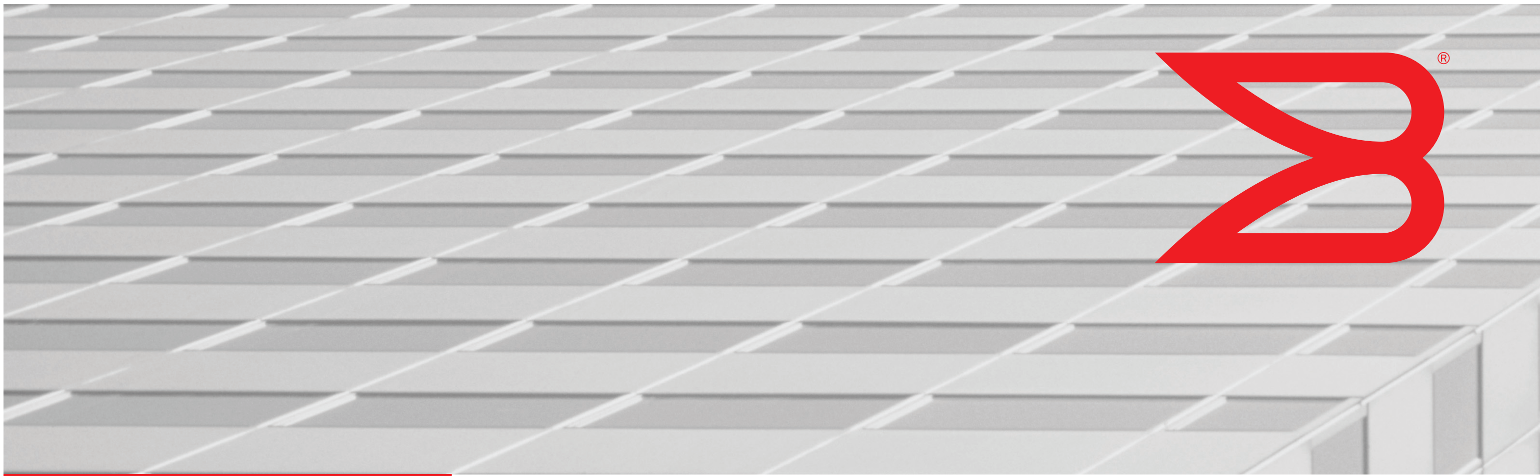

DATA CENTER

Introducing Brocade VCS Technology

Brocade VCS technology is designed to revolutionize the way data center networks are architected and how they function.

BROCADE

Not that long ago, information was stored predominately in text form. Now, graphical data is combined with audio and video, and all of it needs to be stored, indexed, and archived. Consumer demand is also increasing, and current economic conditions make it challenging to balance budgets. End users demand fast and reliable access to information around the clock. Business leaders require that IT functions meet and exceed stringent service levels with minimum downtime. In turn, IT must be able to move at the speed of business to capitalize on new opportunities and respond to increasing global competition.

Brocade® VCS™ technology is designed to meet these challenges by enabling next-generation virtual data center and private cloud computing initiatives. VCS technology comprises three main technology pillars: Ethernet Fabric, Distributed Intelligence, and Logical Chassis. Dynamic Services extends the capabilities of the VCS architecture, providing the highest level of functionality and investment protection. This paper reviews key data center network challenges, including server virtualization, and how VCS technology addresses them.

INTRODUCTION

Data centers continue to grow as digital assets increase and more applications are deployed. Businesses expect agile application deployment—in minutes, not months—to keep their competitive edge as markets and competitors become global in scale. And data center resources such as rack space, power, and cooling are growing more scarce and costly. For these reasons, IT organizations are aggressively deploying server virtualization in data centers to consolidate applications and improve resource utilization. The cost savings of server virtualization resulting from increased asset utilization, higher availability, and on-demand application deployment have met the business mandate to "do more with less." However, the underlying limitations in current network technologies have often prevented organizations from meeting the performance, availability, security, and mobility requirements of server virtualization.

These business pressures on the data center are significant, especially for the networking team, but there is a ray of light for data center architects thanks to significant technology innovation and price reductions. For example, network technologies such as 10 Gigabit Ethernet (GbE) and Brocade VCS technology will allow data centers to meet these challenges. See Figure 1 for a forecast of 10 GbE port growth (Dell'Oro Group, "Ethernet Switch Report, 4Q09," February 2010).

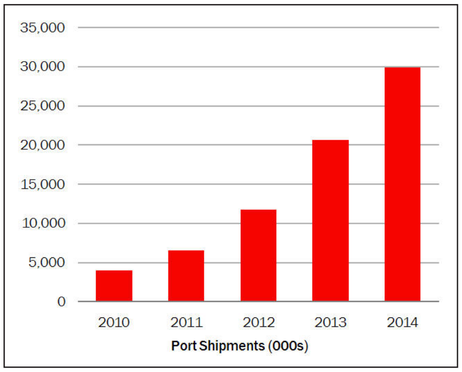

Figure 1.
Increase in 10 GbE ports deployed in large part for server and storage connectivity.

NETWORK CHALLENGES TODAY

Among the other challenges IT groups face, they need to scale virtual server environments, provide application mobility, and handle infrastructure complexity and management overhead.

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, shown in Figure 2), the growing number of GbE connections per server, low utilization, and link failure recovery.

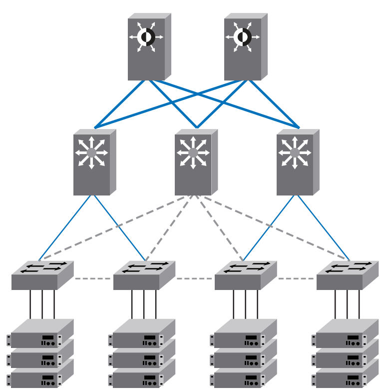

Figure 2.
STP designates redundant paths as standby (shown with dashed lines), which limits network utilization.

Enabling virtualization capabilities, such as Virtual Machine (VM) mobility, requires VMs to migrate within a single Layer 2 network, since non-disruptive migration of VMs across Virtual LANs (VLANs) using Layer 3 protocols is not supported by virtualization hypervisors. In traditional Layer 2 Ethernet networks, to create a highly available network, organizations 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 is not well utilized. Since one of the goals of server virtualization is to increase utilization of the physical server, increased 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 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 is network behavior when links fail. When failures occur, the spanning tree needs to be redefined. This can take anywhere from five seconds with Rapid Spanning Tree (RSTP) up to several minutes with STP—and this convergence can vary unpredictably even with small topology changes. The demands for non-stop traffic flow increases with server virtualization, and consequently network convergence times have to shrink. STP does not provide an adequate solution for these requirements. Finally, when a spanning tree is reconverging, broadcast storms can occur and result in network slowdown. All of these limitations of STP are why Layer 2 networks are typically kept small in the data center.

In contrast, consider the benefits of a 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 reconfiguration
- Makes latency deterministic and is lossless
- Can transport IP and mission-critical storage traffic over the same wire

Brocade VCS will enable efficient scaling of virtual server environments without VM mobility constraints and potential network downtime.

Application Mobility

When an application is running in a VM rather than on a physical server, it is not tied to a specific physical server. This allows a VM to move between physical servers when application demands change, when servers need to be maintained, and when a quick recovery from site disasters is necessary.

VM mobility can occur within a cluster of physical servers that are in the same IP subnet and Ethernet VLAN. This is required for the migration to be non-disruptive to client traffic as changes in the IP subnet are necessarily disruptive. As described in the review of STP limitations, the sphere of VM migration can be further constrained. The solution for flexible VM mobility is a more scalable and available Layer 2 network with higher network bandwidth utilization.

For a VM to migrate from one server to another, many server attributes must be the same on the origination and destination servers. This extends into the network as well, requiring VLAN, Access Control List (ACL), Quality of Service (QoS), and security profiles to be the same on both the source and destination access switch ports. If switch port configurations differ, either the migration pre-flight will fail or network access for the VM will break, as shown in Figure 3. Organizations could map all settings to all network ports, but that would violate most networking and security best practices. The distributed virtual switch in VMware vSphere 4 addresses some of these issues, but at the cost of consuming physical server resources for switching, added complexity in administering network policies at multiple switch tiers, and a lack of consistent security enforcement for VM-to-VM traffic.

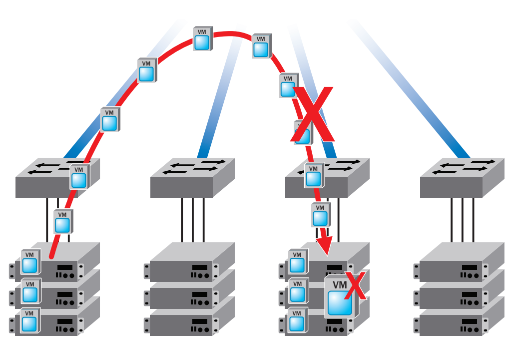

Figure 3.
VM network access can break if port configurations are not properly configured on the destination switch.

With automated VM migration, network administrators will have limited visibility to the location of applications. This makes troubleshooting a challenge, and pinpointing issues to a specific VM will be like finding a needle in a haystack.

Now, consider again a Layer 2 network that:

- Places no physical barriers in the way of VM migration
- Is aware of VM locations and consistently applied network policies
- Does not require manual intervention when a VM moves
- Removes the overhead of switching traffic from the hypervisor for maximum efficiency and functionality
- Supports heterogeneous server virtualization in the same network

Brocade VCS technology allows organizations to broaden the sphere of application mobility, provide VM awareness, and optimize server resources for applications.

Network Management

Similar to data center LANs today, multi-tier architectures involve considerable complexity (as shown in Figure 4), paired with the long list of Layer 2 and 3 protocols with which administrators have to be familiar. And the network has gotten 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 includes multiple stages of switching that extend from the software switch in the hypervisor (called a "softswitch") 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.

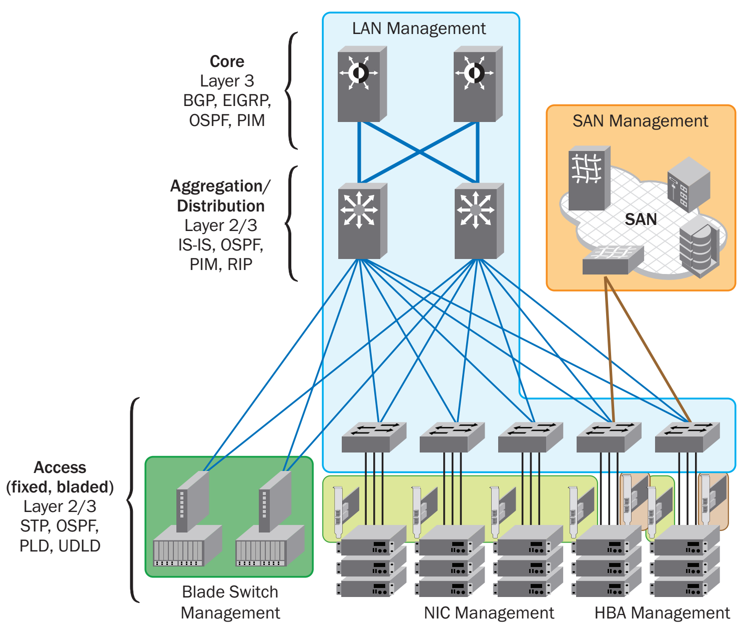

Figure 4.
Multi-tier network architectures and many Layer 2 and 3 protocols increase complexity and drive up management costs.

Contributing to management complexities are the separate tools used to manage the LAN, SAN, blade server connectivity, Network Interface Cards (NICs), and Host Bus Adapters (HBAs). Often administrators can see only what is in their direct line of responsibility and do not have the overall view of the entire network environment.

Now, imagine if they could:

- Logically eliminate 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
- Provide a single, customized view of network status available to server, network, and storage administrators

Brocade VCS technology allows organizations to simplify network architecture, more rapidly scale the network, and significantly reduce management overhead.

BROCADE VCS TECHNOLOGY

Brocade VCS technology is a revolutionary Layer 2 Ethernet capability that improves network utilization, maximizes application availability, increases scalability, and dramatically simplifies the network architecture in next-generation virtualized data centers. As shown in Figure 5, VCS technology comprises three pillars of innovation technology: Ethernet Fabric, Distributed Intelligence, and Logical Chassis. The VCS architecture is also designed to incorporate a set of Dynamic Services for the highest level of functionality and investment protection for data centers, making it a core building block for virtualizing data center networks.

Figure 5. Brocade VCS technology.

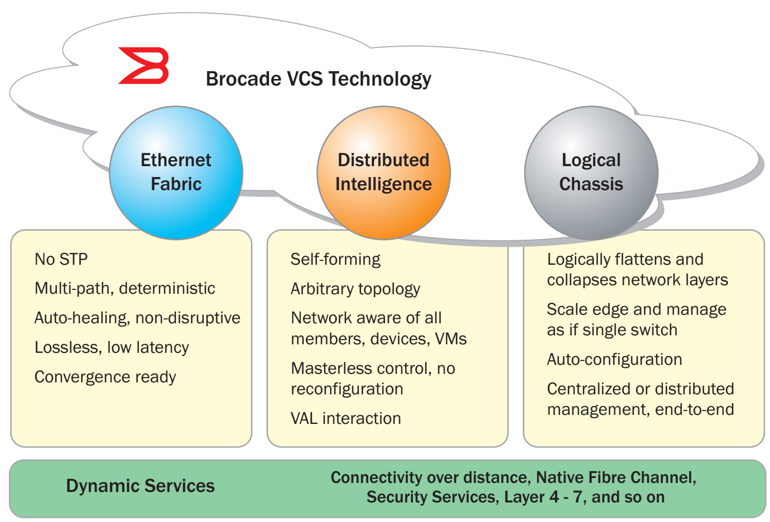

Ethernet Fabric

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.

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. Also, Multi- Chassis Trunking (MCT) capabilities in aggregation switches enable a logical one-to-one relationship between the access (VCS technology) and aggregation layers of the network.

The 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 the equal-cost paths. In this optimized environment, traffic automatically takes the shortest path for minimum latency without manual configuration.

Events such as added, removed, or failed links are not disruptive to the Ethernet Fabric and do 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 reconverge, helping to ensure that no traffic is negatively impacted by an isolated issue (see Figure 6).

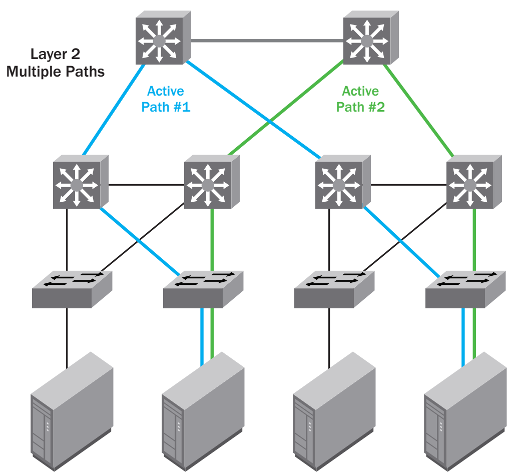

Figure 6.
TRILL enables multiple active paths through the Ethernet fabric.

The Ethernet Fabric is designed to be network-convergence-ready and to include advanced Ethernet technology for greater utilization and performance. With built-in Data Center Bridging (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 for Tier 2 and 3 applications. The combination of TRILL and DCB features will enable multi-hop convergence of IP and storage traffic.

Distributed Intelligence

With 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 7, allows the Ethernet Fabric to be "self-forming." When two switches with VCS capabilities 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 7.
Fabric configuration and end device information is automatically distributed throughout the fabric.

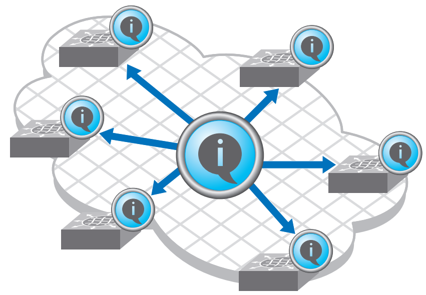

The Ethernet Fabric does not dictate a specific topology, so it does not restrict oversubscription ratios. As a result, network architects can create a topology that best meets specific application requirements. Unlike other technologies, VCS technology enables different end-to-end subscription ratios 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.

Distributed Intelligence also supports a more virtualized access layer. Instead of distributed software switch functionality residing in the virtualization hypervisor, access layer switching is performed in the switch hardware. This approach improves performance, helps ensure consistent and correct security policies, and simplifies network operations and management. Automatic Migration of Port Profiles (AMPP) supports VM migration to another physical server, ensuring that the source and destination network ports have the same configuration for the VM. This key technology helps enable Brocade Virtual Access Layer (VAL) capabilities (see the white paper, *A Vision for the Virtual Access Layer in the Data Center,* on www.brocade.com).

Logical Chassis

All switches in an Ethernet Fabric are managed as if they 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, whether the fabric contains as few as 48 ports or thousands of ports, as shown in Figure 8.

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 does not need to be configured, and a switch can be added to the Ethernet Fabric just as easily. When a switch with VCS capabilities is connected to the fabric, it inherits the configuration of the fabric and the new ports become available immediately.

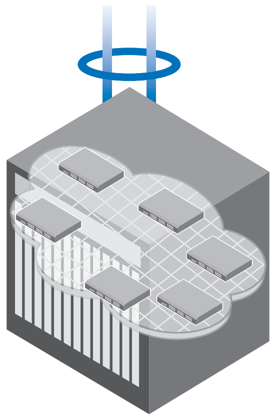

Figure 8.
The Ethernet fabric is managed as if it were a single logical chassis and appears as a Layer 2 switch to the rest of the network.

The Ethernet Fabric is designed to scale to over 1,000 ports per Logical Chassis. Consequently, 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. 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 will further enable a cloud computing model.

Dynamic Services

Dynamic Services extends the capabilities of VCS technology for maximum investment protection and to incrementally incorporate new network services. A Dynamic Service behaves like a special service module in a modular chassis. Examples of these services include fabric extension over distance, native Fibre Channel connectivity, Layer 4 - 7 services such as the Brocade Application Resource Broker, and enhanced security services such as firewalls and data encryption. Switches with these unique capabilities can join the Ethernet Fabric, adding a network service layer that is available across the entire fabric, as shown in Figure 9.

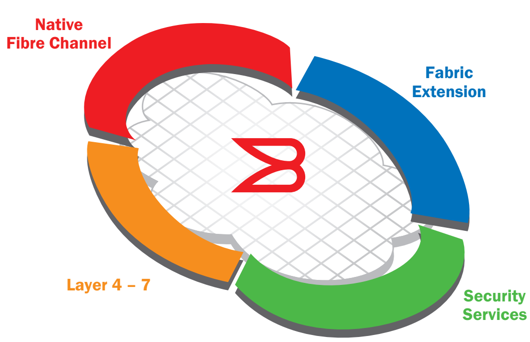

Figure 9.
Dynamic Services extends the capabilities of VCS technology.

THE BROCADE VCS ARCHITECTURE

The VCS architecture, shown in Figure 10, flattens the network by collapsing the traditional access and aggregation layers. Since the fabric is self-aggregating, there is no need for aggregation switches to manage subscription ratios and provide server-to-server communication. For maximum flexibility of server and storage connectivity, multiple protocols and speeds are supported—1 GbE, 10 GbE, 10 GbE with DCB, and Fibre Channel. Since the Ethernet Fabric is one Logical Chassis with Distributed Intelligence, the VM sphere of mobility spans the entire VCS. Mobility extends even further with the VCS fabric extension Dynamic Service.

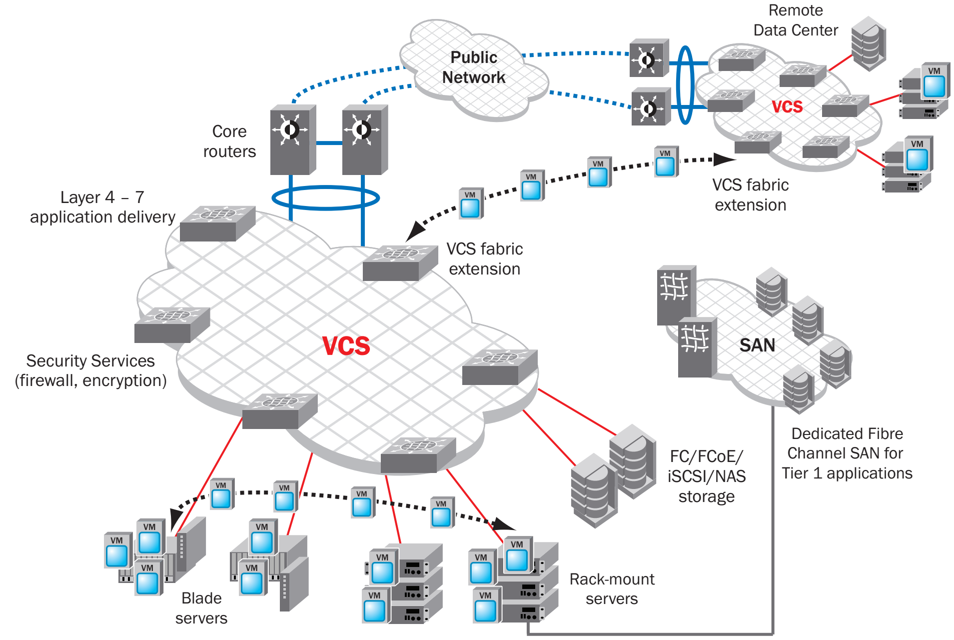

Figure 10.
VCS reference architecture.

At the core of the data center, routers are virtualized using MCT and provide high-performance connectivity between Ethernet Fabrics, inside the data center or across data centers.

Servers running high-priority applications or other servers requiring the highest block storage service levels connect to the SAN using native Fibre Channel. For lower-tier applications, FCoE or iSCSI storage can be connected directly to the Ethernet Fabric, providing shared storage for servers connected to that fabric.

BROCADE DATA CENTER VISION

Brocade VCS technology is one of the new and exciting innovations in the strategic Brocade data center vision. It allows the edge network to be more virtualized, with greater efficiency and agility. VCS technology, Brocade Virtual Access Layer (VAL) technology, and management orchestration initiatives, comprise the data center vision, as shown in Figure 11.

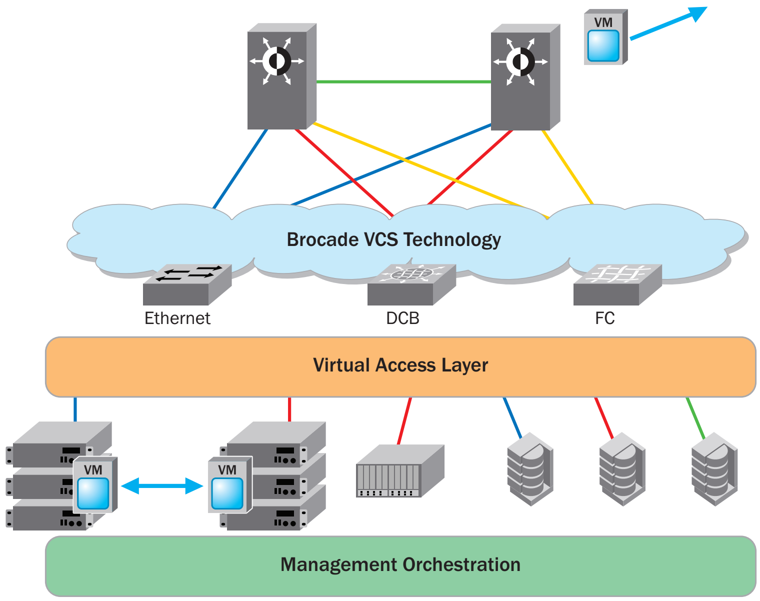

Figure 11.
The strategic Brocade data center vision.

When this new VCS intelligent edge combines with the high-performance, multi-protocol Brocade core, it creates a revolutionary network for next-generation data centers and cloud architectures.

ABOUT BROCADE

Brocade provides innovative, end-to-end network solutions that help the world's leading organizations transition smoothly to a virtualized world where applications and information can reside anywhere. These solutions deliver the unique capabilities for a more flexible IT infrastructure with unmatched simplicity, non-stop networking, optimized applications, and investment protection. As a result, organizations in a wide range of industries can achieve their most critical business objectives with greater simplicity and a faster return on investment.

For more information about Brocade products and solutions, visit www.brocade.com.

Corporate Headquarters
San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters
Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters
Singapore
T: +65-6538-4700
apac-info@brocade.com

BROCADE