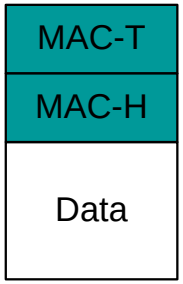

Utilizing ECMP Paths & Reordering

- TRILL supports up to 64 ECMP paths
  - Packet (frames) ordering maintained within flows
- RBridges are required to maintain frame ordering internally
- When multi-pathing is used, all frames for an order-dependent flow must be sent on the same path if unicast or the same distribution tree if multi-destination
- Re-ordering can occur when
  - A destination address transitions between being known and unknown
  - A topology change occurs

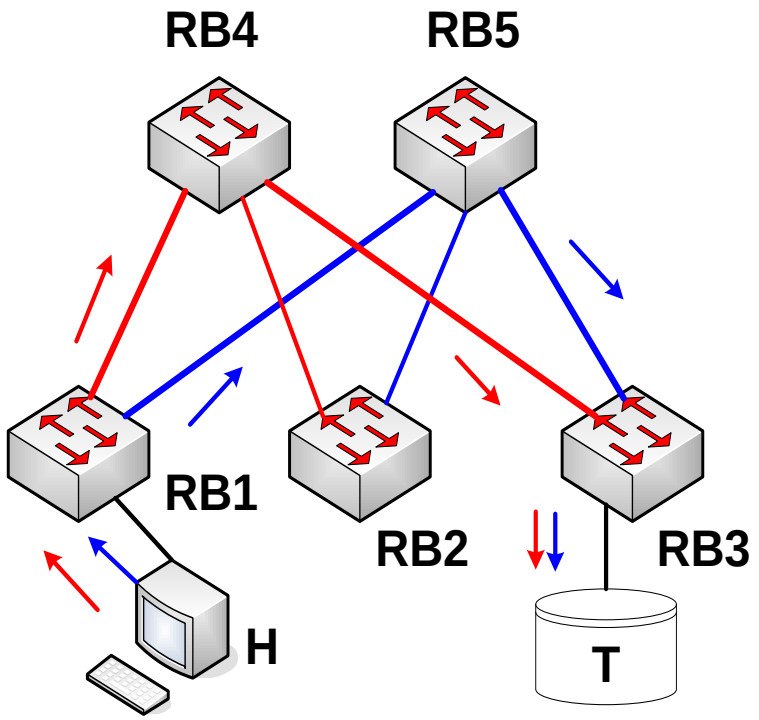

CORE TECHNOLOGY

Brocade Virtual Cluster Switching (VCS)

ETHERNET FABRIC	DISTRIBUTED INTELLIGENCE	LOGICAL CHASSIS	DYNAMIC SERVICE INSERTION

- First data center Ethernet fabric
- No Spanning Tree Protocol
- Multi-path, deterministic
- Auto-healing, non-disruptive
- Lossless, low latency
- Built for convergence

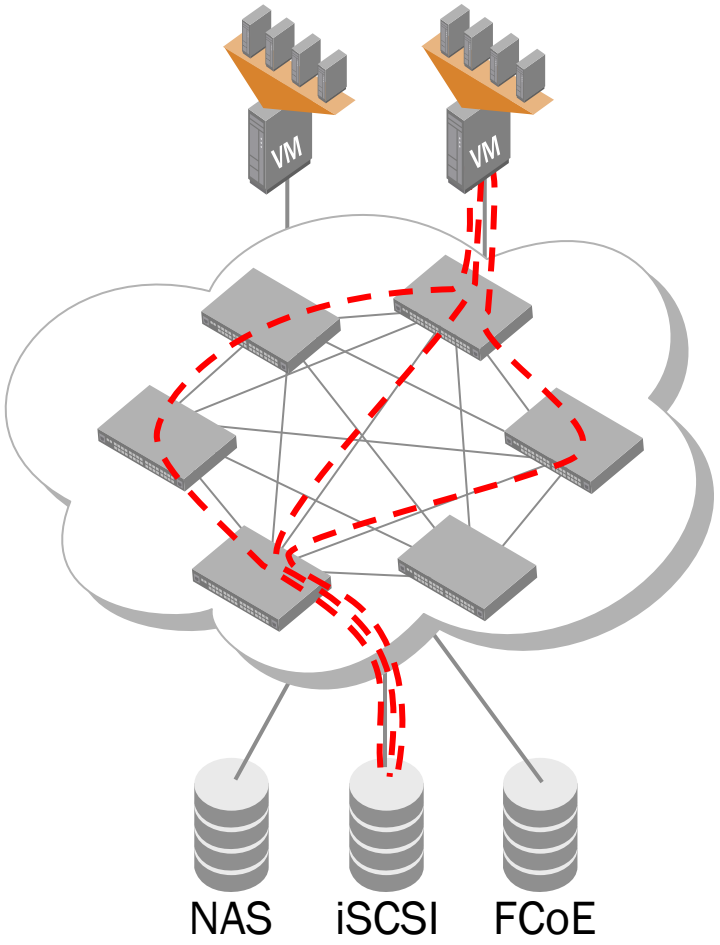

CORE TECHNOLOGY

Brocade Virtual Cluster Switching (VCS)

ETHERNET FABRIC	DISTRIBUTED INTELLIGENCE	LOGICAL CHASSIS	DYNAMIC SERVICE INSERTION

- Fully distributed control plane
- Arbitrary topology, self-forming
- Network-wide knowledge of all members, devices, VMs
- Automatic Migration of Port Profiles (AMPP)

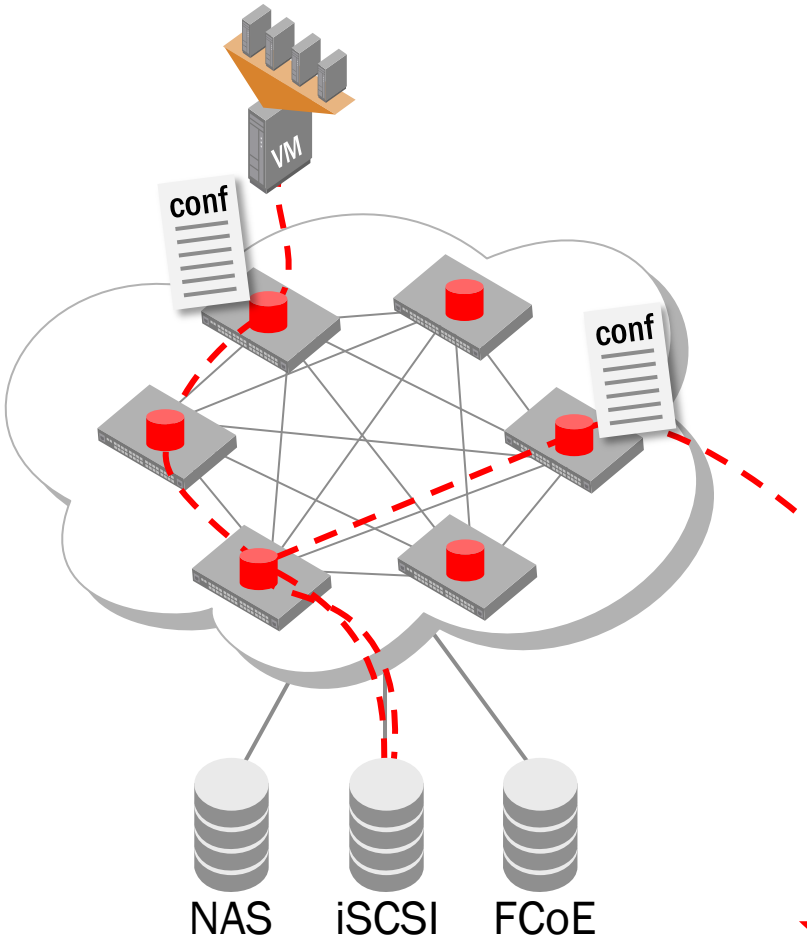

CORE TECHNOLOGY

Brocade Virtual Cluster Switching (VCS)

ETHERNET FABRIC	DISTRIBUTED INTELLIGENCE	LOGICAL CHASSIS	DYNAMIC SERVICE INSERTION

- Managed as a single switch
- Logically collapses network layers
- Auto-configuration for new devices
- Centralized or distributed management
- Radically reduces managed elements

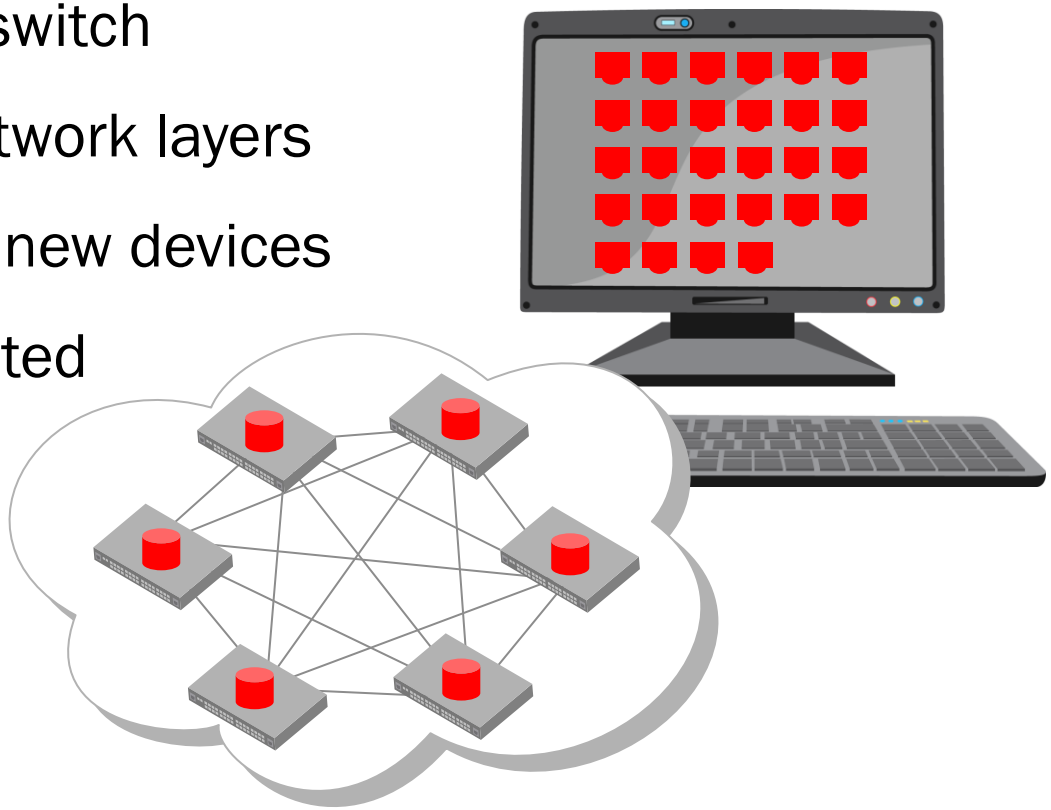

CORE TECHNOLOGY

Brocade Virtual Cluster Switching (VCS)

ETHERNET FABRIC	DISTRIBUTED INTELLIGENCE	LOGICAL CHASSIS	DYNAMIC SERVICE INSERTION

- Reconfigure network via software
- Hardware-based flow redirection
- Incorporation of partner services
- Non-stop service insertion
- Minimizes cost and physical moves

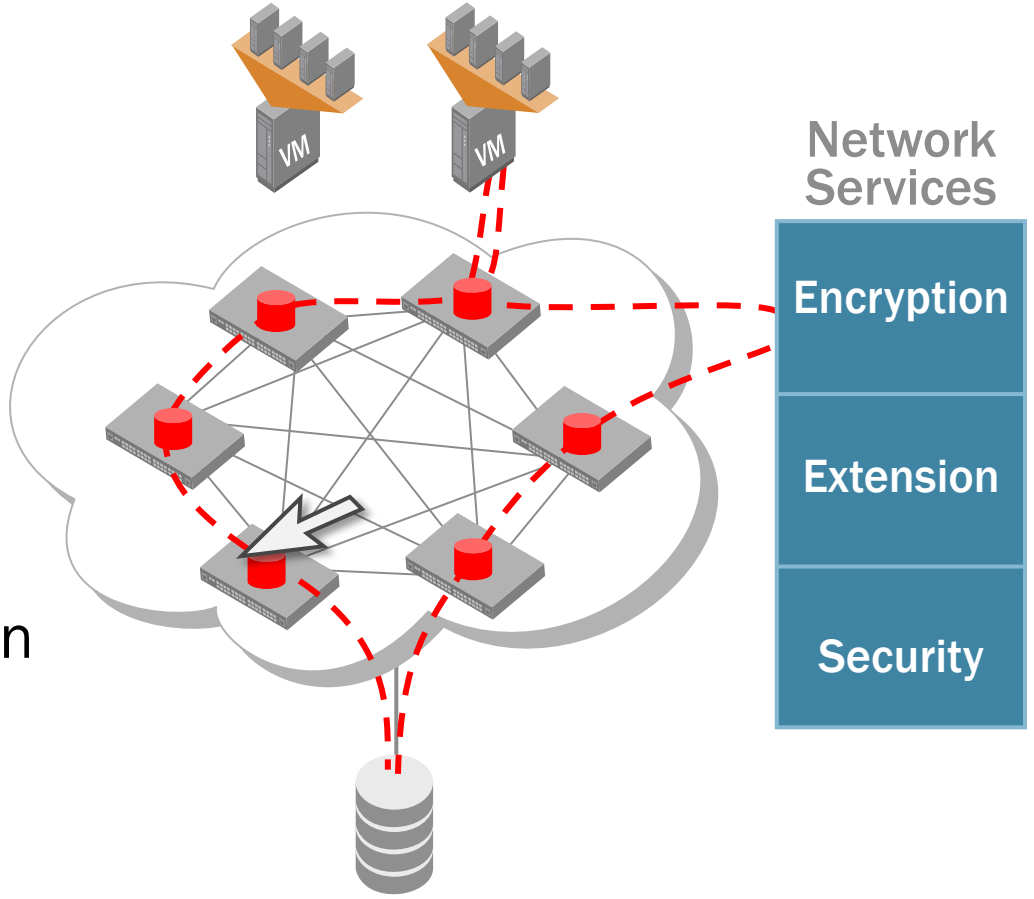

Virtualizing the IO

IO technology used for Virtual Server

Figure 7. Virtualized Server/External Controller-Based Disk Storage Interface Protocols

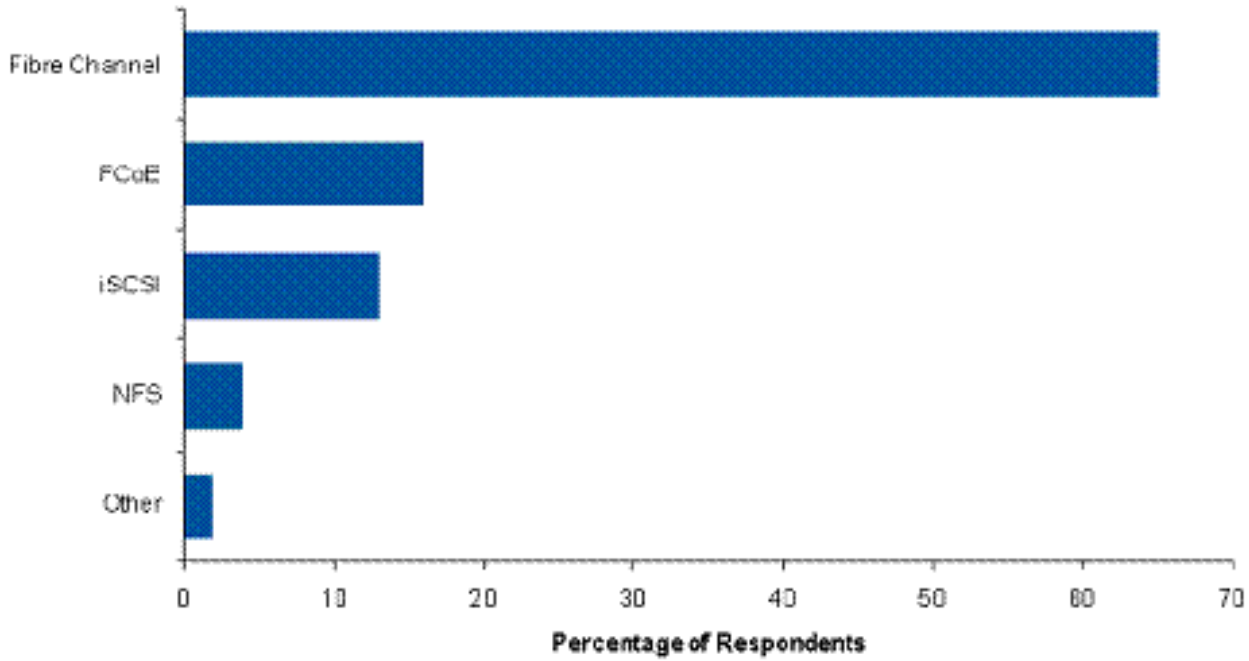

Source: Gartner (January 2010)

Question: Which external controller-based disk storage feature do you consider to be the most important in a virtualized server infrastructure? (Select only one.)

Unified IO

Deployment Use Case

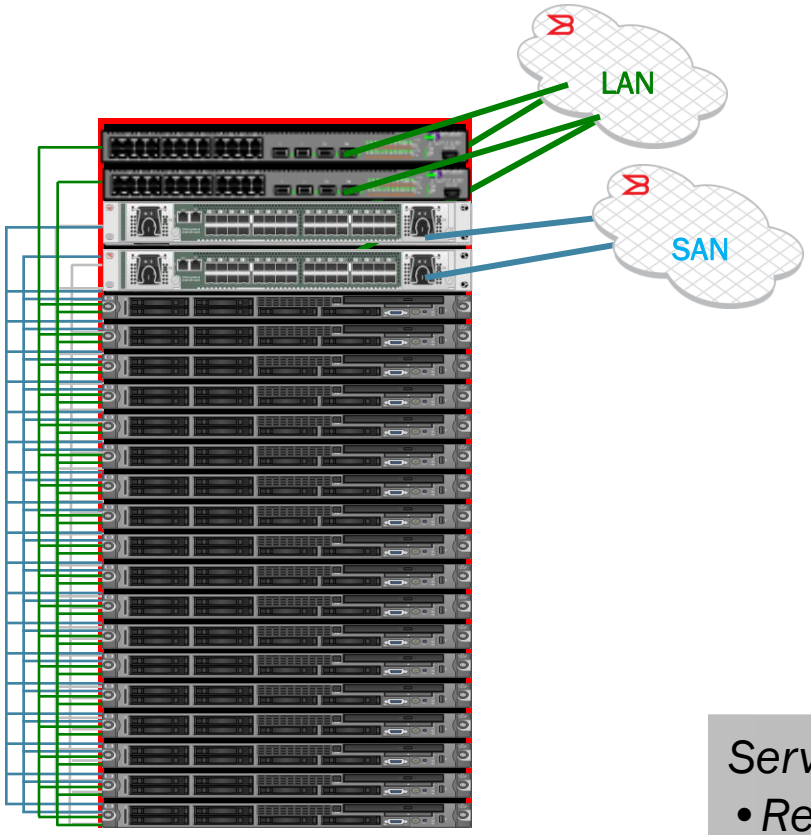

Issue

- Reduce server connectivity cost
- Proprietary solutions and tagging
- Integration with existing networks

Solution

Converged technology:

- Defining CEE/FCoE standards
- Delivering switches and cards certified by all major system vendors
- Convergence at desired pace, level
- Integration into new blade servers

Server I/O convergence:

- *Reduces cables and I/O cards up to 50%*
- *Is compatible with installed networks*
- *Requires no rip and replacement, helping to preserve network investments*

Fibre Channel Gets a New Transport Layer

Encapsulation

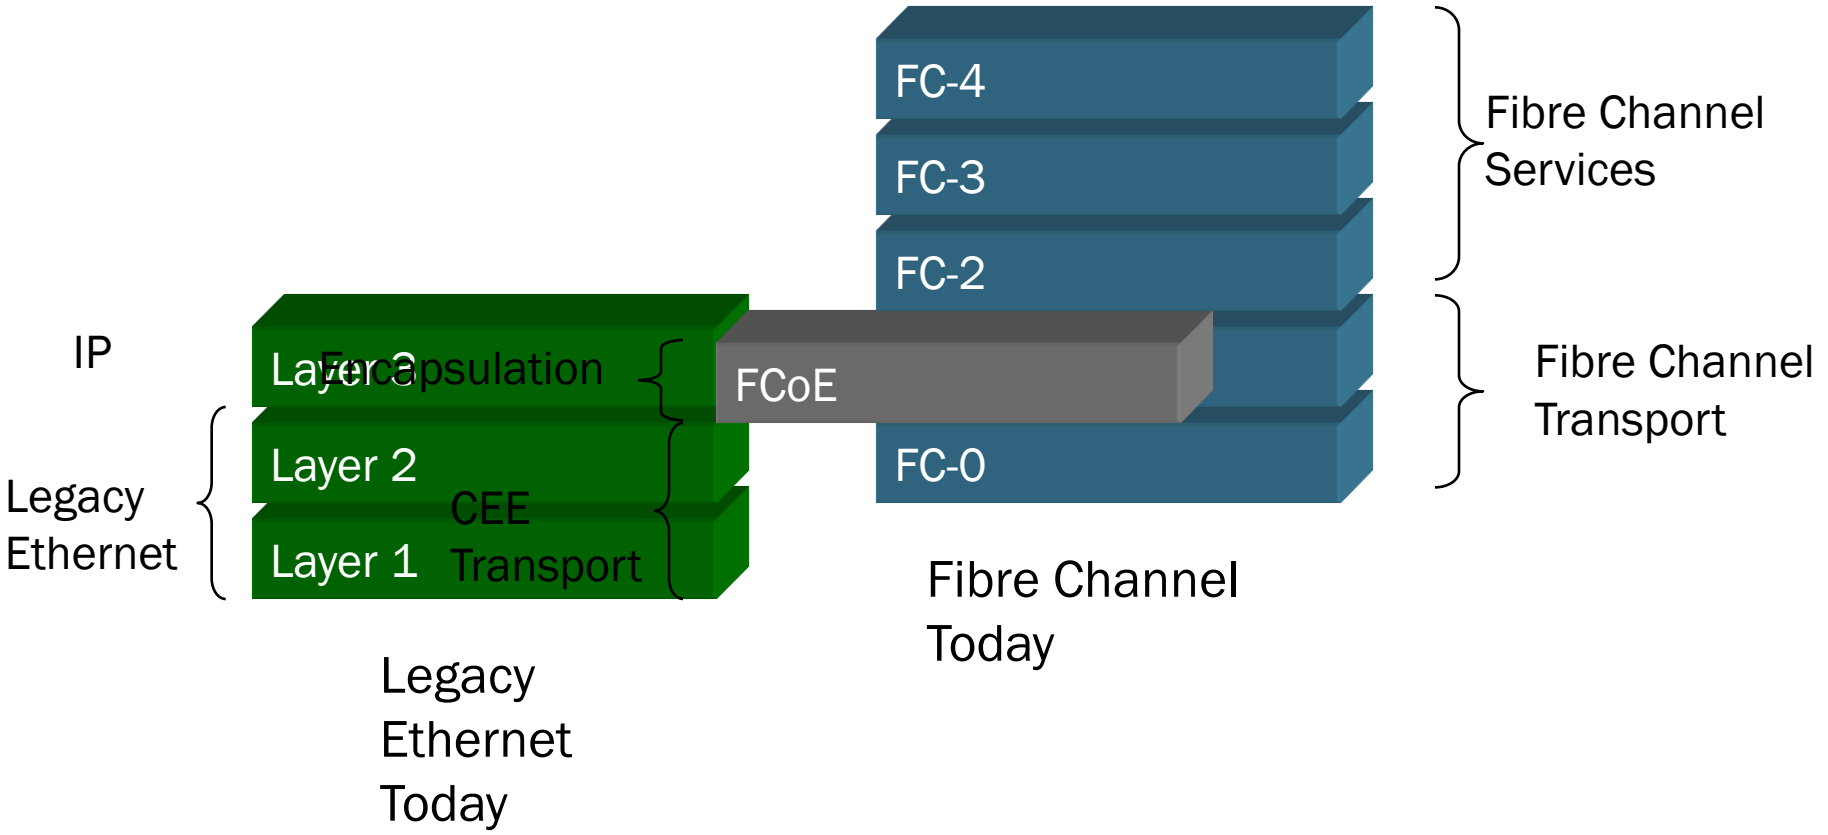

Fibre Channel over Ethernet (FCoE)

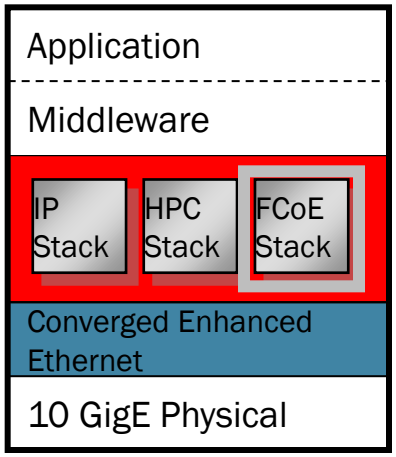

- A mechanism to carry Fibre Channel protocol on top of Converged Enhanced Ethernet
- Encapsulates FC frames over CEE
- Leverages the rich set of FC fabric services for storage connectivity
  - Name services, zoning, WWN structure, management, multi-pathing
- Preserves investment in SAN infrastructure
- Converged server connectivity through one physical connection

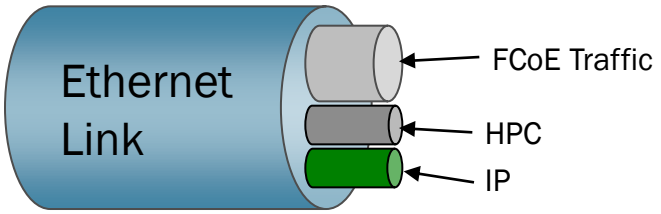

Priority Flow Control - 802.1Qbb

- Converged network environments (CEE) carry several different classes of traffic including FCoE and TCP/IP across a single link
- Priority Flow Control (PFC) provides link-level flow control independently for different classes of traffic
- This allows for coexistence of lossy and lossless traffic on the same physical interface

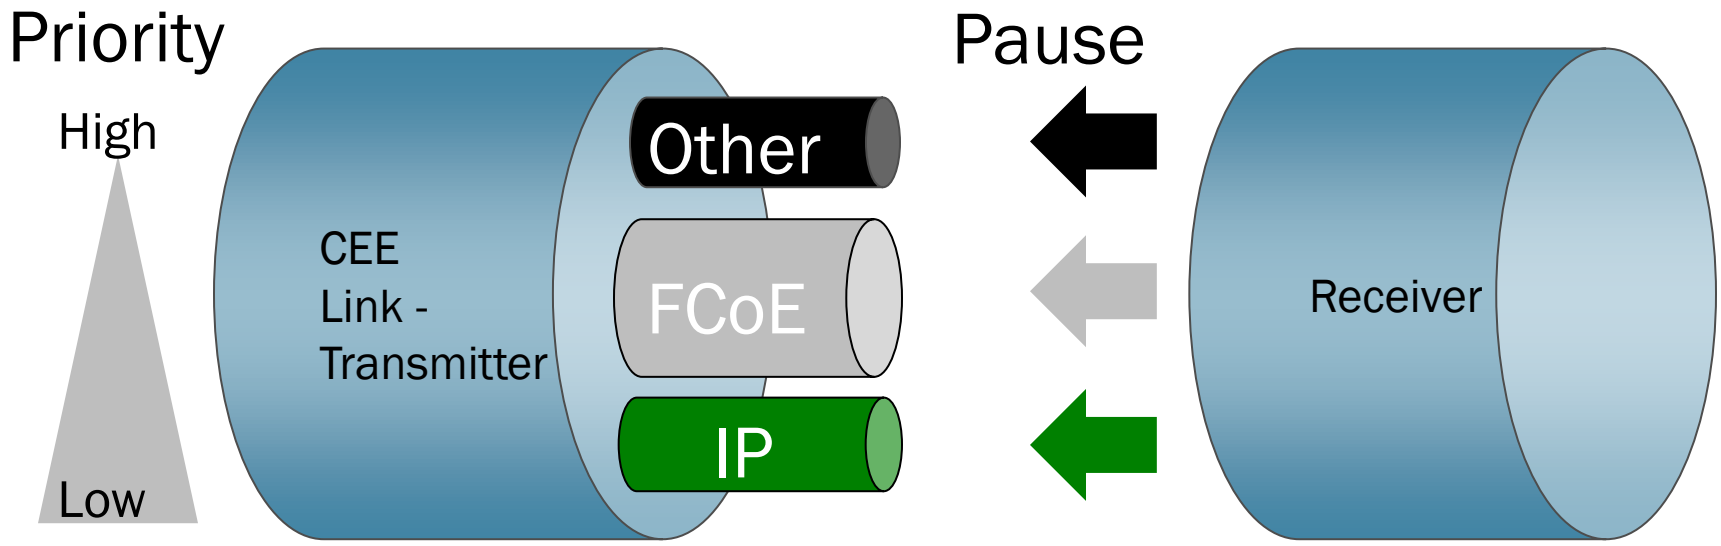

Enhanced Transmission Selection - .1Qaz

- Provides a consistent management framework for assigning different traffic classes (priorities) to different priority groups
- Allows for configurable bandwidth allocation to each priority group
- Data Center Bridging eXchange Protocol (DCBX) is used to negotiate ETS and PFC

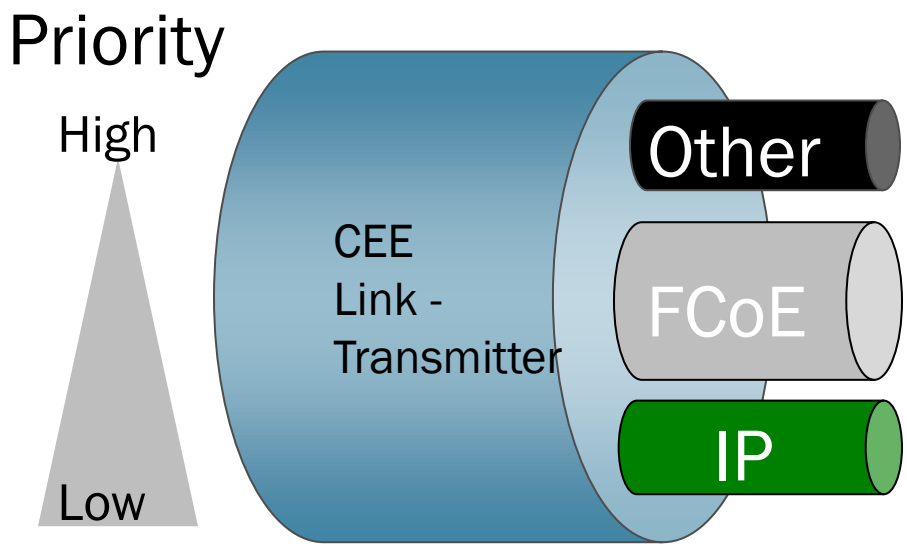

Summary

Summary

- There is no „one size fits all“ – your differentiation is your competitive advantage and hence you need to pick the best solution in each technology area
- Open standards and best-of-breed are key to success
- There are many virtualization-tools that help or can help to provide a service more robust, more scalable and/or more efficient
- Not just a vision, but a reality that is coming together

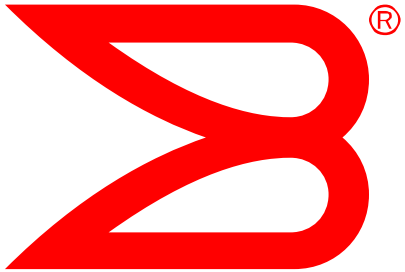

BROCADE

EXTRAORDINARY
NETWORKS