



DATA CENTER

Scaling Virtualized Data Centers with Cloud-Optimized Networks

Describes the network challenges in a virtualized data environment and describes reference architecture for a solution that includes VMware software and Brocade VCS technology, also discusses a use case for deploying the new Brocade VDX 6720 Data Center Switch in a VMware environment.



CONTENTS

Executive Summary	3
Introduction	3
Network Challenges	4
Scaling Virtual Server Environments	4
Mobility Inhibitors	5
Network Management	6
Cable Management	8
Brocade and VMware Solution	8
VMware vCenter Server	9
Brocade Virtual Cluster Switch (VCS) Technology	10
Ethernet Fabric	10
Distributed Intelligence	11
Logical Chassis	13
Hardware Features	13
Deploying the Brocade VDX 6720 in a VMware Environment	14
Summary	16

EXECUTIVE SUMMARY

With the introduction of virtualization, Information Technology (IT) has undergone a radical transformation. Virtualization enables server consolidation, application isolation, and workload migration. And data centers in turn have benefited immensely from virtualization, which enables a flexible server infrastructure, high availability, live migration, and faster deployment times.

VMware, the current market leader in virtualization software for x86 platforms, leads the virtualization industry in technology innovation and product development. VMware vSphere, the company's flagship product, delivers the virtualization infrastructure and application services that organizations need to transform their IT strategy and deliver IT as a service. To manage the VMware environment, VMware vCenter Server provides a central management platform that monitors, controls, and provisions server resources in an automatic, secure and efficient fashion. Coupled with VMware vSphere, VMware vCenter Server in tandem with the other VMware products allows IT to build a virtualized data center for delivering cloud services to its users.

Along with the benefits it brings to data centers, virtualization also introduces new challenges to IT organizations, namely: infrastructure scalability, management complexity, Virtual Machine (VM) mobility, and support for convergence. Company executives require IT organization to reap the benefits of virtualization while still meeting or exceeding negotiated service-level agreements. IT organizations need to explore new technology and products to meet the challenges of virtualization.

Brocade® Virtual Cluster Switching (VCS™) technology is designed to meet these challenges. Brocade VCS is a Layer 2 Ethernet technology that incorporates new and emerging standards as well as advanced features to provide better bandwidth utilization, increased network scalability, seamless support for network convergence, and ease of management. At the core of Brocade VCS technology are three “pillars”: Ethernet fabric, Distributed Intelligence, and Logical Chassis.

Ethernet fabric utilizes the emerging IETF [Transparent Interconnection of Lots of Links \(TRILL\)](#) proposal to remove the need for Spanning Tree Protocol (STP) with multiple paths between VCS-compatible switches for [Data Center Bridging \(DCB\)](#) traffics. All paths in the fabric are active, and if a link fails, traffic is automatically distributed to other available equal-cost paths with minimal delay. With *Distributed Intelligence*, configuration and network topology information is automatically distributed to each of the member switches in the fabric. With TRILL, the fabric delivers unmodified frames from a source port to a destination port in the fabric as if the entire fabric were a single logical switch chassis and each Brocade VCS-compatible switch a port module in the chassis. The fabric is scalable to over 1000 ports in the *Logical Chassis*.

Brocade VCS technology enables IT to meet the network challenges of a virtualized environment. Brocade VCS technology in VMware environments also provides IT with a scalable and easy-to-manage infrastructure for cloud computing. This paper will first describe the challenges of the network. A reference architecture solution that includes VMware software and Brocade VCS technologies will be discussed in details. Following the discussion, two use cases are presented: VCS supports for DCB traffics over two cables from each vSphere server; and VCS connects to existing Fibre Channel SAN.

INTRODUCTION

Today's company executives requires their IT organization to deploy new services quickly and efficiently, to provide fast and reliable data access, to meet or exceed stringent service level requirements at reduced costs. In short, IT must move at the speed of business to capitalize on new opportunities and respond to increasing global competition.

Cloud Computing is a new computing paradigm for the data center to meet these IT requirements. It relies on server virtualization to reduce capital and operating costs, to provide a more secure and highly available computing environment, to reduce application deployment time, and to increase capacity elasticity for

changing demands. Server virtualization is a dramatic shift in computing technology where a virtual machine (VM) encapsulates all the software required to provide specific services. VMs decouple hardware from application so that services can run on any available hardware that provides a hypervisor to run the VM. Cloud computing also changes business model of IT requiring service-based computing where shared resources are connected to applications and cost is based on resource consumption, not asset purchase.

However, as IT organizations deploy server virtualization and transition to cloud computing data center managers and administrators are also facing new challenges that the paradigm shift brings with it. To overcome these challenges, IT organizations need to understand the new technology and products, so that they can explore available solutions and implement the best new networking practices. This paper helps IT organizations in their planning for cloud computing by proposing a reference architecture with VMware software products and Brocade VCS technology.

NETWORK CHALLENGES

Among the many challenges that server virtualization introduces, network orchestration and provisioning are becoming high-priority infrastructure components for IT organizations. The network needs to be able to:

- Scale for virtualized environments
- Optimize application mobility
- Provide simplified management
- Reduce costs through simplified cable management

Scaling Virtual Server Environments

When organizations scale virtual server environments, the network presents challenges and limitations, such as the shortcomings of Spanning Tree Protocol (STP), low utilization, and link failure recovery (see Figure 1).

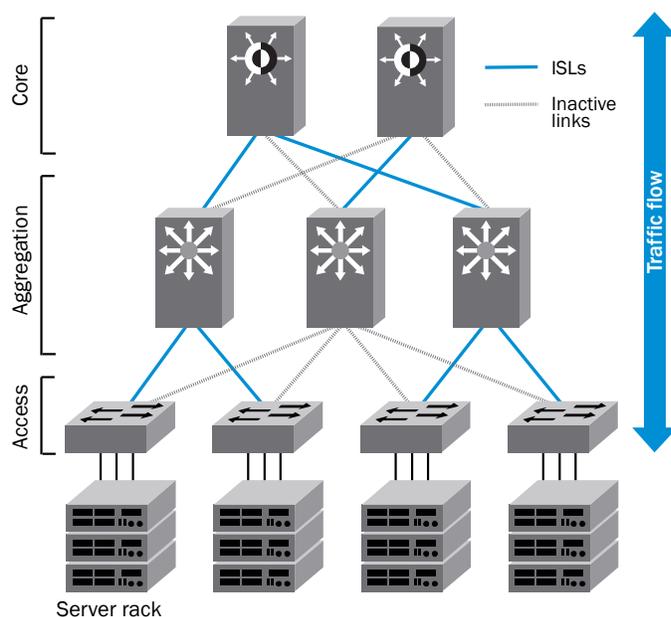


Figure 1. STP designates redundant paths as standby (show with dotted lines), which limits network utilization

In traditional Layer 2 Ethernet networks, to create a highly available network, organization designate paths through the network as active or standby using STP. While this provides an alternate path, *only one path can be used at a time*, which means that network bandwidth cannot be load balanced across multiple paths leading to under-utilization of resources. Since one of the goals of VMware is to increase utilization of the physical server, optimal utilization of network bandwidth should also be expected.

To increase network utilization, Multiple Spanning Tree Protocol (MSTP) and similar protocols allow for separate spanning trees per Virtual LAN (VLAN). While this improves bandwidth utilization, the STP limit of one active path between switches remains. And, because traffic paths are manually configured with MSTP, complexity increases.

Another challenge with STP network behavior occurs when links fail. When failures occur, the spanning tree needs to be redefined and blocks all traffic. Traditional STP can take up to a several minutes, while Rapid Spanning Tree (RSTP) can take about 5 seconds. These convergence times vary greatly depending on the size of the environment and can be very unpredictable. With Gartner predicting that 50 percent of all workload will be virtualized by 2013 and companies moving to the cloud, the demand for non-stop traffic is even more important. In addition, VM consolidation ratios are getting larger and larger every year, thus an STP re-convergence will have an even larger impact of critical workloads that reside on vSphere 4.1 servers. Finally, when spanning tree is re-converging, broadcast storms can occur and result in network slowdown.

In contrast, consider the benefits of a flat Layer 2 network that:

- Is highly available
- Guarantees high-bandwidth utilization over equal-cost paths
- Does not stall traffic when links are added or removed due to failure or network reconvergence
- Makes latency deterministic and is lossless

Mobility Inhibitors

VMware vMotion™ has probably been one of the most significant features of VMware vSphere platform. The ability to migrate a running VM from one physical server to another without downtime has made IT a lot more agile, because it can respond faster to infrastructure demands. Enabling technologies such as vMotion requires that VMs reside within the same IP broadcast domain and Ethernet VLAN, which is required to maintain non-disruptive client traffic to the VM. The ability of the physical network to provide the same characteristic anywhere within the confines of a data center poses a challenge to the mobility of the applications. The fact that you need to reconfigure the network for vMotion to work across the data center results a longer lead time required to scale your business-critical applications, resulting in a negative impact on the efficiency and effectiveness of your IT operations.

The issue of network reconfiguration is exacerbated when two data centers are used as the resource pool for server virtualization. This implies that VMs have the ability to migrate to and from the data centers, presenting the IT network management with yet another challenge of stretching Layer 2 boundaries across data centers. The complexity in extending the networks related to excessive broadcast storms, longer re-convergence time of STP, and increased management overheads have led data center operations personnel to restrict the Layer 2 to the confines of a rack. As companies start bringing new capacity online by adding to their existing data center, building a new data center, or leveraging the cloud, these small L2 subnets will not scale because a VM needs to be able to move without being re-assigned an IP address.



Figure 2. VM network access can break if port configurations are not properly configured on the destination switch.

Now, consider a large flat Layer 2 network that:

- Reduces the physical barriers to enable vMotion
- Has switches network administrators don't need to configure when a VM is deployed
- Allows VMware administrators to deploy VMs at their own pace and not worry about other teams getting in the way

Network Management

Similar to data center LANs today, multi-tier architectures involve considerable complexity (as shown in Figure 3), paired with the long list of Layer 2 and 3 protocols with which administrators have to be familiar. And the network is much more complicated to manage with the introduction of server virtualization and blade servers. The access layer is no longer managed via a single switch, but now stretches from the vSwitch in VMware hypervisor or DVS to the top-of-rack or end-of-row access switch. Each time a new rack of servers is deployed to host VMs, each switching layer has to be configured, driving up cost and complexity. As mentioned above, since VMs are moving around the data center, each physical switch needs to have the ACLs, QoS, and security policies of each VM. So network administrators are constantly updating switch configurations to support the virtualized environment.

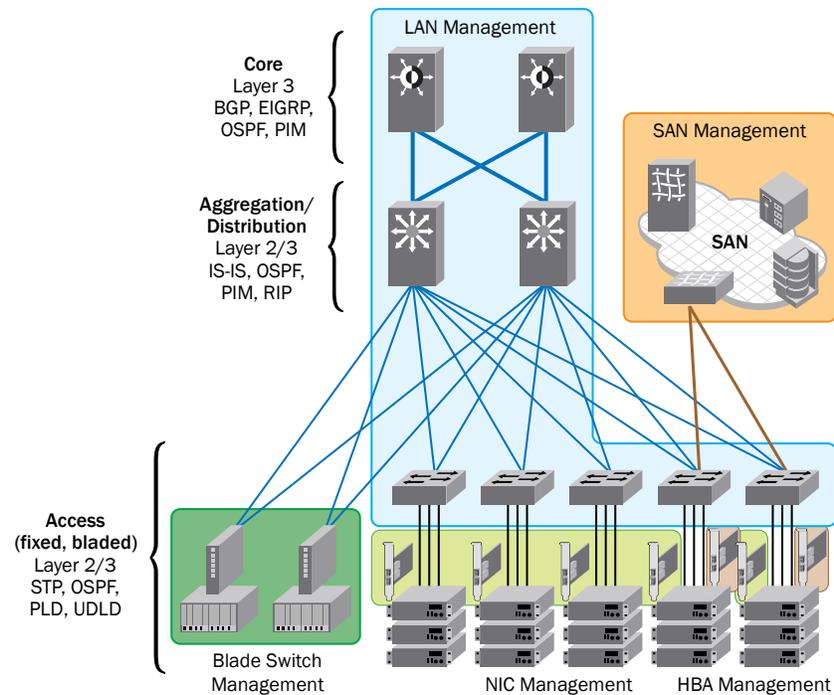


Figure 3. Multi-tier network architectures and many Layer 2 and 3 protocols increase complexity and drive up management costs.

Contributing to management complexity, separate tools are needed to manage the LAN, SAN, servers, NICs, HBAs, and vCenter. Typically the individuals who specialize in each tool don't have a view of the overall environment. Troubleshooting performance problems in such an environment is very challenging and involves two or more administrators to resolve the issue. However, VMware has provided a framework of Application Program Interfaces (APIs) with which vendors can create plugins to give VMware administrators visibility into the physical infrastructure. Currently Brocade has plugins that provide this visibility for both the Fibre Channel SAN and server load balancers.

Brocade network management offers the following advantages to administrators:

- Logically eliminates the management of multiple switching layers
- Apply policies and manage traffic across many physical switches as if they were one switch
- Scale network bandwidth without manual reconfiguration of switch ports and network policies
- Provides a single, customized view of network status to server, network, and storage administrators

Cable Management

One of the biggest challenges in virtualized data centers is the amount of cabling required for each vSphere 4.1 server (see Figure 4). Best practices dictates isolating the different types of traffic, so vMotion traffic, fault tolerance, IP storage, and management, each requires its own dedicated Network Interface Card (NIC)—and in most cases multiple NICs—for redundancy. In addition, a few FC Host Bus Adapters (HBA) are required to connect to Fibre Channel storage. You might need 6 to 10 cables for each vSphere server. However, in most cases these NICs are underutilized, which made configuring each vSphere server more complicated and increased both capital and operating costs.

Cable management issues could be resolved if you could:

- Transport IP and mission-critical storage traffic over the same wire
- Consolidate different traffic classes to a few cables with high availability
- Reduce physical port count with faster, more power-efficient switches

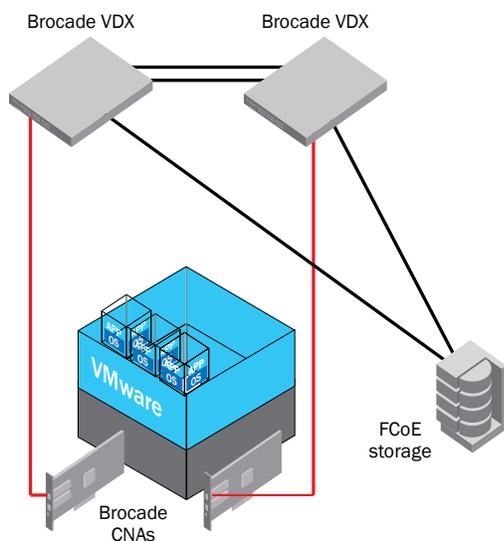


Figure 4. VMware vSphere servers running all traffic over two Brocade CNAs

BROCADE AND VMWARE SOLUTION

Figure 5 shows a reference architecture that addresses the network challenges in a virtualized data center. The architecture includes Brocade VCS technology, VMware vSphere, and VMware vCenter. In Figure 5, a fabric in the access layer connects server and storage to the aggregation layers of the network. The fabric is network convergence ready and IP and storage traffic share the same cable. VMware vCenter is included to manage server, storage, and network resource utilization in the virtualized environment.

The resulting architecture reflects a network that is:

- Efficient in network bandwidth utilization
- Highly available with no single point of failure and fast failover
- Easy to manage as the fabric behave as if it were a single logical switch chassis
- Scalable to over 100 ports per fabric
- Capable of supporting live VM migration to any server in the fabric
- Able to provide fast link reconvergence as STP is no longer required

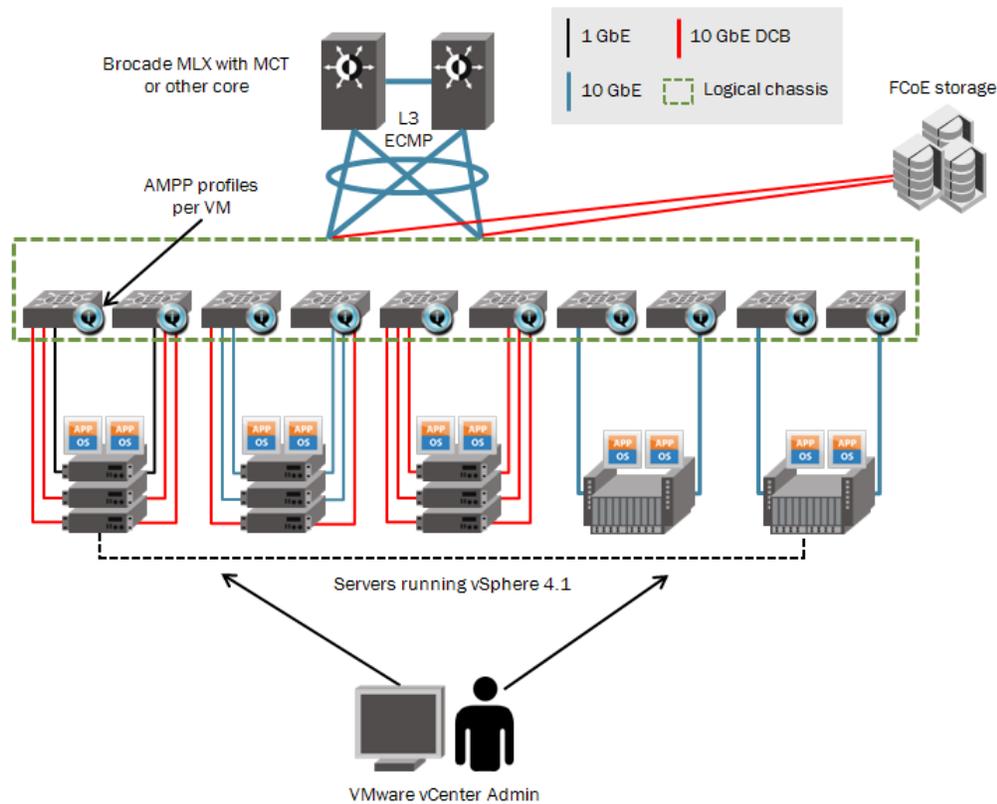


Figure 5. Brocade VCS Ethernet fabric with VMware vCenter™

VMWARE vCENTER SERVER

VMware vCenter Server provides a scalable and extensible platform that forms the foundation for [virtualization management](#). VMware vCenter Server centrally manages [VMware vSphere](#) environments allowing IT administrators dramatically improved control over the virtual environment compared to other management platforms. The key features of vCenter Server include:

- Real-time monitoring of dynamic virtual elements
- Customizable alarm triggers
- Simplified navigation and inventory search
- Rapid provisioning and simplified patch management
- Dynamic resource allocation to ensure Service Level Agreements (SLAs)
- Workflow automation
- High availability
- Scalable architecture

The VMware vCenter Server platform includes several key components that work together to give enterprises a scalable virtualization management hub. Management servers provide central management points for hosts and virtual machines, with inventory and performance information stored in a database. A VMware vCenter agent provides connectivity between the host and management server. Administrators can access VMware vCenter Server from the VMware vSphere client running on any Windows PC, or use the vCenter Web Access portal for remote access from any Web browser. Roles and permissions are replicated across management servers, giving administrators the ability to manage multiple instances of VMware vCenter Server from one console, and a search engine makes it possible to quickly find VMs, hosts, or other inventory objects located anywhere in the enterprise.

Brocade Virtual Cluster Switch (VCS) Technology

Brocade Virtual Cluster Switch (VCS™) is a revolutionary Layer 2 Ethernet technology, purpose built for virtual data centers, which improves network utilization, maximizes application availability, increases scalability, and dramatically simplifies the network architecture in next-generation VMware data centers. Features include: Ethernet fabric, Distributed Intelligence, and Logical Chassis, as shown in Figure 6.

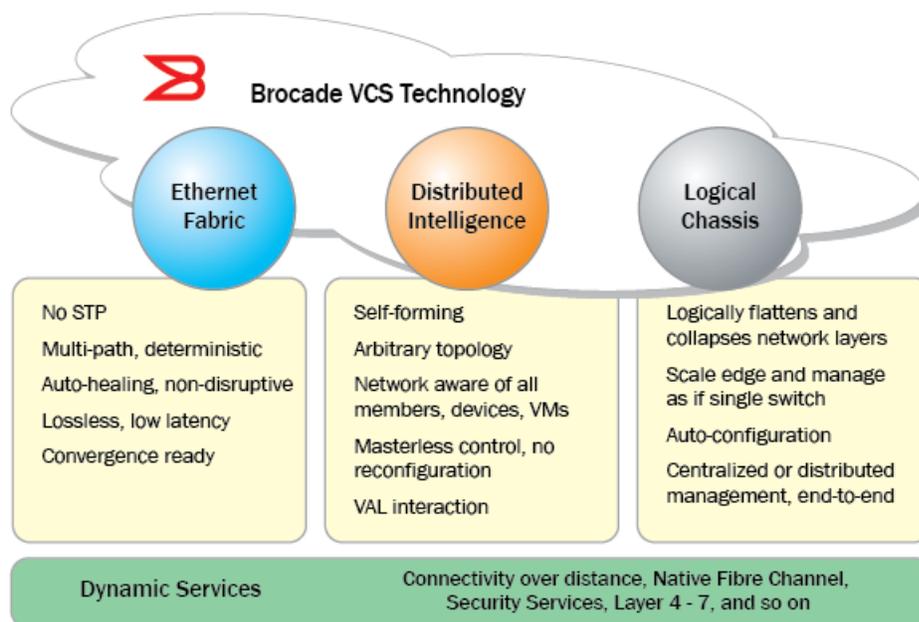


Figure 6. Brocade VCS technology components

Ethernet Fabric

Brocade has pioneered the development, architecture, and deployment of network fabric technology in the data center. In fact, Brocade SAN fabric technology is currently deployed in over 90 percent of the Global 1000 data centers. Now Brocade is bringing the same level of innovation to the data center LAN, combining Ethernet and Brocade fabric technology.

One of the biggest features of STP is no longer necessary, because the Ethernet fabric appears as a single logical switch to connected servers, devices, and the rest of the network. This allows organizations to start flattening network layers and scaling their Layer 2 network at the edge, which in turn makes deploying the VMware environments simpler.

The Brocade Ethernet fabric is an advanced multi-path network utilizing an emerging standard called Transparent Interconnection of Lots of Links (TRILL). Unlike STP, with TRILL all paths in the network are active and traffic is automatically distributed across equal-cost paths. In this optimized environment, traffic automatically takes the shortest path for minimum latency without manual configuration.

Multi-Chassis Trunking (MCT) capabilities in access switches enable a logical one-to-one relationship between the access (with Brocade VCS technology) and core layers of the network. Multi-Chassis Trunking is a Brocade technology that allows multiple switches to appear as single logical switch connecting to another switch using a standard LAG. Since the technology is an enhancement to the standard LAG protocol, a single MCT-unaware server or switch using a standard LAG trunk can connect to two MCT-aware switches—and the traffic is dynamically load balanced.

MCT is beneficial to VMware administrators because now they can provide high availability of hardware when using “Route based on IP Hash” to provide link aggregation between virtual switches and physical switches. This has been a challenge for many users because many of their upstream switches didn’t support MCT, so they weren’t able to provide link aggregation to multiple switches.

If links are added, removed, or fail, then there is no disruption to the Ethernet fabric and does not require all traffic in the fabric to stop. If a single link fails, traffic is automatically rerouted to other available paths in less than a second. Moreover, single component failures do not require the entire fabric topology to re-converge, helping to ensure that no traffic is negatively impacted by an isolated issue (see Figure 7).

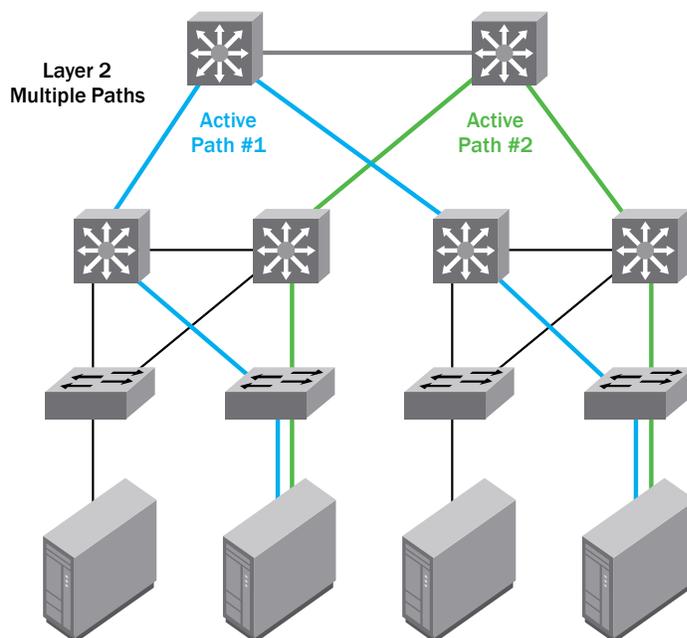


Figure 7. TRILL enables multiple active paths through the Ethernet fabric

The Brocade Ethernet fabric is designed to be network-convergence ready with advanced Ethernet technologies for greater utilization and performance. With built-in DCB capabilities, the Ethernet fabric is lossless, making it ideal for Fibre Channel over Ethernet (FCoE) and iSCSI storage traffic, while enabling LAN and SAN convergence. The combination of TRILL and DCB features will enable multi-hop convergence of IP and storage traffic. This will allow VMware administrators to reduce the number of cables associated with each vSphere server from the typical 6 to 10 cables to a minimum of the 2 cables required for high availability. This will reduce cost, while increasing performance of the VMs.

Distributed Intelligence

With Brocade VCS technology, all configuration and destination information is automatically distributed to each member switch in the fabric. For example, when a server connects to the fabric for the first time, all switches in the fabric learn about that server. In this way, fabric switches can be added or removed and physical or virtual servers can be relocated—without the fabric requiring manual reconfiguration.

Distributed Intelligence, shown in Figure 8, allows the Ethernet fabric to be “self-forming.” When two Brocade VDX 6720 switches are connected, the fabric is automatically created, and the switches discover the common fabric configuration. Scaling bandwidth in the fabric is as simple as connecting another link between switches or adding a new switch as required.

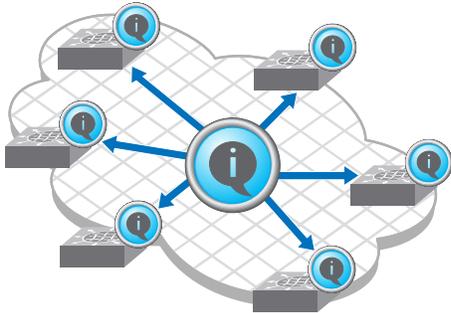


Figure 8. Fabric configuration and end device information is automatically distributed throughout the fabric

The Ethernet fabric does not dictate a specific topology, so it does not restrict oversubscription ratios. As a result, network architects can create a topology that best meets specific VMware virtualization requirements. Unlike other technologies, VCS enables different end-to-end subscription ratios to be created or fine-tuned as application demands change over time.

And, unlike switch stacking technologies, the Ethernet Fabric is masterless. This means that no single switch stores configuration information or controls fabric operations. Any switch can fail or be removed without causing disruptive fabric downtime or delayed traffic while a new master switch is selected.

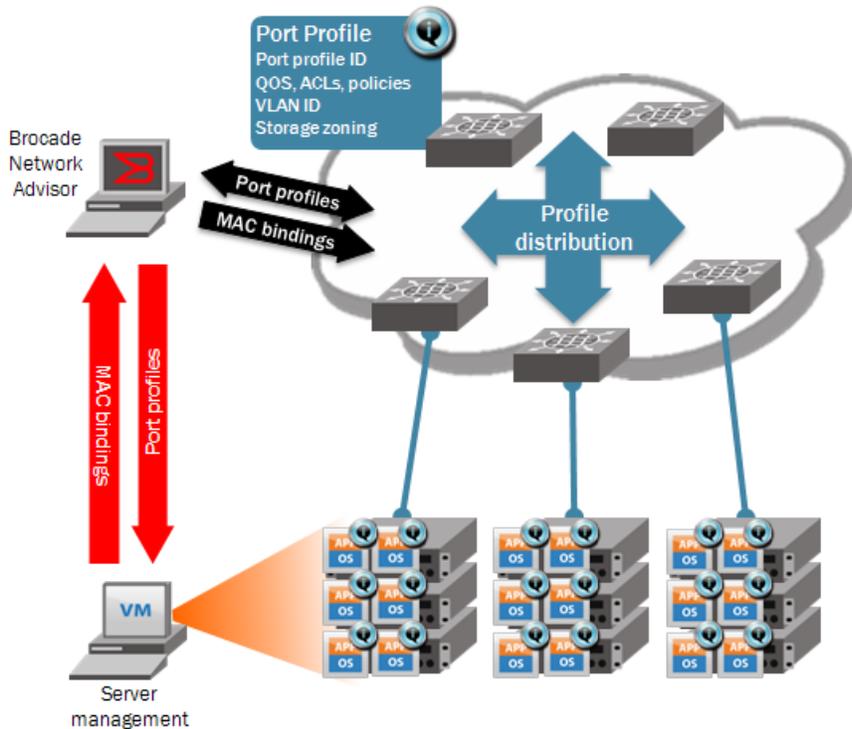


Figure 9. Port profiles follow the VM as it migrates from one switch to another

Logical Chassis

All switches in an Ethernet Fabric are managed as if it were a single logical chassis. To the rest of the network, the fabric looks no different than any other Layer 2 switch. The network sees the fabric as a single switch (shown in Figure 10), whether the fabric contains as few as 48 ports or thousands of ports.

Each physical switch in the fabric is managed as if it were a port module in a chassis. This enables fabric scalability without manual configuration. When a port module is added to a chassis, the module doesn't need to be configured—a switch can be added to the Ethernet fabric just as easily. When a Brocade VDX 6720 switch is connected to the fabric, it inherits the configuration of the fabric and the new ports become available immediately.

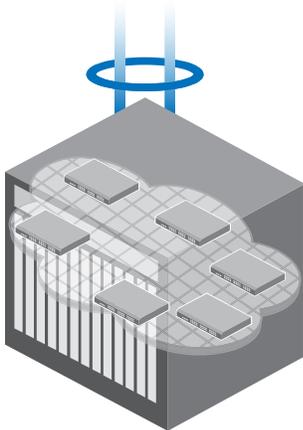


Figure 10. Configuration is accomplished via a single IP address

The Ethernet fabric scales to over 1000 ports per Logical Chassis. Consequently, Brocade VCS technology removes the need for separate aggregation switches, because the fabric is “self-aggregating.” This enables a flatter network architecture, dramatically reducing cost and management complexity. VMware and network administrators have been requesting this because too many switches became as a roadblock in achieving all their virtualization objectives, such as the distance they could migrate a VM. Network architects will be able to move to core/edge architectures, simplifying their designs while reducing capital and operating expenses.

The Logical Chassis capability significantly reduces management of small-form-factor edge switches. Instead of managing each top-of-rack switch (or switches in blade server chassis) individually, organizations can manage them as one Logical Chassis, which further optimizes the network in the virtualized data center and allow companies to scale their private and public clouds as business grows.

Hardware Features

Brocade has a long-standing tradition of building the most robust switches in the industry. They have built their switches by providing best-in-class performance and density, environment flexibility, network convergence, and highly resilient and efficient design. The first product released under the Brocade VCS architecture is the Brocade VDX 6720 in 24-port and 60-port models. Brocade VDX 6720 switches, seen in Figure 11, was built on a foundation of Brocade technology leadership with future-looking features—available now.

These features include:

- Best-In-Class performance and density
 - 24- and 60-port models with Ports On Demand
 - Non-blocking, cut-through architecture at wire speed

- 600 ns port-to-port latency; 1.8 us across port groups
- 80 Gigabits per second (Gbps) Inter-Switch Links (ISLs) between switches
- Environmental flexibility
 - 10 and 1 Gigabit (GbE) supported on every port
 - Twinax, direct-attached optical, and SFP optical connectivity options
 - Switch depth less than 17 inches 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
 - Simple design, optimal power efficiency



Brocade 6720-24 Data Center Switch



Brocade 6720-60 Data Center Switch

Figure 11. Brocade VDX™ Data Center Switches, 24-port and 60-port models

DEPLOYING THE BROCADE VDX 6720 IN A VMWARE ENVIRONMENT

In this configuration shown in Figure 12 the Brocade VDX switches form an Ethernet fabric, which flattens and seamlessly scales out Layer 2 layer at the edge. The servers that run VMware vSphere connect to the switches with two separate cables to provide full redundancy. The fabric connects to the core router with multiple links for high availability and better performance. VMware vCenter Server monitors, manages, and provisions server resources in the fabric.

All IP and storage traffics from each vSphere server are load balancing over two Brocade VDX 6720 switches to take advantage of converged I/O with improved performance. This allows IT to eliminate the complex cabling for each vSphere 4.1 server required in today's VMware environments. It also provides maximum storage flexibility by allowing IT to deploy their preferred IP and network storage protocols as they wish, such as FCoE, iSCSI, or Network-Attached Storage (NAS). Each vSphere server typically has several NICs and FC HBAs to allow for separation of traffic classes such as vMotion, VM traffic, and storage traffic over different networks. This configuration allows for companies to reduce the amount of cables, the number of network ports, and HBAs, all the while providing high-speed performance for both IP and FCoE traffic.

All Brocade VDX 6720 switches in the configuration appear to the network administrator as one Logical Switch, which can include up to 10 Brocade VDX switches. More than 300 Brocade VDX ports are available in a flat Layer 2 network for the configuration to scale. Coupled with VMware vMotion and VMware vNetwork Distributed Switch (VDS), the fabric gives IT the flexibility to scale out the configuration as business demands increase.

With the release of VMware vSphere 4.1, customers can now take advantage of the new features such as Network I/O control and Storage I/O control. These features allow companies to extend the benefits of the convergence with Brocade VDX platforms by allowing them to prioritize their traffic within vSphere.

With VMware Network I/O Control, rules and policies can be configured to specify the business priority of each VM. When I/O congestion is detected, Network I/O Control dynamically allocates the available I/O resources to flow types according to your rules, enabling you to:

- Improve service levels for critical applications
- Virtualize more types of workloads, including I/O-intensive, business-critical applications
- Ensure that each cloud tenant gets their fair share of I/O resources
- Increase administrator productivity by reducing the amount of active performance management required
- Increase flexibility and agility of the infrastructure by reducing the need for storage volumes dedicated to a single application

The use case, shown in Figure 12, allows customers to take advantage of the new and exciting features of Brocade VDX 6720 switches in a virtualized data center running VMware vSphere and vCenter server. In addition, deploying the Brocade VDX 6720 does not require rip and replace of existing IT setup. The combination of Brocade VCS technology and the Brocade VDX 6720 platform allow organizations to easily scale out their virtualized environments while gaining performance, scalability, and taking advantage of features built into vSphere 4.1 such as Network I/O and Storage I/O control.

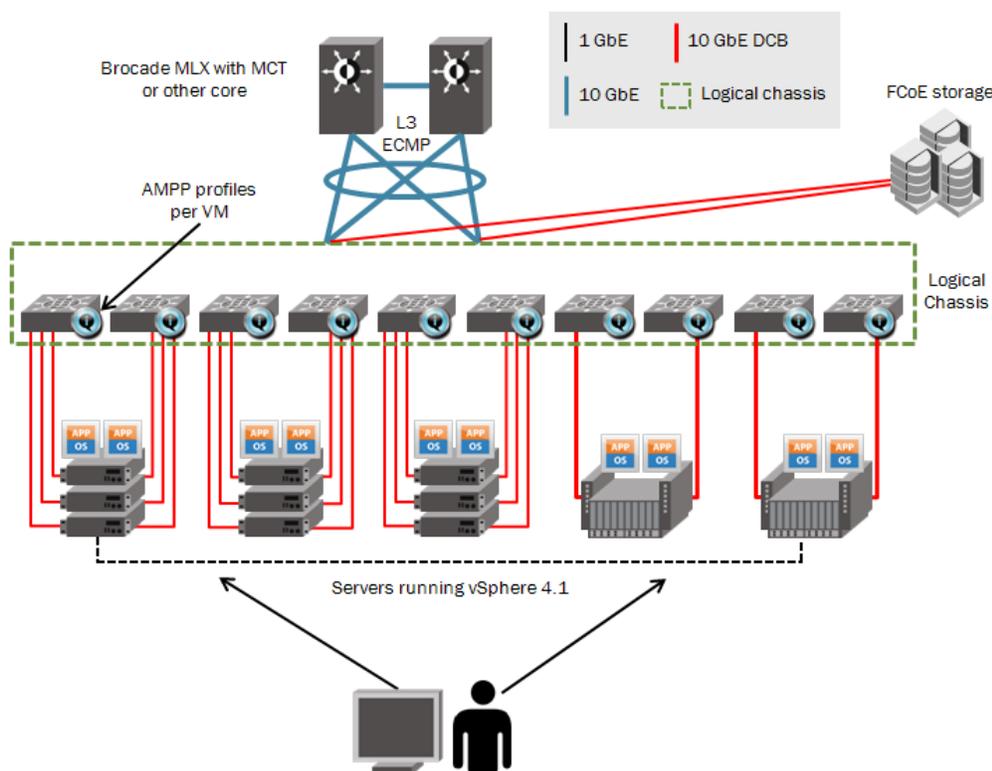


Figure 12. VCS supports DCB traffics over two cables from each vSphere server

SUMMARY

As IT organizations adopt server virtualization and cloud computing, they also face a set of challenges in deploying the new computing paradigm. These challenges include scaling the network for virtualization, VM mobility, management complexity, and support for convergence. This paper presents a reference architecture that leverages Brocade VCS technology in a VMware environment to address these new challenges. The architecture includes VMware vSphere™, VMware Center™ Server, Brocade VCS technology, and Brocade VDX 6720 Data Center Switches.

Brocade VCS technology creates an Ethernet fabric, allowing many switches to concurrently communicate at high bandwidths and low latency in a well-managed environment running DCB traffic. Servers and storage devices can use multiple paths through the Ethernet fabric. The multiple Ethernet fabric paths through the network are used for both fault tolerance and load balancing existing bandwidth. Brocade VCS technology automatically scales when another Brocade VDX 6720 switch is added to the fabric. The new switch logically “plugs” into the fabric and its configuration parameters are shared across all Brocade VCS-compatible switches in the fabric. Thus, network administrators no longer have to manually configure the switch. Brocade VCS technology supports AMPP to ensure that switch port profiles of a VM are the same after live migration of the VM with VMware vMotion.

VMware vSphere provides the virtualization platform in the architecture. Data center administrators use VMware vCenter Server to monitor, manage, and provision server resources. IT organizations can now deliver the virtualized data center as cloud services to its customers. The resulting architecture is a secure private cloud that enables IT organization to consolidate servers, isolate workloads, and provide capacity elasticity to handle demand changes while meets challenges of the network.

© 2010 Brocade Communications Systems, Inc. All Rights Reserved. 11/10 GA-TB-353-00

Brocade, the B-wing symbol, BigIron, DCFM, DCX, Fabric OS, FastIron, IronView, NetIron, SAN Health, ServerIron, Turbolron, and Wingspan are registered trademarks, and Brocade Assurance, Brocade NET Health, Brocade One, Extraordinary Networks, MyBrocade, VCS, and VDX 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.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.