BROCADE

ETHERNET FABRICS ENABLING THE CLOUD OPTIMIZED DATA CENTER

Joerg Ammon (jammon@brocade.com)

Principal 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.

TRILL – Technology Overview

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

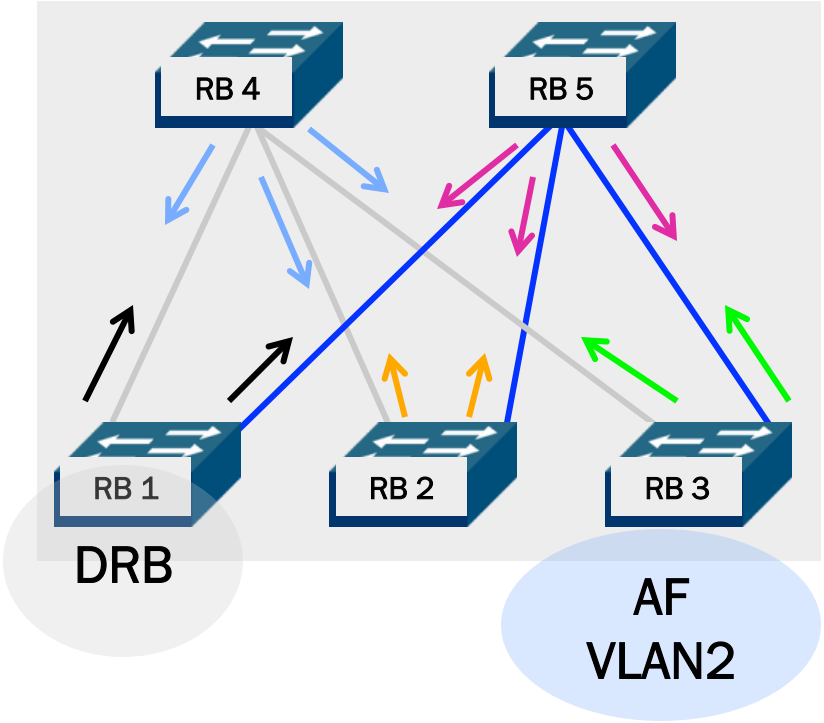

- RBridges discover each other by exchanging TRILL IS-IS (or FSPF) Hello frames
  - TRILL Hellos are sent to the All-IS-IS-RBridges multicast address
- Using link state protocol (IS-IS or FSPF), a single Designated RBridge (DRB) is elected from among all RBridges on the LAN
  - The DRB specifies the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

How RBridges Work?

Appointed Forwarder - AF

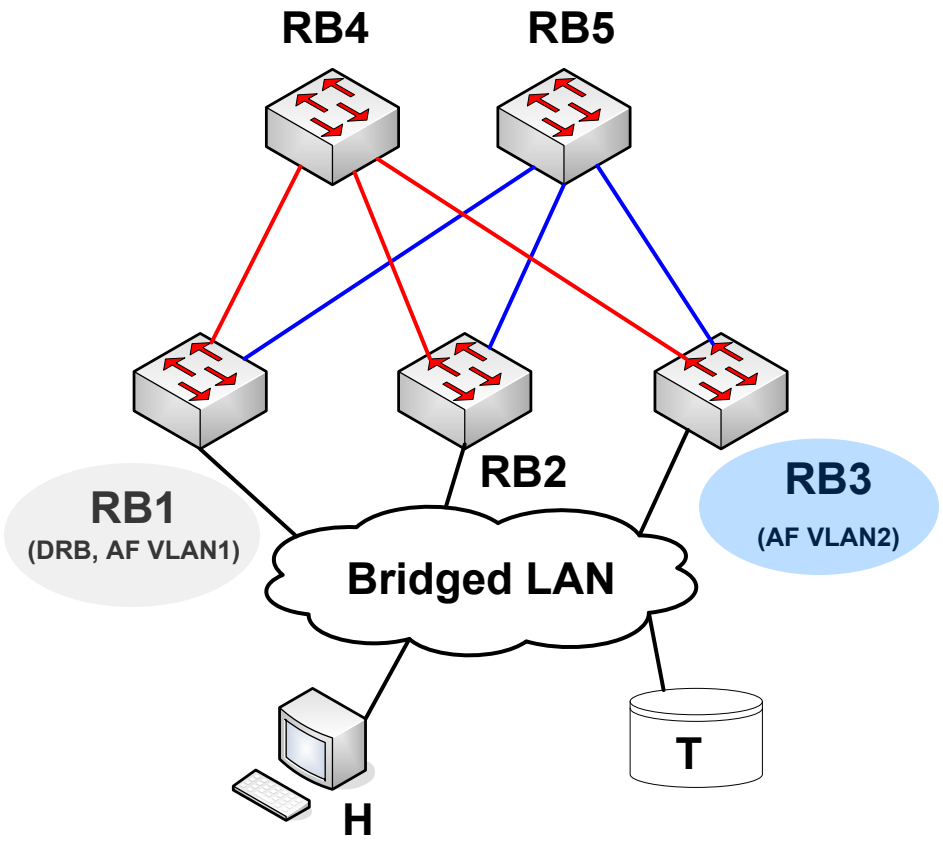

- The DRB specifies the Appointed Forwarder (AF) for each VLAN
  - DRB can also be the AF
- Only *ONE* AF can be appointed per VLAN; One VLAN - One AF
- The AF is in charge of handling all native frames in the VLAN
  - Ingress RBridge function: Encapsulates TRILL data frame
  - Egress RBridge function: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

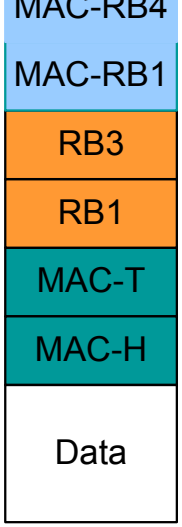

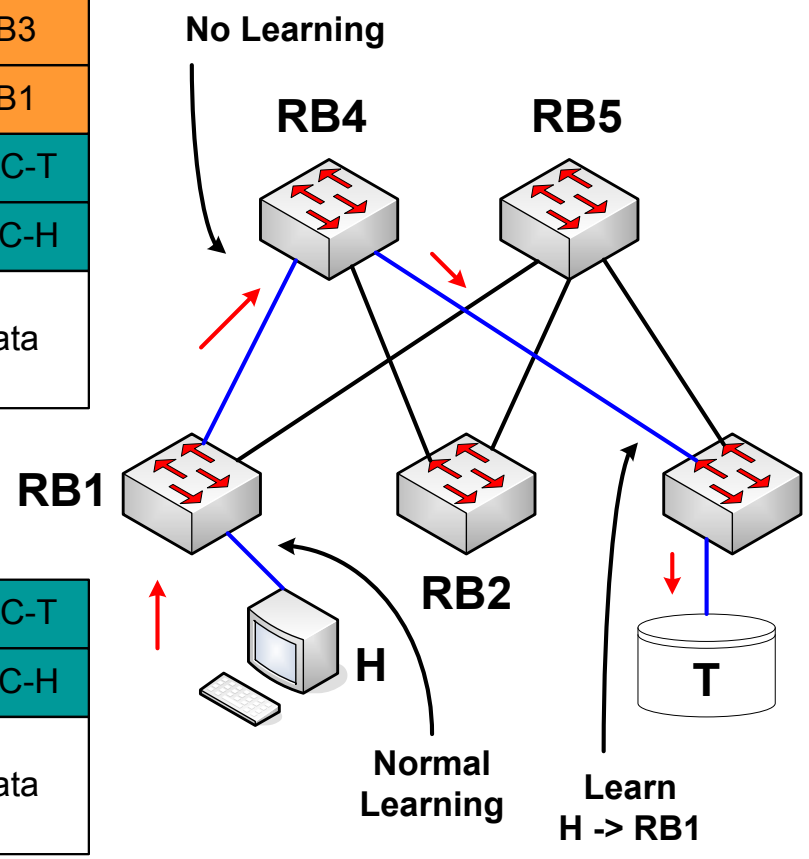

3. RB3 removes TRILL header and associates MAC-H with RB1

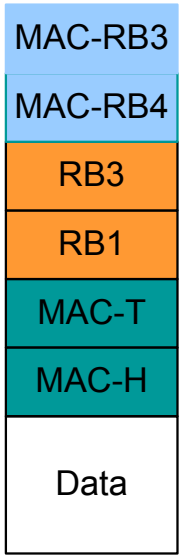

1. Host sends frames

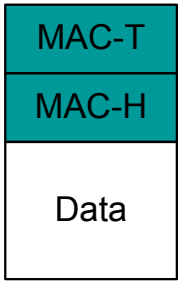

4. Target receives frames

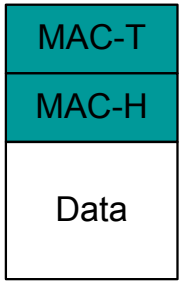

Utilizing ECMP Paths & Reordering

- 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

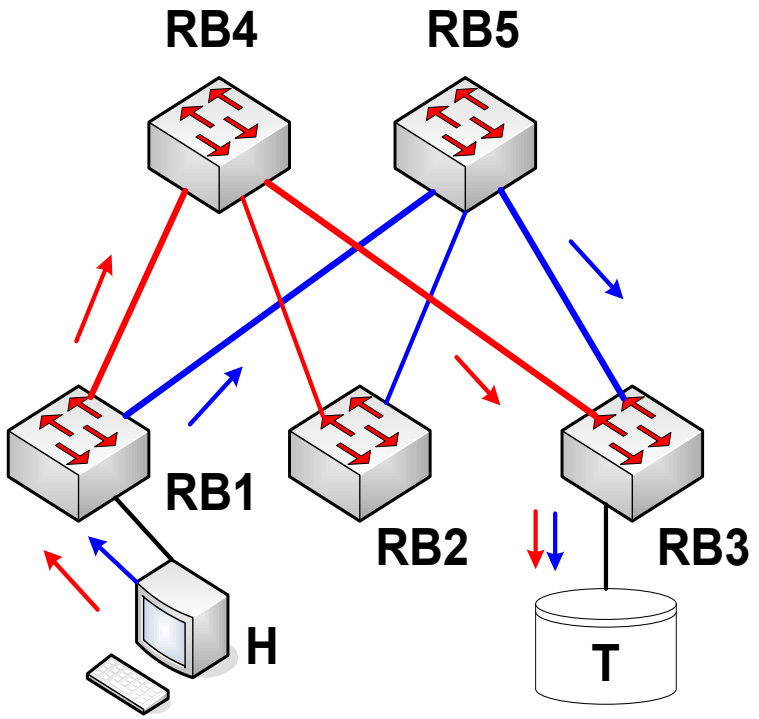

Advantages over conventional L2 with STP

- Multipathing and shortest optimal path
- TTL for loop mitigation
- Client MAC learning only on edge
- vLAGs

WHAT IS AN ETHERNET FABRIC?

- Resilient architecture for today's consolidated environments
- Flexible topology to meet application and storage requirements
- Massive scalability while remaining elastic
- Inherently flat architecture for virtualization and clouds

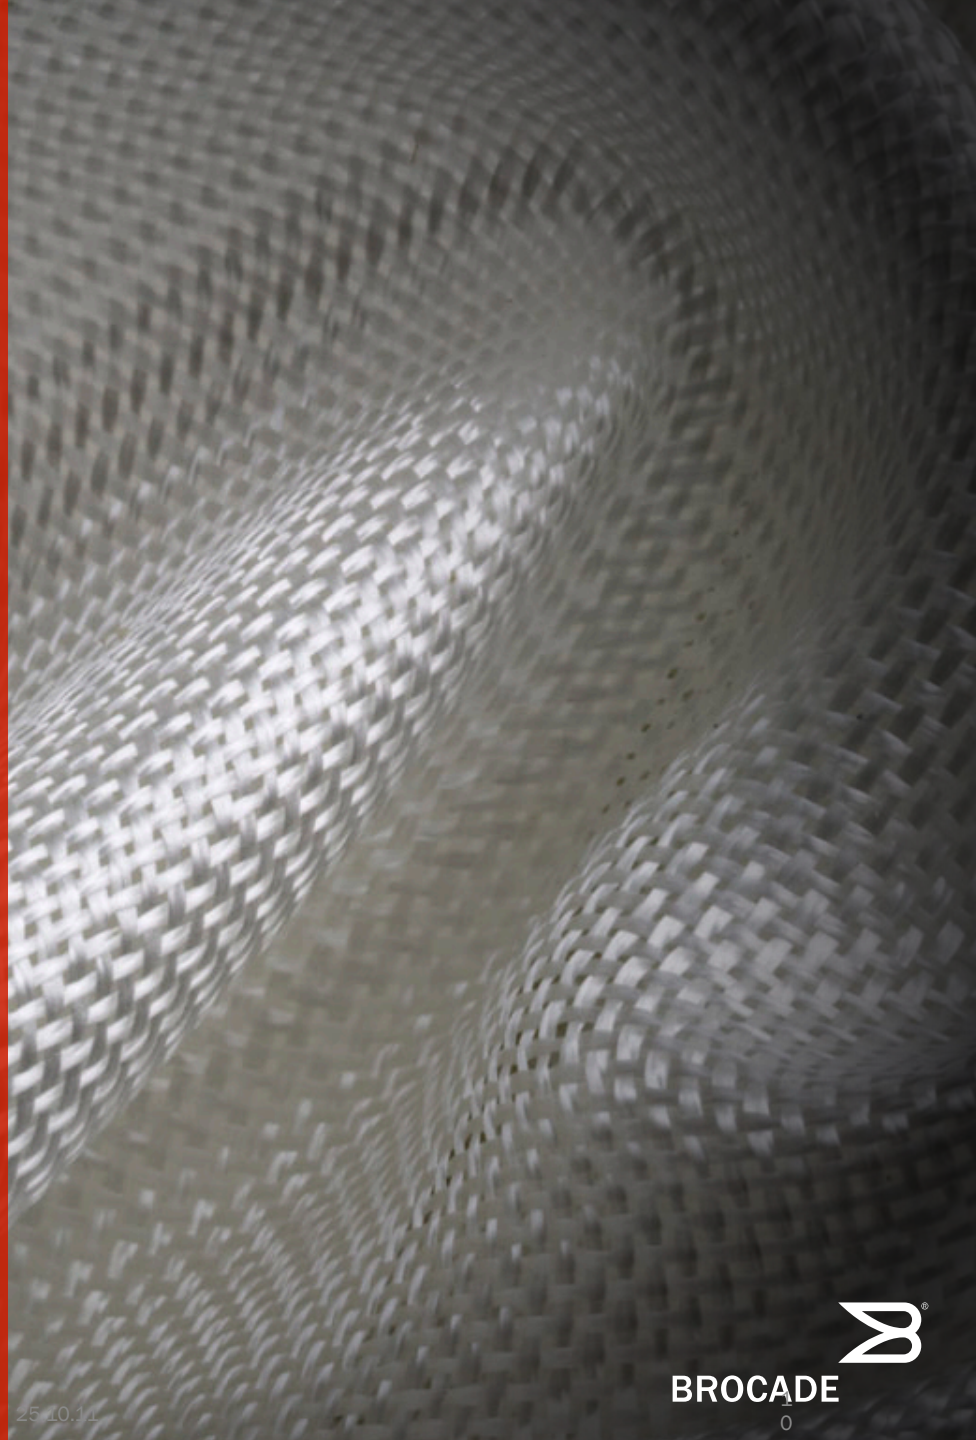

BROCADE

CORE TECHNOLOGY

Brocade Virtual Cluster Switching (VCS)

ETHERNET FABRIC	DISTRIBUTED INTELLIGENCE	LOGICAL CHASSIS	DYNAMIC SERVICE INSERTION

- No Spanning Tree Protocol
- Multi-path, deterministic
- Auto-healing, non-disruptive
- Lossless, low latency
- Built for convergence

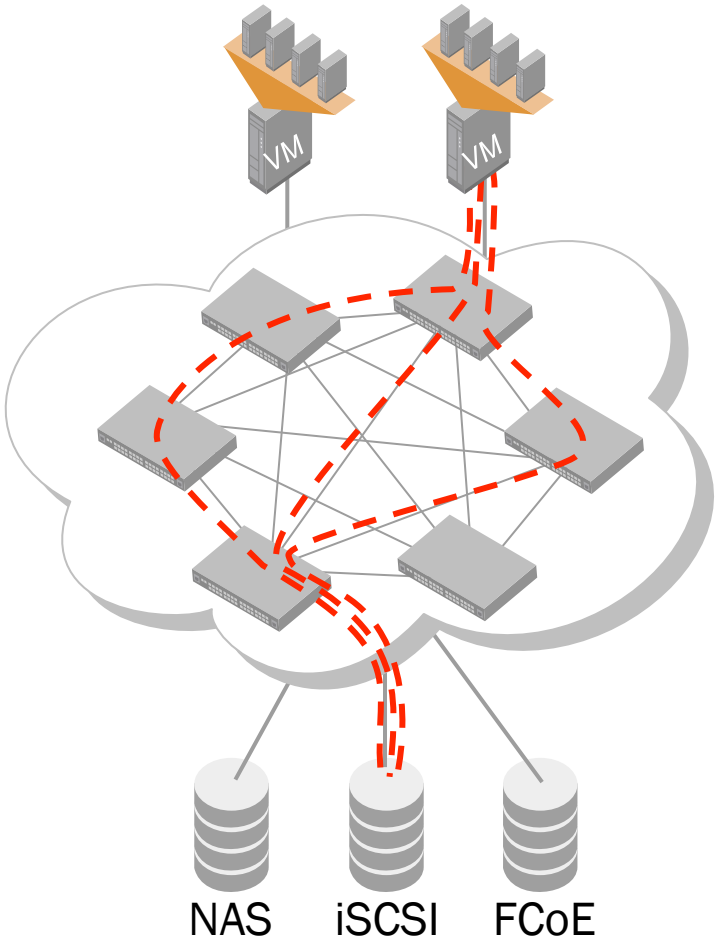

CORE TECHNOLOGY

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)

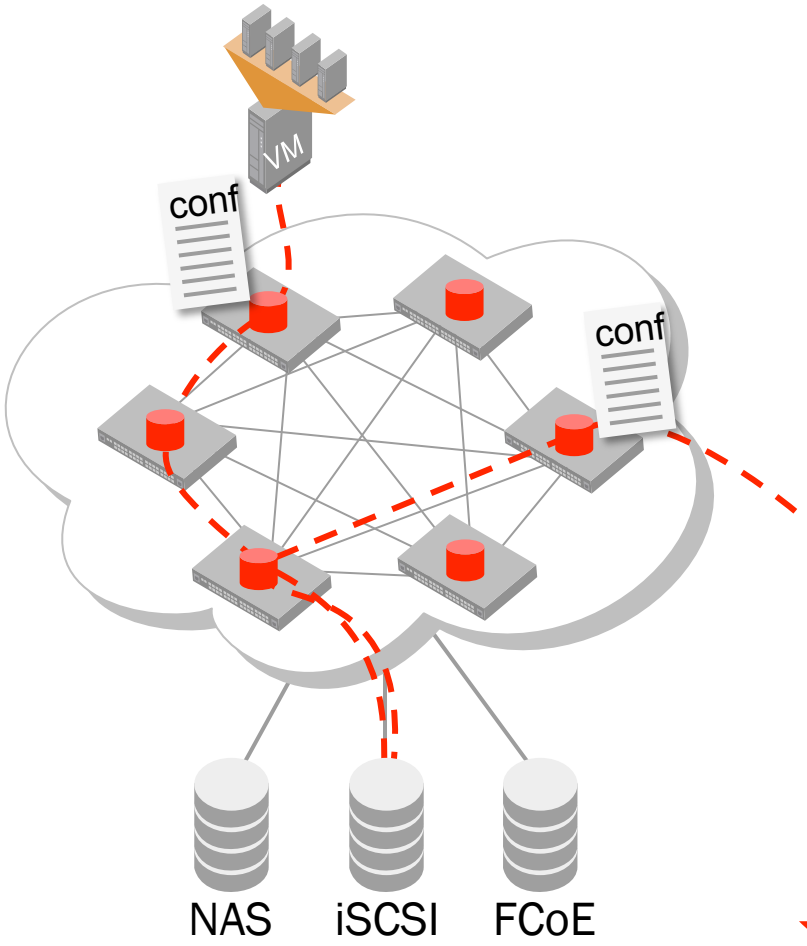

CORE TECHNOLOGY

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

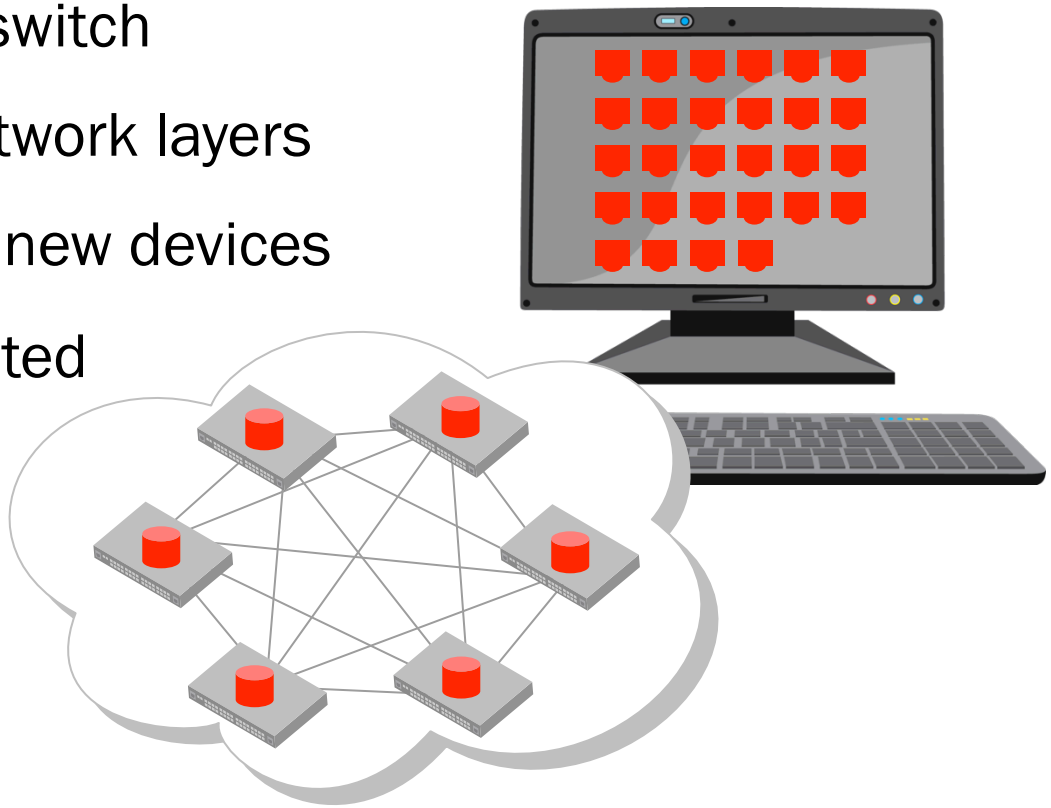

CORE TECHNOLOGY

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

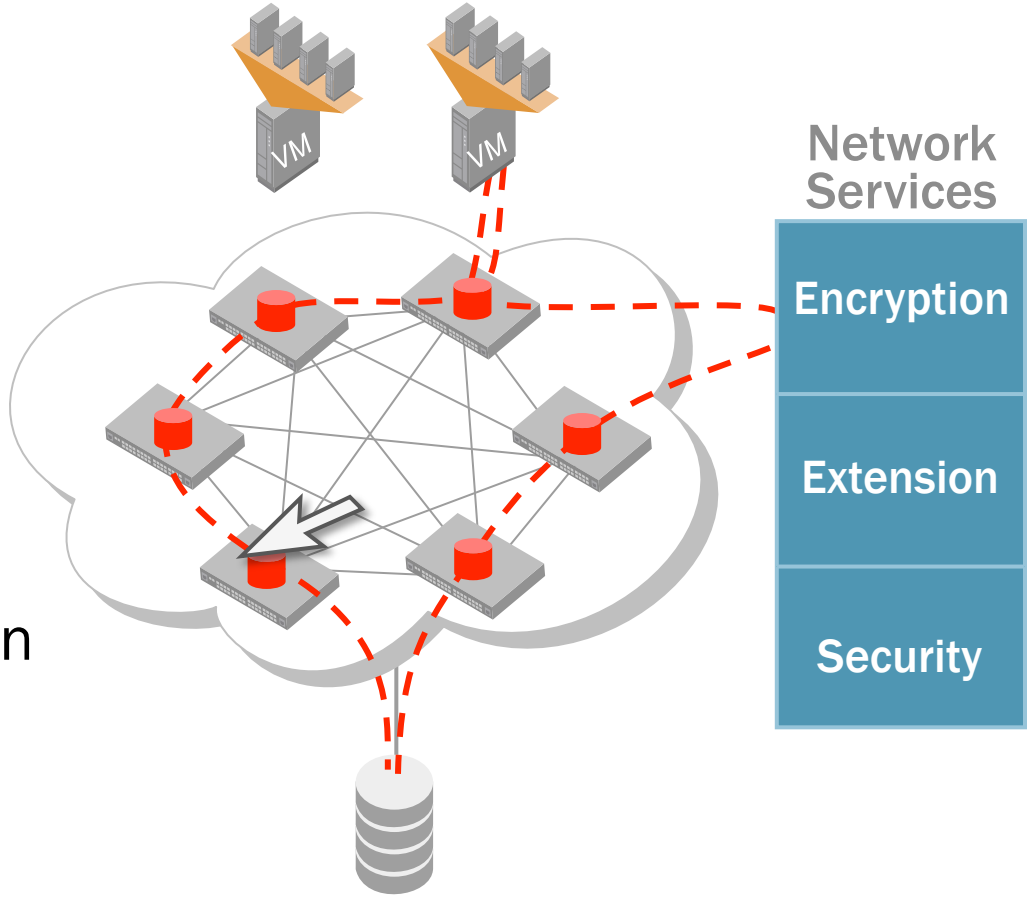

Some more technology components

Priority Flow Control - 802.1Qbb

- Converged network environments (CEE) carry several different classes of traffic including iSCSI, 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

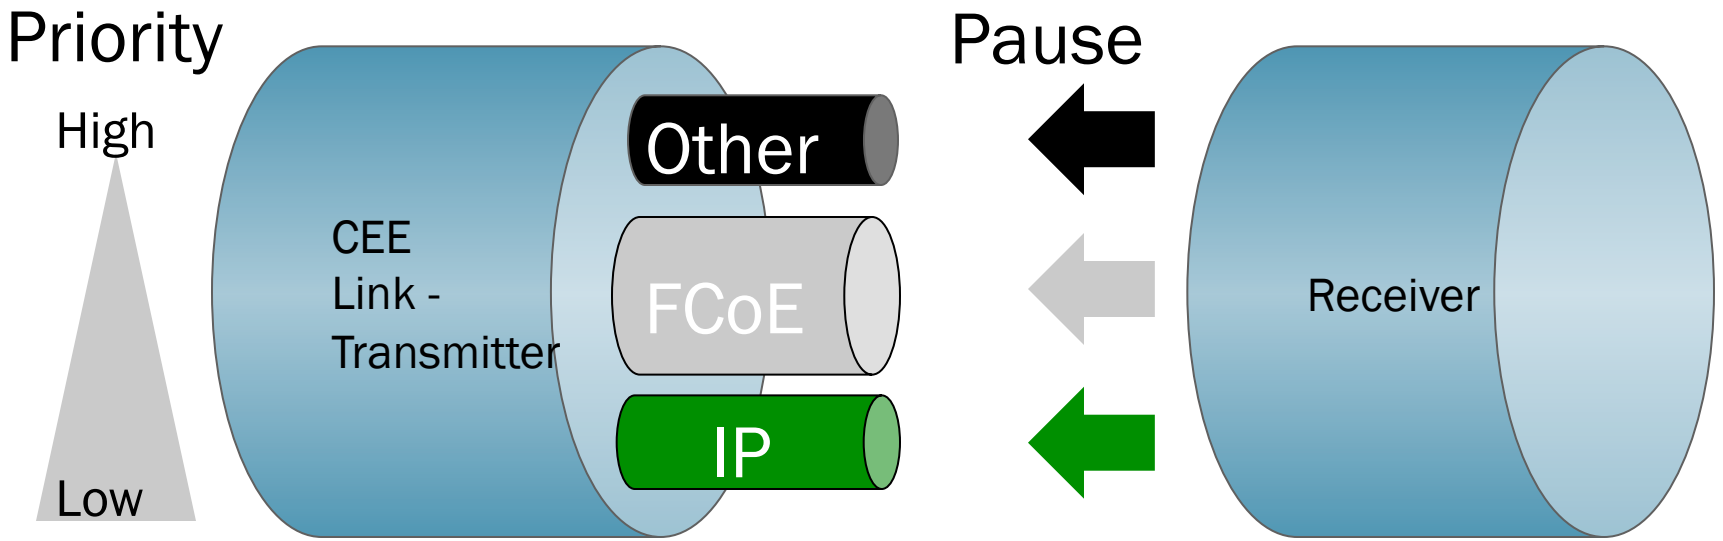

Enhanced Transmission Selection - .1Qaz

- 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

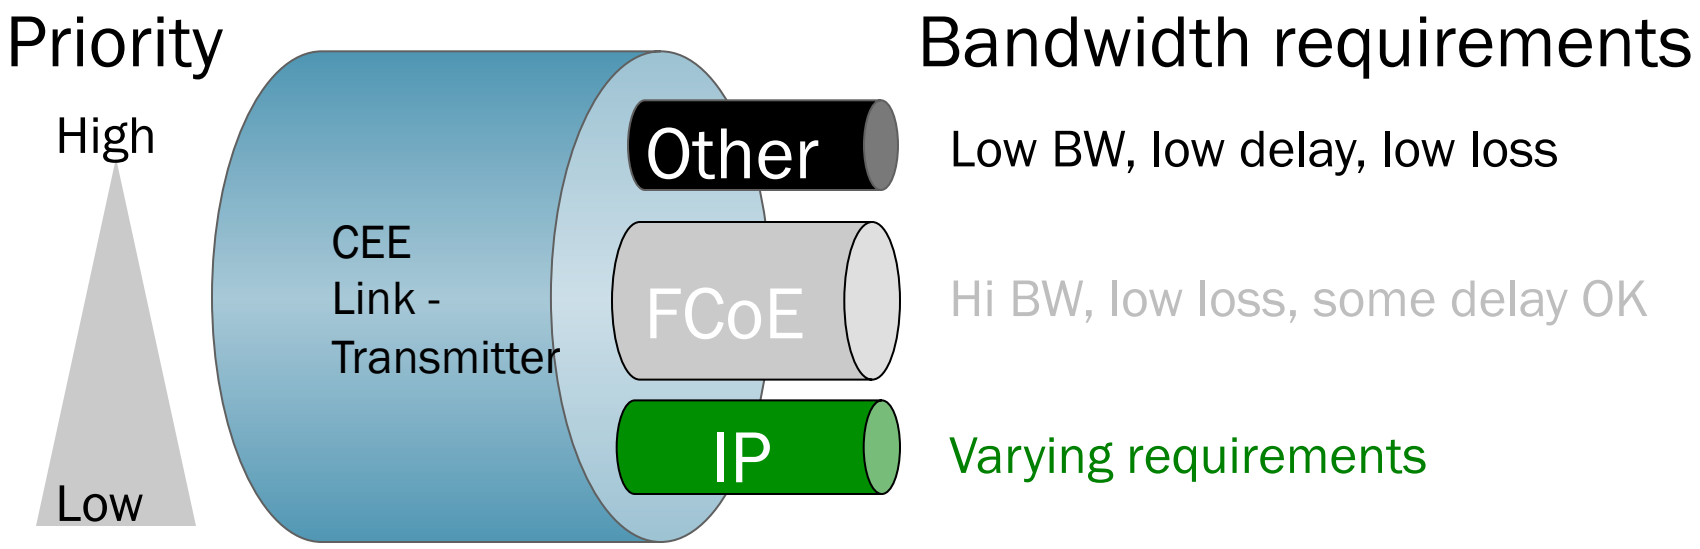

Why another L2-solution?

MARKET TRENDS: DATA CENTER

Data Center Construction

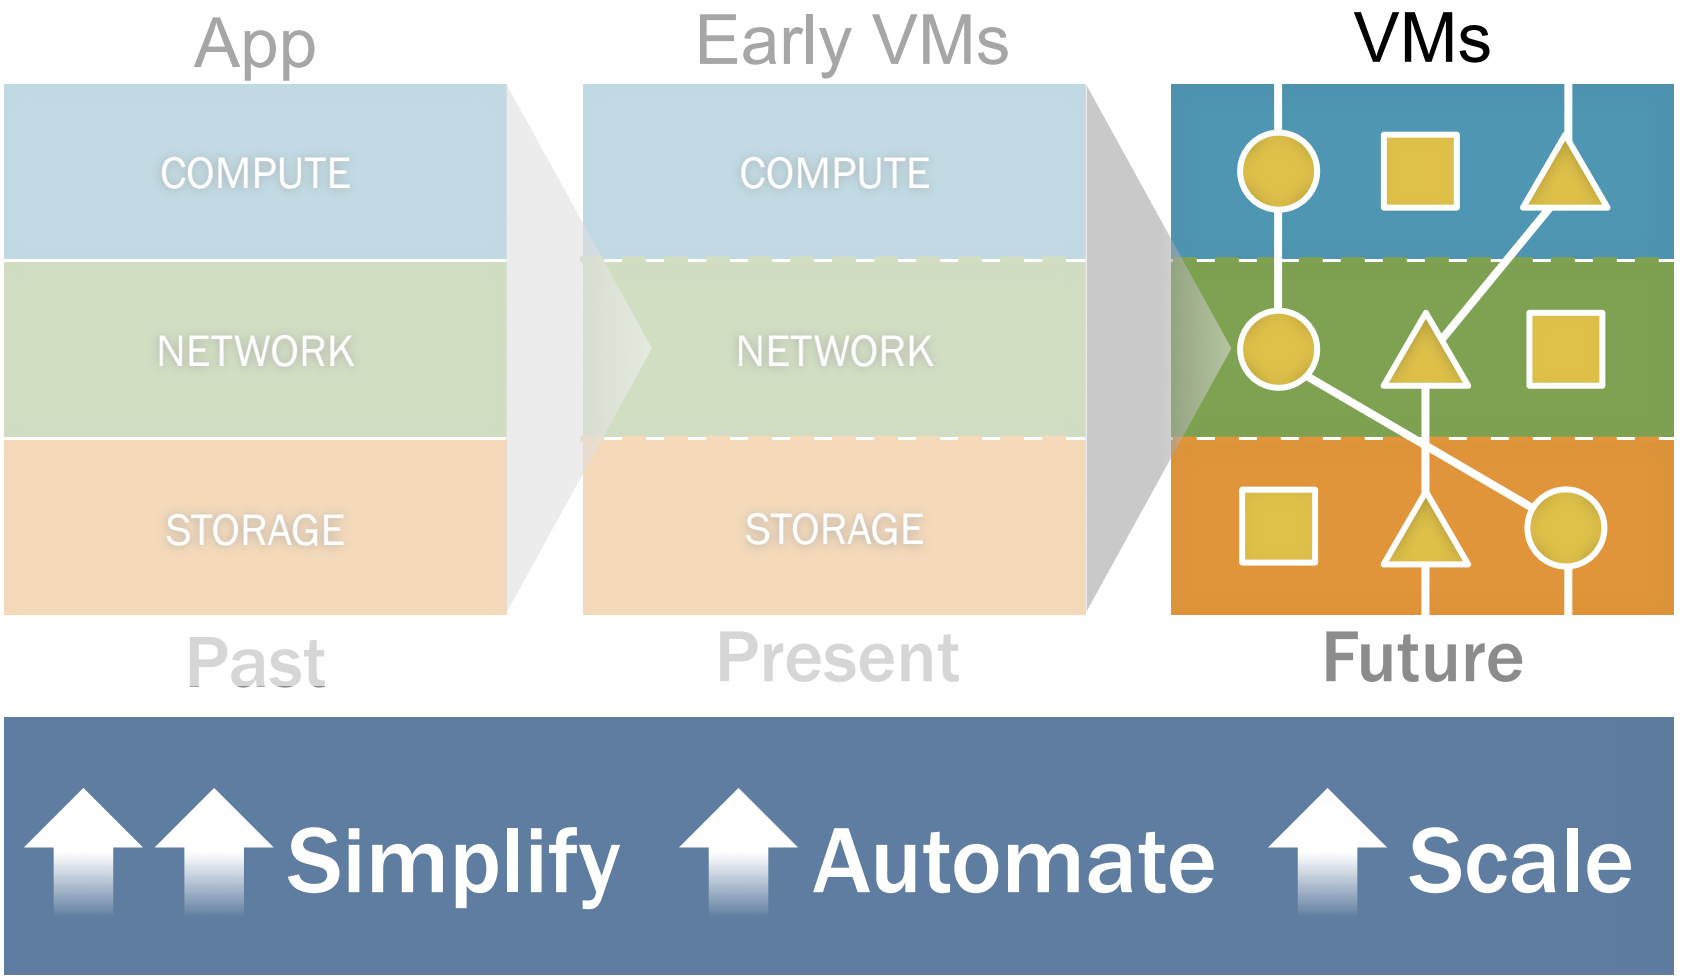

The Network is Central to the Cloud

Virtualization Brings New Demands

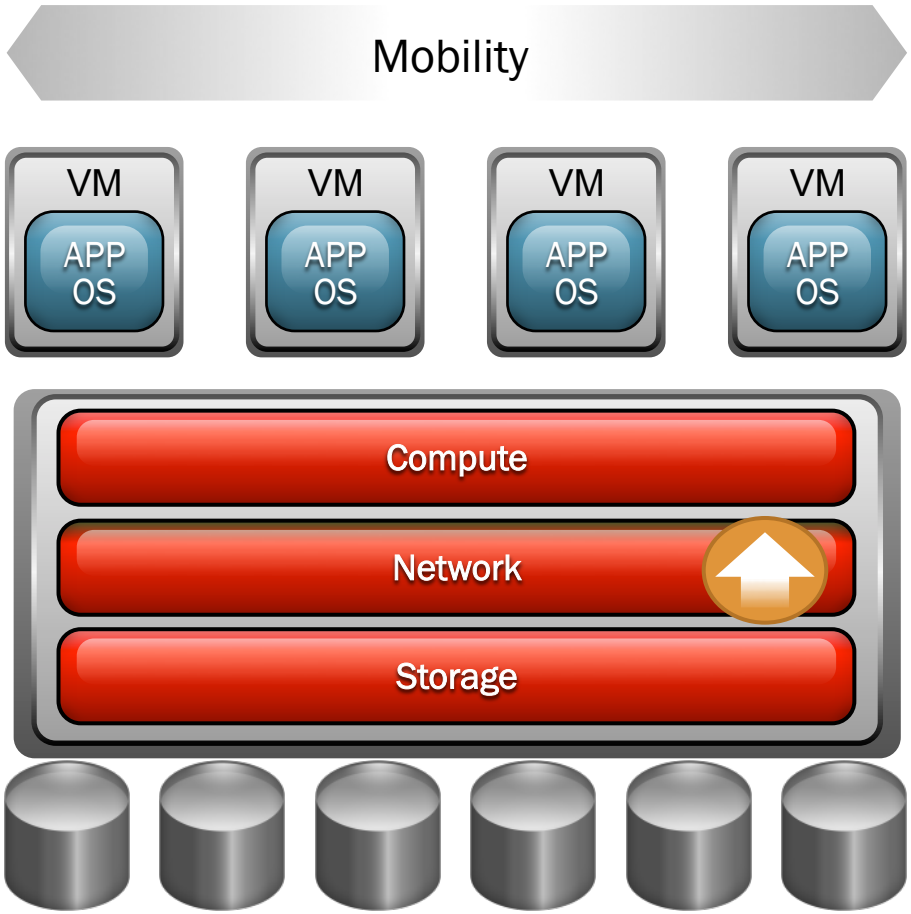

Challenges

Network performance/scalability constraints

Application resiliency and performance under load

VM mobility limits

Infrastructure complexity

Management silos

Thank you