



DATA CENTER

Architecture Brief: The Brocade Data Center Fabric Reference Architecture for VMware vSphere 4

The guiding principles driving a new architecture for the data center include consolidation, virtualization, service-level resource allocation, and policy-based data management. These principles simplify how business exploits its data for competitive advantage in a cost-effective manner—as detailed in this paper.

BROCADE

CONTENTS

Executive Summary	3
Overview of the Brocade Data Center Fabric Strategy	4
The Brocade Reference Architecture	5
1. VMware vSphere 4.....	6
2. Brocade Data Center Fabric.....	7
3. Blade Server Chassis.....	7
4. Brocade Solutions at the Server Edge.....	7
Access Gateway.....	7
Host Bus Adapter.....	7
Converged Network Adapter.....	8
5. Brocade DCX Backbone Family.....	8
6. Brocade Adaptive Networking Services.....	8
7. Virtual Machine SAN Boot.....	8
8. Storage Virtualization and Tiered Storage.....	8
9. Synchronous Storage Replication over xWDM Optical Networks.....	8
10. Asynchronous Storage Replication over TCP/IP.....	9
11. VMware Cluster Replication over TCP/IP.....	9
12. Integrated Management.....	9
Using the Reference Architecture in the Virtual Data Center	9
Collapsing the Blade Server Edge with Brocade Access Gateway.....	10
Adding Adaptive Networking with the Brocade HBA.....	11
VMware Cluster Storage Architecture with FC QoS.....	13
Scaling the Virtual Infrastructure.....	15
Evolution at Server Edge.....	16
Brocade Unified Fabric Management.....	17
Components of the Reference Architecture	17
Brocade DCX Backbone.....	17
Brocade Access Gateway.....	17
Brocade HBA.....	17
Brocade CNA.....	18
Blade Servers.....	18
VMware vSphere 4.....	18

EXECUTIVE SUMMARY

There is a fundamental transformation going on in the data center Infrastructure. Companies are intently focused on data as a strategic asset—accelerating the amount of digital data they manage. Innovations in application architecture have transformed discrete applications into software services logically integrated using an Enterprise Service Bus (ESB) to improve business agility. The guiding principles driving a new architecture for the data center include: consolidation, virtualization, service-level resource allocation, and policy-based data management. These principles simplify how business exploits its data for competitive advantage in a cost-effective manner. Networks in the data center will be transformed from a number of network "silos" into a single, logical, high-performance, low-latency data center fabric.

NOTE: An Enterprise Service Bus (ESB) refers to a software architecture construct typically implemented by technologies found in middleware infrastructure products. It is usually based on recognized standards, which provide foundational services for more complex architectures via an event-driven and standards-based messaging engine (the bus).

This new fabric will accelerate deployment of virtualization, simplifying management and ensuring that infrastructure performance/cost ratios keep up with Moore's Law (twice the performance at the same cost, every 18 months). And all infrastructure transformations have to integrate existing assets with new technology. New virtualization-aware products, which push policy-based management into the infrastructure and extend application service levels across the infrastructure, must non-disruptively co-exist with existing products.

The strategy for the Brocade® data center fabric was created with these real-world requirements in mind as a roadmap for the evolving data center. Using this strategy, we have developed a reference architecture that adds significant value to VMware vSphere 4 data center technology. Customers who are planning to deploy VMware or who already use VMware can use this reference architecture to transform today's Storage Area Network (SAN) fabrics into tomorrow's data center fabrics.

The reference architecture detailed in this paper describes Brocade's strategy for using the power of the Brocade DCX® Backbone (at the fabric core) connected to Brocade Access Gateway in a blade server chassis and Brocade Host Bus Adapters (HBAs) and Fibre Channel over Ethernet (FCoE) connectivity via Converged Network Adapters (CNAs) in the blade server—all elements of the Brocade architecture—to simplify configuration and management of VMware vSphere and to extend fabric services into the Virtual Machine (VM). These products easily integrate with a variety of blade server and storage products supplied by Brocade OEM partners. The Brocade Access Gateway feature of Fabric OS® (FOS) and ISA technology non-disruptively connect to existing SAN fabrics. This provides an "on ramp" to a high-performance, scalable, continuously available data center fabric that converges server-to-server clusters, SANs, and storage replication networks.

The Brocade DCX Backbone family represents a new class of fabric infrastructure platform. Essential requirements for the core of the Brocade data center fabric include extraordinary performance (bandwidth, port count, and power efficiency), non-disruptive scalability, continuous availability, and end-to-end security—all achieved by the Brocade DCX architecture. The Brocade DCX design supports advanced technology and protocols such as Fiber Channel over Ethernet, Converged Enhanced Ethernet (CEE), encryption services, and policy-controlled data mobility. Rather than compete with applications for control of their data, the Brocade DCX extends fabric intelligence, logically connecting it to application service levels and data management policies—moving them into the fabric. This improves infrastructure utilization; simplifies provisioning, capacity planning, and management; and reduces infrastructure costs. All of these efficiencies make more dollars available for investment in business growth.

The Brocade DCX Backbone performance—up to 384 x 8 Gigabits per second (Gbps) ports—4x or greater switching performance compared to today's SAN directors and innovative 512 Gbps Inter-Chassis Links (ICLs) are required to deploy VMware on blade servers at the fabric edge. Growing demand for server virtualization and new blade servers designed for 8-core Central Processing Units (CPUs) and higher mean that the fabric core must scale to keep up with much higher ratios of VMs per server. The performance of the HBA and Access Gateway platforms at the fabric edge will scale to 8 Gbps performance levels as well.

OVERVIEW OF THE BROCADE DATA CENTER FABRIC STRATEGY

The Brocade strategy extends the recognized success of SAN technology (fabrics) to the entire data center. Five strategic directions guided the design:

- Consolidation
- Virtualization
- Management partitioning
- Application-aware service levels
- Policy-based data management

As the amount of digital data continues to grow at a rate of 50 to 70 percent per year, consolidation has become an essential strategy to scale the server and storage infrastructure, boosting resource utilization and keeping costs in check. The growth of separate networks in the data center (many providing "best effort" rather than "guaranteed" delivery) constrains the rate of growth, add complexity, and increase risk. A new network architecture is needed to achieve much higher levels of performance, availability, and scalability. Since these qualities are the hallmark of SAN fabric technology, it is the logical foundation for building a data center fabric.

An effective consolidation tool is virtualization of servers and storage. Server virtualization of commodity servers using VMs creates an agile computing infrastructure that scales cost effectively. Server virtualization can triple resource utilization with a corresponding reduction in Total Cost of Ownership (TCO). Storage virtualization, combined with transparent data mobility in the fabric, enables Information Lifecycle Management (ILM). Transparent data mobility allows data management policies to be extended into the fabric, dramatically reducing data storage TCO.

Recognizing the impact of server virtualization on the fabric, Brocade developed its virtual channel technology to provide fine-grained resource allocation managed by service-level policies. Logically connecting virtual servers to virtual storage using virtual channels means that application service levels are extended into the fabric, which improves the agility of fabric provisioning and configuration. Combining Virtual Channels with HBA support for N_Port ID Virtualization (NPIV) means that VM provisioning connects directly with fabric services simplifying management as workloads migrate across virtual servers.

As multiple applications share physical server, network, and storage infrastructure, logical partitioning becomes an essential tool to align application priorities and service levels with organizational structures and budgets. Virtual fabrics are a tool to partition the management of the physical resources used by virtual servers and virtual storage.

Infrastructure management can no longer assume that an application runs on a particular server or that its data is stored in a single physical storage location. In today's virtualized environments, VMs are unaware of physical ports, since all underlying hardware has been abstracted. Binding Virtual Machines to virtual ports enabled by NPIV and classifying application workloads via Quality of Service (QoS) groups ensure that fabric service levels move with the VM as workload managers migrate it across the physical server resource pool. Brocade's Adaptive Networking services provide this logical link between application service levels and fabric services.

Data management has relied on the assumption that an application's data was captive to a physical storage location over its lifecycle. With storage virtualization and ILM policies, data management, security, and protection are no longer tied to a single physical storage location. Instead, data management policies become properties of the data itself. As data migrates between storage tiers, the fabric needs to access the data management policies to consistently protect, secure, and retrieve data regardless of the physical storage tier in which it currently resides.

The five strategic principles—consolidate, virtualize, partition, allocate resources at the required service level, and configure policy-based data management—are well suited to work with the VMware vSphere 4 strategy. The remainder of this paper shows how current and planned products designed to implement the Brocade data center fabric could be deployed in VMware environments to transform today's physical data center into tomorrow's virtual data center.

THE BROCADE REFERENCE ARCHITECTURE

(Key elements of the architecture are shown by the numbered circles in the diagram and are described in the numbered sections that follow.)

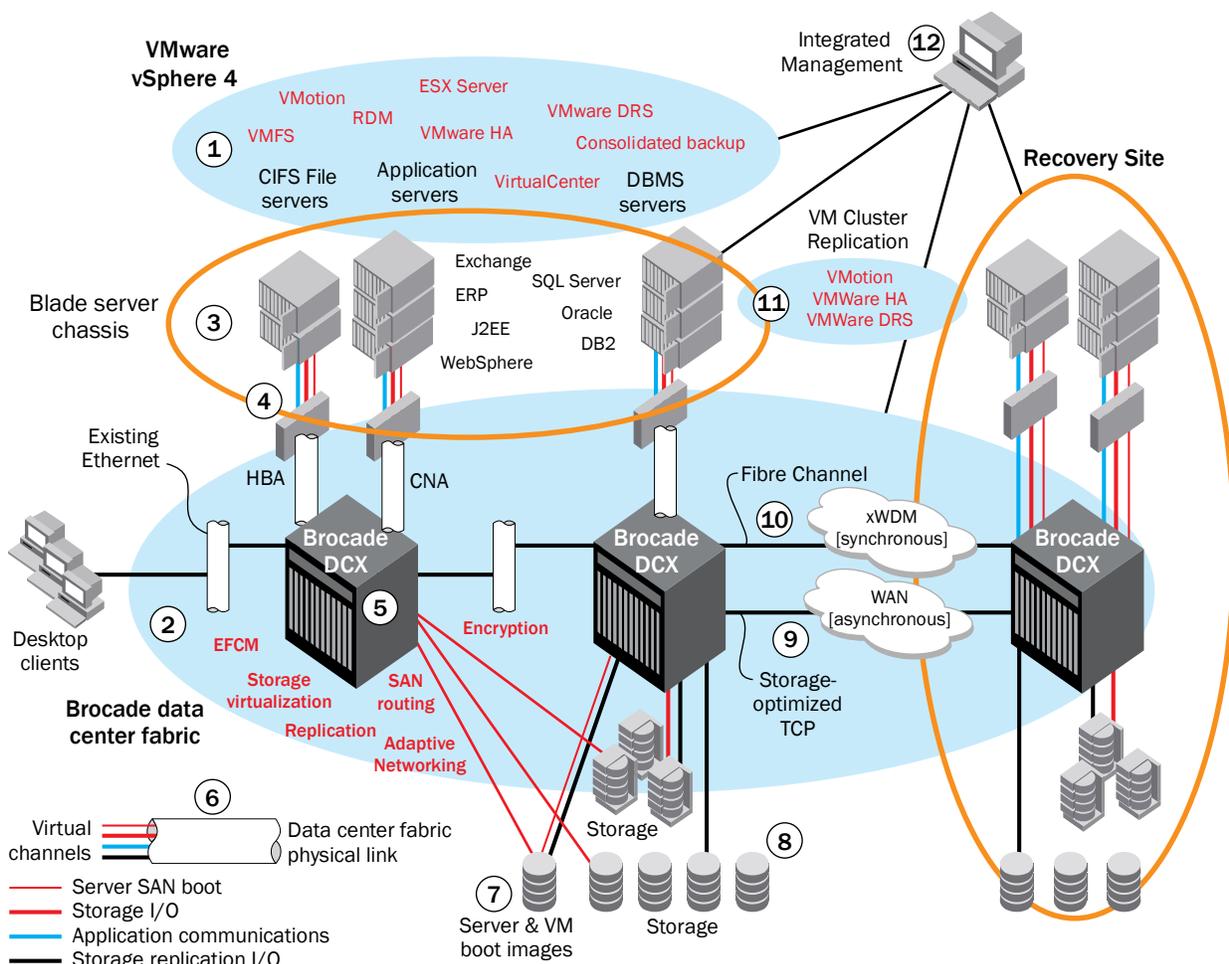


Figure 1. VMware virtual data center with Brocade data center fabric

1. VMware vSphere 4

VMware vSphere 4 is software designed to consolidate commodity server technology into a resource pool, which can efficiently and cost effectively host mission-critical business applications. Servers and storage can be added to the resource pool to support multiple applications. This dramatically boosts server and storage utilization, improves application availability (since application workloads can be rapidly moved from one physical server to another), and eliminates application outages as servers are replaced at the end of their useful economic life.

With vSphere 4, VMware introduces a new architecture to enable third party vendors to integrate their unique capabilities with the VMware platform. vSphere 4 defines abstraction layers for storage, network connectivity, and computational services—called vCompute, vStorage and vNetwork respectively. These open interfaces enable third party hardware and software vendors to expose unique differentiated services in each of these areas.

Software components available in vSphere 4 are described in the following sections.

ESX Server 4.0

ESX Server is software that emulates the physical resources of a server. It creates a virtual machine (memory, CPU, peripheral devices, network and storage connections) containing an operating system and business applications. Components of the ESX Server are deployed on the bare metal commodity servers, creating a virtualization layer that separates the server hardware from the operating system and applications. ESX Servers can be centrally provisioned, configured and managed using the VMware VirtualCenter management server and the VI client.

VMware VirtualCenter and VI Client

VMware VirtualCenter server installs as a service on a Windows server allowing management and control of virtual machines and ESX Server software. The VI client installed on a Windows machine provides the user interface for managing the VI3 environment.

VMFS

Virtual Machine File System (VMFS) is a clustered file system with integrated volume management, which stores VM configurations and application data. All physical servers in a VMware cluster have access to the data stored in VMFS. SANs are commonly used to provide many-to-many connectivity between the physical servers and the storage LUNs holding the VMFS data. When VMs and their applications are moved between physical servers in the resource pool, the clustered file system ensures that the application maintains non-disruptive access to its data.

RDM

Raw Device Mapping (RDM) provides direct block access for applications that access their data at the block level rather than file level. RDM can be configured so that VMFS can control access to block storage. This allows virtual machines to maintain block access to storage as they migrate between physical servers in the resource pool.

VMotion

VMotion is the tool used to migrate a VM with its operating system and applications between physical servers. This migration is transparent to the application, maintaining connectivity non-disruptively between the application and its data.

VMware HA

VMware HA (High Availability) is a management application that protects a VM and its applications from server hardware failures. It is used with a server cluster and monitors the status of the server hardware. Should the server go down, VMware HA instructs VMotion to migrate all the VMs running on the downed server to different physical servers in the cluster. VMware HA can be used with application clustering, such as Microsoft Cluster Server software.

VMware DRS

VMware DRS (Distributed Resource Scheduler) is a management application used with a VMware cluster to monitor the resource use of a VM. If the application workload requires more resources than currently allocated to the VM, VMware DRS can move the VM to another server with more resources. VMware DRS can also be used with VMware HA. In this configuration, if a physical machine in a cluster goes down, VMware DRS can automatically migrate the VMs on the failed server to other servers, so applications continue to be processed while continuing to meet service levels.

2. Brocade Data Center Fabric

The Brocade data center fabric provides integrated connectivity for SAN, server cluster, storage-to-storage, and Ethernet traffic. At the storage core of the data center fabric is the Brocade DCX Backbone; at the edge is the Brocade HBA for Fibre Channel and Brocade Access Gateway inside the blade server chassis or CNA for FCoE and Converged Enhanced Ethernet (CEE) at the blade server.. The Brocade DCX is designed to non-disruptively connect to existing SAN fabrics to enable scalability when business requirements mandate data center growth. Further, the Brocade DCX seamlessly supports technologies, such as FCoE and Converged Enhanced Ethernet (CEE) for those customers who want I/O convergence. Consequently, integration of the Brocade DCX into existing fabrics creates a stepping stone for a multiprotocol data center fabric.

3. Blade Server Chassis

Blade server chassis are at the fabric edge. They contain blade servers and slots for FC and Ethernet switches in a dense, energy-efficient package. Blade server technology is available from all the primary suppliers of server technology and blade servers are offered in a variety of configurations. They typically include integrated chassis management, can be stacked up to four high in standard 19" racks, and are easy to attach to both Ethernet networks and FC fabrics. With new technology using 10 Gbps CEE with FCoE, a single physical server connection can be used to attach blade servers to existing IP networks and storage SANs..

4. Brocade Solutions at the Server Edge

Brocade innovation simplified management of the edge of the fabric with the Access Gateway feature in FOS and the release of Brocade Fibre Channel HBAs and Brocade CNAs,

Access Gateway

The Brocade Access Gateway aggregates Fibre Channel traffic from all the blade servers in a chassis and provides non-blocking frame trunking with transparent failover. It greatly improves scalability of blade server environments by eliminating low-port-count-switches. It improves manageability by cleanly separating server administration from SAN administration, allowing server administrators to configure the Access Gateway with no adverse affect on the fabric. With F_Port Trunking, high utilization and automatic link failover simplify management, increase availability, and reduce costs.

Host Bus Adapter

Brocade offers 4 and 8 Gbps Fibre Channel HBAs, which can deliver the performance required by 4-core CPU servers with their increased VM workloads. The HBA provides SAN boot functionality and is designed to extend Adaptive Networking services into the VM. Application service levels are linked to QoS priorities via the Brocade virtual channel technology, which extends application service levels into the fabric and enables QoS and traffic management services to detect and avoid congestion. This simplifies provisioning, as QoS priorities are defined by the VM and move with the VM when it migrates to another physical server. Application priority is not the same as application workload. The HBA extends VM configuration to include application priority, assuring fabric resources are allocated accordingly and reducing the need to migrate VMs.

Converged Network Adapter

The Brocade 1010 (single port) and Brocade 1020 (dual ports) 10 Gbps FCoE/CEE CNAs provide a new level of server I/O consolidation through hardware capabilities that integrate seamlessly with existing Fibre Channel and Ethernet environments. This new class of CNAs is designed to help organizations reduce the number of adapters used in servers, which in turn reduces the number of cables and corresponding switch ports to lower overall costs. Leveraging IEEE standards for Data Center Bridging (DCB), the Brocade 1000 Series CNAs provide a highly efficient way to transport Fibre Channel storage and IP traffic over Ethernet links and meet the demands of increased traffic created by server virtualization.

5. Brocade DCX Backbone Family

At the storage core of the Brocade data center fabric is the Brocade DCX Backbone, comprising the Brocade DCX and the Brocade DCX-4S. The Brocade DCX delivers the demanding performance required by server consolidation and provides "no-compromise" virtual server deployment. With 4x or greater the switching performance of today's SAN directors, 384 ports of 8 Gbps Fibre Channel, the Brocade DCX can scale to meet large increases in bandwidth created by the ever-growing number of VMs per server. The Brocade DCX architecture is designed to non-disruptively integrate FCoE and CEE for VMware cluster connectivity with 24-port 10 Gbps FCoE/CEE blades (Brocade FCOE10-24 Blade).

6. Brocade Adaptive Networking Services

Virtualization of servers and storage place new demands on the fabric to keep up with dynamic workloads and avoid congestion caused by unexpected resource bottlenecks. Adaptive Networking services include QoS, Traffic Management, Fabric Dynamic Profiling, and Buffer Credit Recovery. The Brocade HBA extends these services to the application layer. Brocade Adaptive Networking services create multiple logical connections over a single physical link, which efficiently isolates traffic by QoS priority (high, medium, low) and efficiently allocates fabric resources when connecting virtual servers to virtual storage. Adaptive Networking features include Ingress Rate Limiting, Congestion Detection, Traffic Isolation, and fabric dynamic profiles such as "Top Talkers." These services simplify management and ensure high utilization of fabric resources. Multiple virtual channels per physical link can be used to isolate the dynamic behavior of different application service classes when they share the same physical connection so that they do not disrupt each other. Adaptive Networking services make the fabric more agile and ensure application data continues to flow according to service levels as application workloads are moved between physical servers.

7. Virtual Machine SAN Boot

VMware supports the ability to store VM configurations as files and supports HBA SAN boot configurations. This can dramatically reduce costs as blade servers do not require disk drives when configured with SAN boot HBAs. Further, change control and OS upgrades are easier to manage with a centralized repository of master boot images, which can be cloned and used by multiple VMs.

8. Storage Virtualization and Tiered Storage

The performance of the Brocade DCX increases the scalability of network-hosted storage virtualization. With a data mobility engine controlled by storage virtualization applications, the Brocade DCX simplifies adoption of tiered storage architectures, a principle benefit of ILM strategies. ILM migrates data between storage tiers over the lifecycle of the data based on data management policies, which reduces storage and management costs and increases storage scalability.

9. Synchronous Storage Replication over xWDM Optical Networks

Data protection is a fundamental requirement of every business. With the advent of regulatory requirements mandating data protection and timely data recovery, the use of high-performance, array-based synchronous data replication has grown. Typically this is deployed over Layer 1 optical networks using either Dense or Coarse Wave Division Multiplexing (xWDM). These optical networks are now capable of delivering 10 Gbps performance in a cost-effective manner.

The Brocade DCX performance supports 10 Gbps xWDM links and integrates Fibre Channel over IP (FCIP) with per-port SAN Routing to efficiently move data between the principal data center and a recovery site. It also supports an application services platform integrated with *synchronous data replication applications* to speed synchronous data replication traffic for heterogeneous array replication.

10. Asynchronous Storage Replication over TCP/IP

When the distance between sites exceeds 60 to 80 mi (100 to 120 km), asynchronous storage replication over TCP/IP is commonly used. To reduce bandwidth costs and latency, the Brocade DCX integrates Storage Optimized TCP over FCIP and supports both 1 GbE and 10 GbE links. These enhancements to TCP improve efficiency reducing WAN bandwidth (and cost) while reducing the asynchronous replication outage window. Additional optimization features include FastWrite and Tape Pipelining for both FCIP (SCSI) and FICON traffic. The Brocade DCX supports an application services platform integrated with *asynchronous data replication applications* to speed asynchronous data replication traffic for heterogeneous array replication.

11. VMware Cluster Replication over TCP/IP

VMware HA provides cluster failover when a physical server goes down or is taken offline. This can be used with VMware DRS to load balance application workloads from failed servers to the remaining servers in the cluster. VM configurations are stored as "images" in files, which are replicated to the recovery site using synchronous or asynchronous storage replication. The Brocade DCX application services platform support for storage replication means that VMware cluster replication can take advantage of the ability of the Brocade DCX to speed replication traffic for both the application image and its storage—simplifying disaster recovery designs.

12. Integrated Management

VMware environments dramatically increase the number of data flows and bandwidth per physical server. Combined with the dynamic movement of workloads between servers, complex traffic patterns are created and unexpected congestion can occur. This complicates management, can cause premature VM migration, and can lower infrastructure availability and performance. The solution lies in connecting fabric services with virtualization management tools, which improves application control of the data rather than letting the fabric compete for control. Fabric services need to be fine-grained to monitor resource use at the individual application data flow level, not just the fabric port. Open management standards, such as SMI-S, are important for integrating virtualization management tools with the fabric's management services. Configuration management, capacity planning, service level policies, and VM provisioning tools can be integrated with fabric services such as Adaptive Networking, encryption, and security policies.

USING THE REFERENCE ARCHITECTURE IN THE VIRTUAL DATA CENTER

Data center infrastructure, by the nature of its construction, never allows "big bang" transitions. Data centers host business applications and corporate data, which increasingly must be continuously available. Individual IT projects are often the way new technology is adopted and deployed. Consequently, technology transformations, such as virtualization, must co-exist with existing products and technology.

The Brocade data center fabric requires products such as the Brocade DCX Backbone to interoperate with existing SAN technology. When business determines that you need to make the transition from SAN to data center fabric, you can do so non-disruptively by upgrading SANs with future products designed to enable the data center fabric. The following example illustrates how to apply this reference architecture to integrate the Brocade DCX Backbone and Access Gateway, with server adapters in data centers using blade servers, storage arrays, and VMware vSphere 4—to achieve high-performance, scalable, cost-efficient server virtualization.

Collapsing the Blade Server Edge with Brocade Access Gateway

The blade server chassis includes space for Ethernet and FC switches, which provide cluster communications and SAN storage connections to 8 to 16 blade servers per chassis. Early in the adoption of blade servers, Brocade was among the first to understand and appreciate their impact on the fabric. The issue was this: many low-port-count switches quickly exhaust the FC switch, or domain ID, addresses, which are limited to 255. Brocade introduced the Access Gateway feature in FOS to improve performance, reduce management complexity, and boost fabric scalability at the blade server edge.

NOTE: The term “Access Gateway” is used in the following sections to refer to a Brocade embedded switch in Access Gateway mode.

As shown in Figure 2, this software feature can be used with existing 4 Gbps embedded blade server FC switches, converting them from a switch to an Access Gateway or N_Port cluster. On the back-end of the Access Gateway, an F_Port connects each blade server to the Access Gateway. A second Access Gateway is connected to each blade server providing high availability. On the front-end, the Access Gateway presents multiple N_Ports, which are aggregated into a load-balanced F_Port trunk connected to the core backbone or optionally an edge switch.

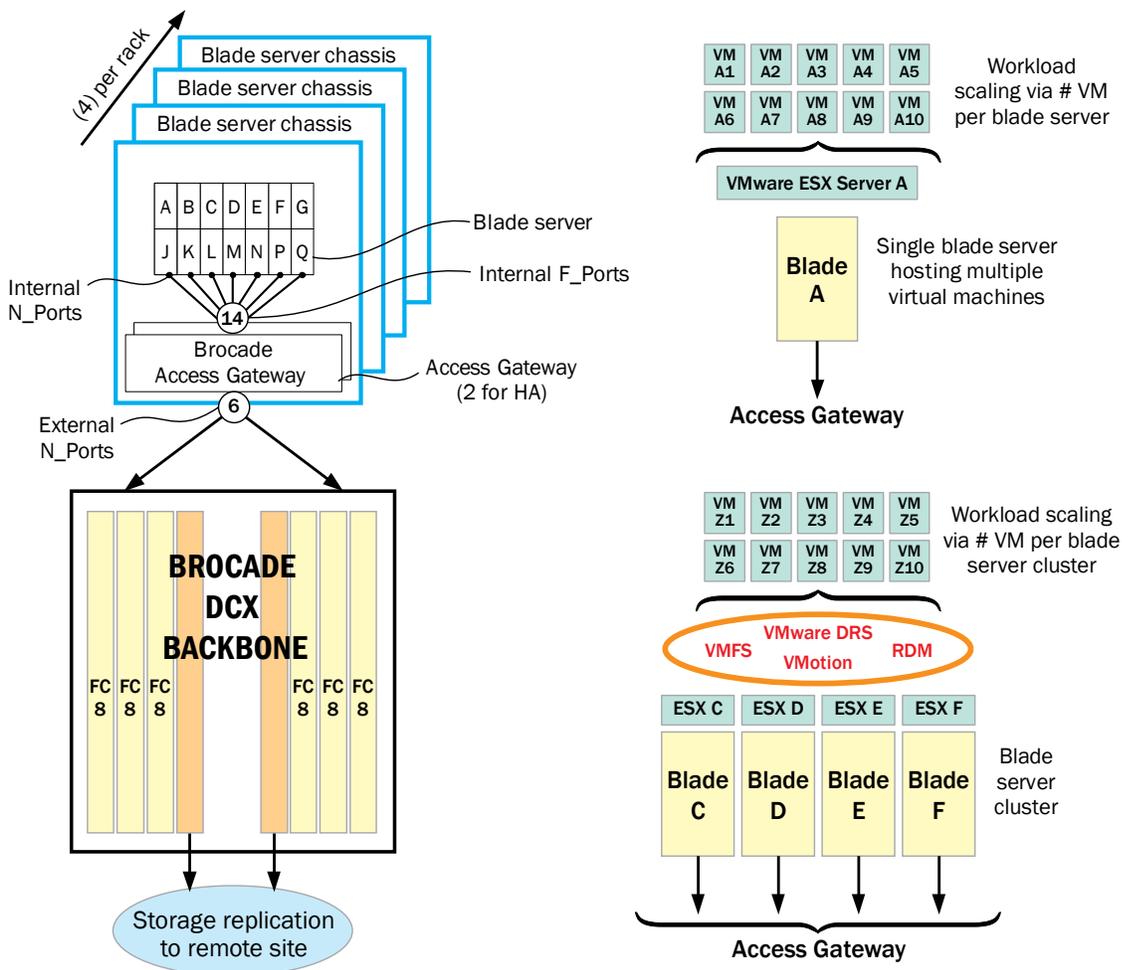


Figure 2. VMware virtualization with the Brocade DCX Backbone and Access Gateway

There are several advantages to using an Access Gateway to collapse the fabric edge. First, the FC address space is more efficiently used greatly improving scalability. Rather than consume an address from the smaller domain address space (limited to 255 addresses) for every 8 to 16 servers, the much larger device address space (approximately 65,536 addresses) is used, greatly increasing scalability. Second, server administrators can configure the Access Gateway inside the blade server chassis independent of the storage administrator's configuration of the fabric. And third, the Access Gateway provides load balancing between backend and front-end ports using F_Port trunking between the Access Gateway and the Brocade DCX. This delivers high bandwidth utilization without over-provisioning the number of physical links.

VMware can be deployed on the blade servers, as shown at the top right of Figure 2, to subdivide a single blade server into multiple VMs. Each VM runs an instance of an operating system and appropriate application(s). Or, as shown at the bottom right, multiple blade servers can be combined into a cluster. The VMware cluster file system, VMFS, and/or the cluster "raw" block interface, RDM, can be used to connect VMs to their storage in the SAN fabric.

Adding Adaptive Networking with the Brocade HBA

With the VMware ESX Server software loaded on the blade servers, a cluster can be created to pool the computing resources (CPU, memory, network, and I/O) of the individual blade servers. VMs can be defined and resources allocated from the pool based on the expected application workload. This means a single blade server has multiple logical fabric connections, one for each VM. Provided the fabric HBA within the blade server supports N_Port ID Virtualization (NPIV), multiple virtual N_Port connections on the backend are created, one (or more) for each VM.

Brocade Adaptive Networking technology includes QoS, Traffic Management, Fabric Dynamic Profiling, and Buffer Credit Recovery services. Application-aware service levels can be connected to Adaptive Networking services in the fabric via HBAs on the server edge. This is an important feature designed to integrate the fabric's intelligence with the application, rather than use it to compete with the application for control of its data. Logically connecting fabric services to the VM simplifies fabric provisioning and improves utilization of fabric resources.

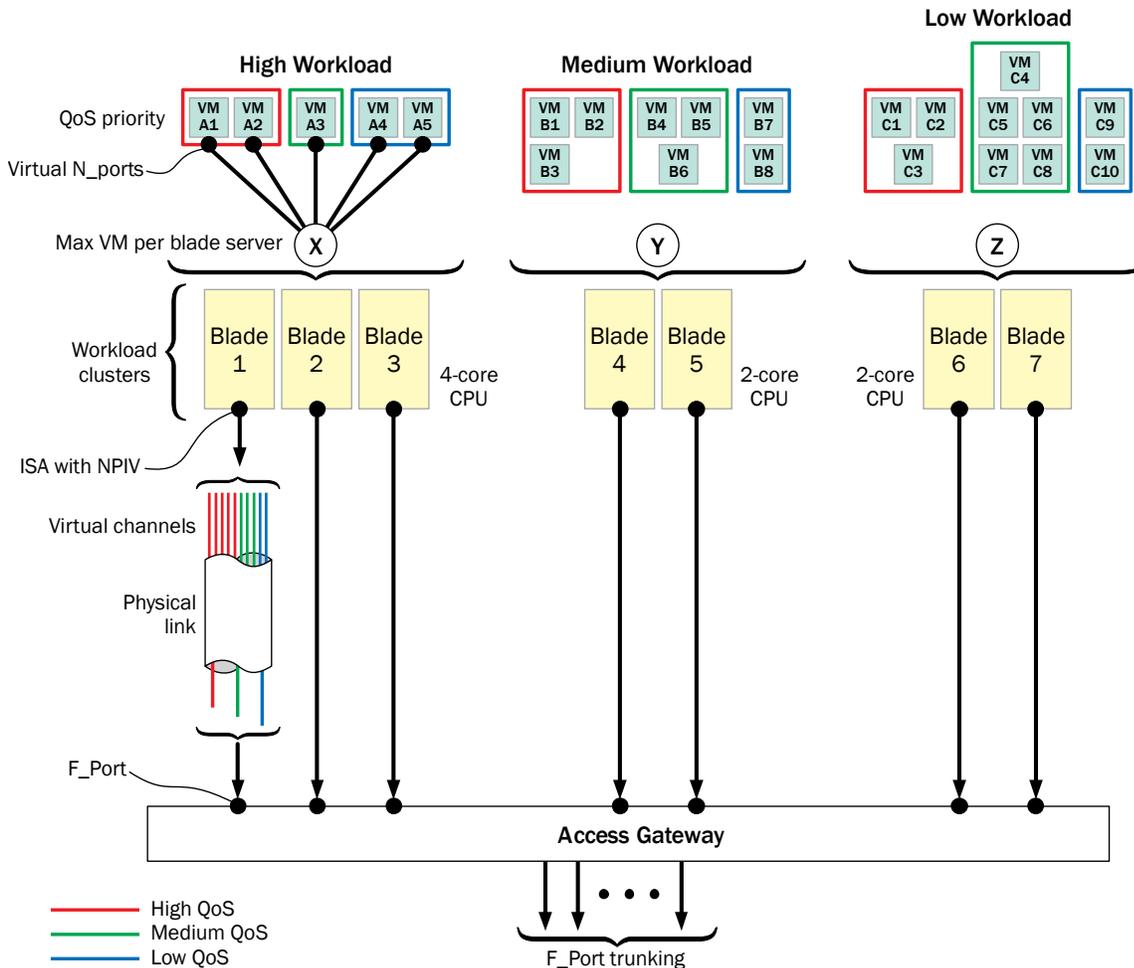


Figure 3. VMware Infrastructure resource allocation with Brocade application-aware QoS

Figure 3 shows a blade server edge with VMware deployed. Three different VMware clusters are configured, each for a different class of application workload (high, medium, low). Applications that have high workloads (and correspondingly high I/O requirements) are deployed on the High Workload cluster. Today, this cluster might use blades with four core CPUs. Similarly, the other two clusters are configured for medium and low application workloads, perhaps using dual-core CPUs due to the lower workloads and I/O requirements.

Notice that VMs in each cluster have been assigned different QoS priorities. This is an important capability of the Brocade reference architecture. Since the priority of an application is not the same as its workload, the Brocade data center fabric allows independent management of the QoS priority from the cluster resource needed to support an application's workload. For example, backup application workloads are often very high, but since they are usually less critical than other business applications, they can be assigned a lower QoS priority and higher cluster resources.

VMware Cluster Storage Architecture with FC QoS

VMware clusters combine multiple servers and their resources into a resource pool and allocate resources from the pool to a VM. The cluster can use a resource manager, VMware DRS, to automatically monitor resource use and move VMs between physical servers to accommodate variations in application workloads.

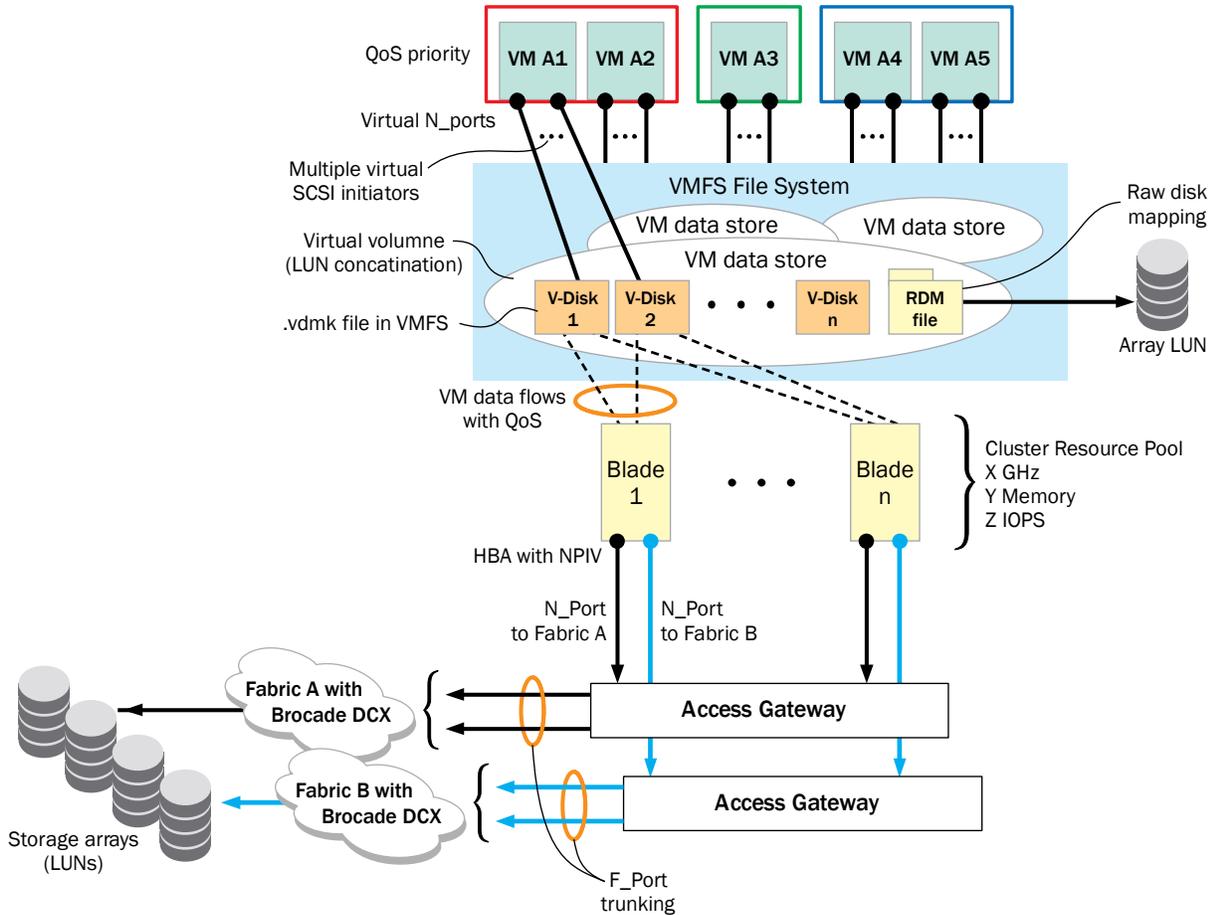


Figure 4. Brocade Adaptive Networking QoS applied to a VMware cluster with VMFS file system

As shown in Figure 4, VMware has a clustered, journaling file system, VMFS, with file locking and volume management. One-way storage can be presented to VMs via virtual disks (v-disks). V-disks are special files stored in the VMFS file system. Storage capacity (LUNs) is attached to a VM data store. A data store can span storage types (SCSI, FCP, iSCSI, Network-Attached Storage (NAS)) and presents a single interface to the virtual SCSI initiators used by a VM for storage access. In addition to the application's data files, VMFS also contains the VM resource definitions, running state and boot files.

To a virtual machine, the VMFS disk file appears to be a physical disk. Via VMFS, a VM continues to access its storage as it moves from one blade server to another. VMware DRS monitors VM resource use and can use VMotion to transparently move a VM and its application workload to a different physical server. For example, in Figure 4, if virtual machine VM-A1 moved from blade server 1 to blade server *n*, VMware DRS ensures that the move is transparent to the application running on it. This requires that all VMs in the cluster have access to the same disk files (v-disk). When VMotion moves the VM to blade server *n*, it migrates the virtual SCSI initiators allocated to VM-A1 from blade server 1 to blade server *n*. Thus, the data that flowed through the physical HBA attached to blade server 1, now flows through the physical HBA attached to blade server *n*. Since Brocade Adaptive Networking assigned a QoS priority of high to the virtual N_Port attached to VM-A1, and the virtual N_Port is part of the resource definition file for VM-A1, it moves when the VM moves. This ensures the application's QoS priority stays with the VM as it moves.

The VMFS file system also supports "raw" I/O to a storage LUN and migrates that I/O with the VM. A special RDM file is used by the VMFS. The RDM file points to the storage LUN. This allows VMotion to migrate the VM without disrupting the raw I/O to the storage LUN.

Finally, the virtual volume function of the VM data store allows storage LUN concatenation to expand the storage a v-disk can access without application disruption. And, the virtual volume manager supports array-based snapshots for creating test data, Point in Time (PIT) copies, and remote replication applications.

Scaling the Virtual Infrastructure

Blade servers today are using 4-core CPU technology and 4 Gbps Fiber Channel connections. Their performance will increase as 8-core CPU and 8 Gbps Fibre Channel technologies are integrated. The boost in server performance will fuel an increase in the number of VMs running on a physical server generating larger demands on fabric performance.

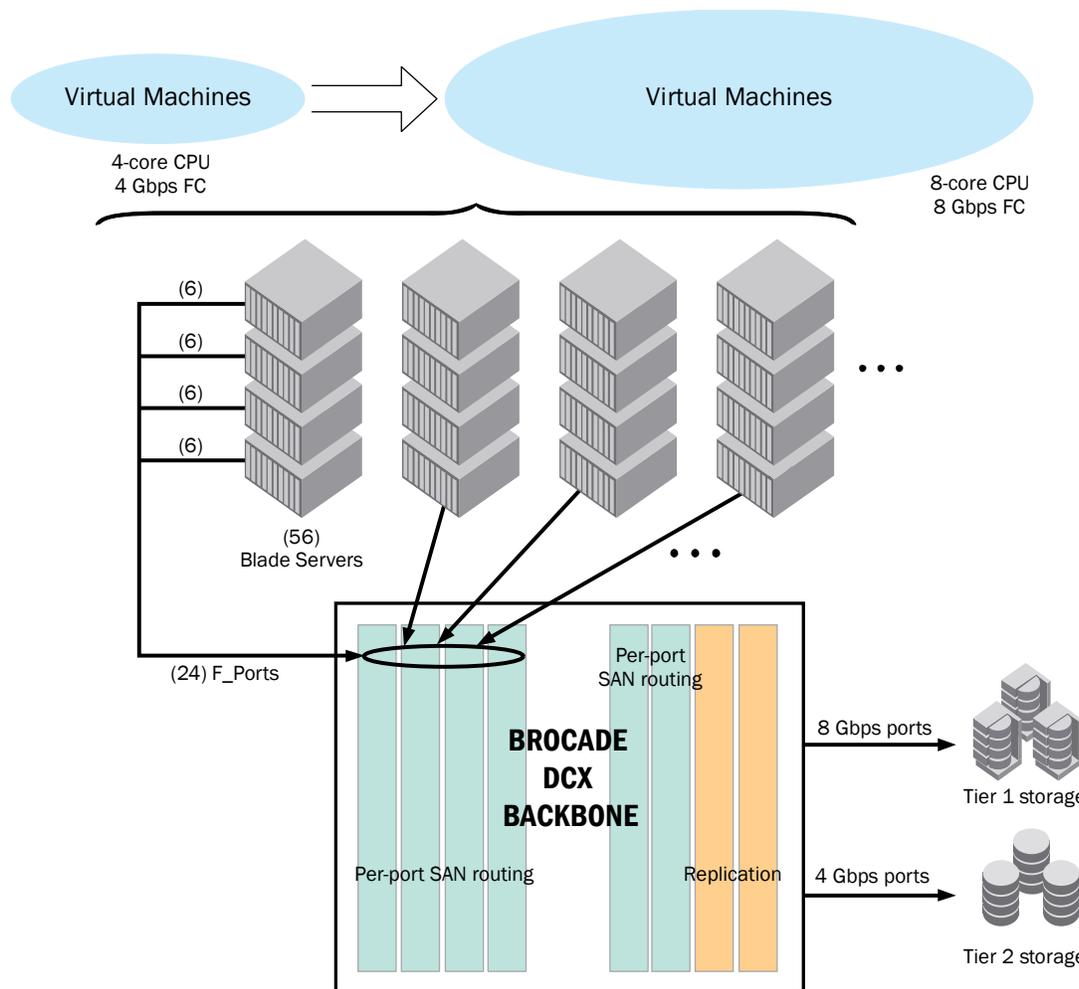


Figure 5. Scaling VMware with the Brocade DCX Backbone 8 Gbps technology

The next-generation blade server chassis will increase VM density by a factor of two, three, or more beyond the current levels. With its 8 Gbps technology and Adaptive Networking services, the existing Brocade DCX can easily scale its performance to meet this demand. Application-aware congestion management rather than port-based approaches will become important as more VM workloads dynamically move across blade servers. Trunking between the blade server edge and the core will ensure high-bandwidth utilization without costly over provisioning. These features are integrated into the Brocade DCX Backbone, Access Gateway, and ISA.

Figure 5 shows a Brocade DCX Backbone initially connecting to today's 4 Gbps, 4-core CPU blade servers. Since the Brocade DCX supports up to 256 virtual N_Ports per physical switch port, and the Access Gateway provides F_Port Trunking with load balancing, the original design can scale from the current ratio of 20 to 30 virtual machines per blade server to the much larger ratios possible with 8-core CPUs and 8 Gbps connectivity.

The Brocade DCX Backbone integrates data mobility services to accelerate storage virtualization and replication via its application services platform. This simplifies use of tiered storage architectures commonly deployed for ILM strategies. Combined with per-port SAN Routing, the Brocade DCX connects storage replication applications (synchronous or asynchronous) over xWDM or TCP/IP WAN connections.

Evolution at Server Edge

Transforming a SAN into a data center fabric begins by adding a Brocade DCX Backbone to the core. (Note that to support the evolving data center fabric, Brocade is investing and actively contributing to the development of the Enhanced Ethernet standard required to enable FCoE.) The Brocade DCX was designed to non-disruptively add FCoE technology using the underlying 10 Gbps CEE to transport Fibre Channel.

With the introduction of FCoE, the core of the fabric needs a high-performance, multi-protocol switch engine able to handle FCoE and CEE connectivity. As shown in Figure 6, either the Brocade 8000 Switch (top of rack) or the FCOE10-24 Blade in the Brocade DCX chassis (end of row) can be used to integrate FCoE solutions in the server with the existing SAN. With the Brocade 8000, FCoE traffic from servers is separated into IP and SAN storage traffic. With the FCoE blade in the Brocade DCX, FCoE traffic from servers enters via CEE ports in the blade, SAN traffic is routed over the backplane to Fiber Channel SAN storage ports, while IP traffic is sent over 10 GbE ports in the blade connected to routers in the IP network. The Brocade DCX Backbone architecture supports consolidation of multiple traffic types. It switches multiple protocols using a switch engine architecture that is agnostic to the type of frame (FC, CEE, and so on). Using the Brocade DCX at the core delivers cost-effective convergence of FC, FCoE, and CEE and traffic to leverage its backbone level of performance.

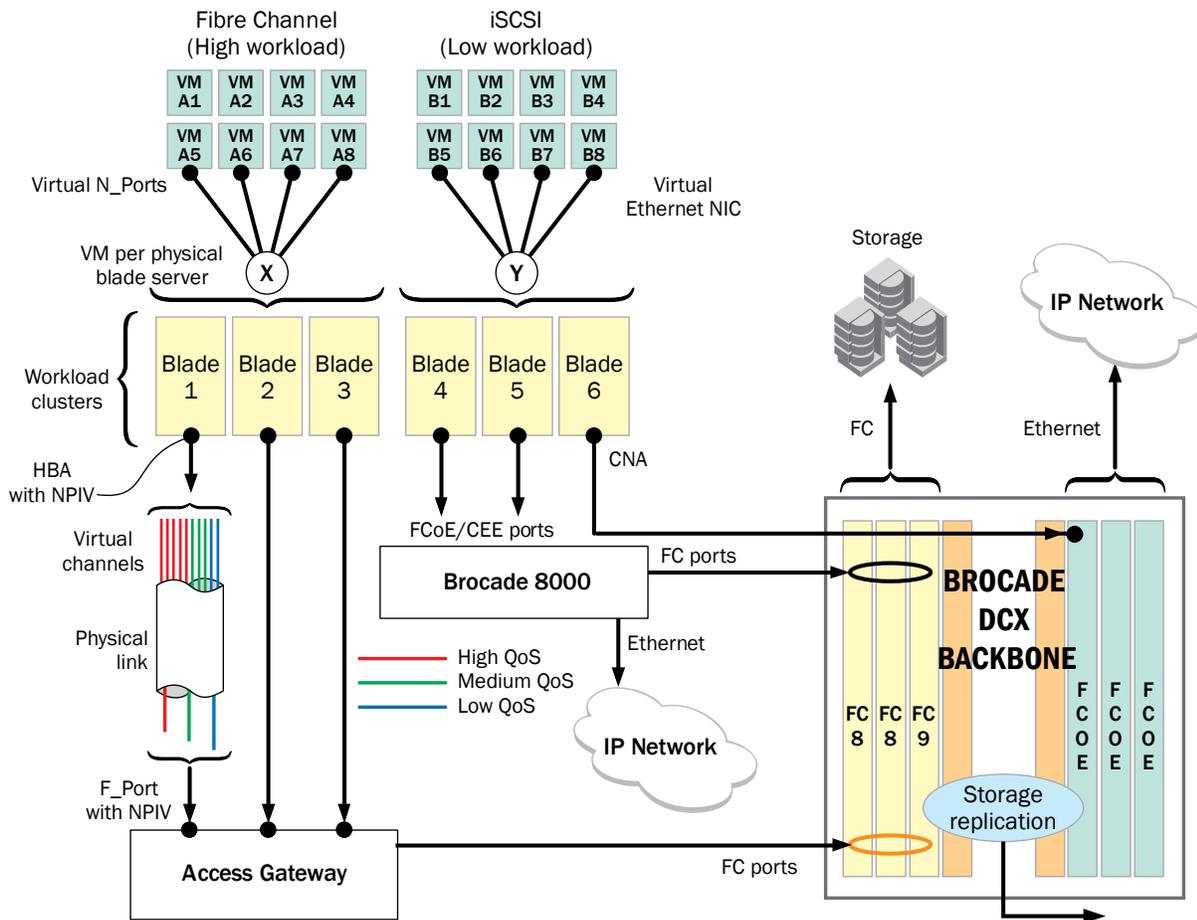


Figure 6. Transitioning from a SAN to a Brocade data center fabric

Blade servers are on track to integrate 10 GbE technology with FCoE support, enabling storage and application traffic consolidation onto a single physical interface. As customers refresh their blade servers, the SAN will seamlessly evolve into a Brocade data center fabric.

Brocade Unified Fabric Management

Brocade® Data Center Fabric Manager (DCFM™) Enterprise is a comprehensive network management application that enables end-to-end management of data center fabrics. Brocade DCFM Enterprise manages and secures the flow of data across multiple fabrics. As a key component of the architecture, Brocade DCFM Enterprise enables unified management of data center fabrics—from storage ports all the way to the HBAs, both physical and virtual. It configures and manages the Brocade DCX family along with Brocade directors, routers, and switches. Moreover, it supports Brocade encryption capabilities for data-at-rest and HBA products, as well as storage networking technologies such as Virtual Fabrics and protocols such as FCoE and CEE.

COMPONENTS OF THE REFERENCE ARCHITECTURE

The following summarizes the components supported for this reference architecture. Brocade works with our partners to certify these components for use in the data center fabric architecture and to support new releases and upgrades as they become available. To find out more about the Brocade components, visit www.brocade.com/products-solutions.

Brocade DCX Backbone

Key features of the Brocade DCX Backbone used in the reference architecture include:

- Trunking (E_Port and F_Port) with automatic link fail-over
- Virtual Fabrics
- Adaptive Networking services—Fabric and application QoS, Traffic Management, Fabric Dynamic Profiling, and Traffic Isolation
- Fibre Channel and FCoE Connectivity to VMware ESX hardware
- Per-port SAN Routing
- Fabric-assisted storage replication over distance
- 1, 2, 4, 8, and 10 Gbps Fibre Channel
- 1 and 10 Gbps Ethernet including CEE enhancements

Brocade Access Gateway

Key features of the blades in Access Gateway mode used in the reference architecture include:

- N_Port clustering
- F_Port Trunking
- Path failover

Brocade HBA

Key features of the Brocade HBA used in the reference architecture include:

- 8 Gbps FC
- Adaptive Networking services

Brocade CNA

Key features of the Brocade CNA used in this reference architecture include:

- Server network connectivity for server I/O consolidation
- Powerful 10 Gbps solution with reduced TCO
- Server provisioning and SAN boot with fabric-based LUN discovery, authentication, and virtualization mobility

Blade Servers

Key features of blade server switches used in the reference architecture include:

- Integrated chassis management
- Brocade embedded FC switches
- Brocade FC HBA with SAN boot support
- Brocade CNA to support convergence at the server edge
- Brocade 8000 Switch
- CDD switch
- Intel and AMD processors

VMware vSphere 4

Key features of VMware vSphere 4 used in the reference architecture include:

- ESX Server 4
- Virtual Machine File System (VMFS)
- Raw Device Mapping (RDM)
- VMotion
- VMware High Availability (HA)
- Distributed Resource Scheduler (DRS)
- Consolidated backup

To find out more about these features, visit the VMware corporate Web site, www.vmware.com.

© 2009 Brocade Communications Systems, Inc. All Rights Reserved. 09/09 GA-AB-071-01

Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, IronPoint, IronShield, IronView, IronWare, JetCore, NetIron, SecureIron, ServerIron, StorageX, and Turbolron are registered trademarks, and DCFM, Extraordinary Networks, and SAN Health are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services 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.