



DATA CENTER FABRIC

Using the Brocade DCX Backbone in a Virtualized Server Environment

Describes scenarios that demonstrate how to leverage the Brocade DCX Backbone with other Brocade products to meet three challenges of virtual server environments: scalability, performance, and manageability.

BROCADE

CONTENTS

Introduction.....	3
Brocade DCF Technology.....	3
NPIV	3
Access Gateway	4
Adaptive Networking.....	4
The Brocade DCX Backbone	4
Use Cases Scenarios	4
Use Case #1: Physical Server Consolidation at the Edge.....	4
Use Case #2: Virtual Server Consolidation at the Edge.....	5
Use Case #3: Scalable Performance at the Edge	7
Use Case #4: Adaptive Networking Services.....	8

INTRODUCTION

There is a fundamental transformation going on in the data center. Companies are intently focused on converting data into a strategic asset accelerating the amount of digital data they manage. This growth in data is changing data center architecture. The principles guiding this transformation include consolidation, virtualization, resource allocation based on service levels, and policy-based automation of data management.

The Brocade® Data Center Fabric (DCF) strategy was developed with these realities in mind. This strategy extends existing Storage Area Networks (SAN) fabrics into converged multiprotocol fabrics that simplify management of virtual infrastructure. The Brocade DCX® Backbone is a new class of fabric infrastructure. Essential requirements for the core of the Brocade DCF 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. The Brocade DCX design anticipates advanced technology such as Fibre Channel over Ethernet (FCoE), 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 to the application, logically connecting it to application service levels and data management policies. This improves infrastructure utilization, simplifies provisioning, capacity planning and management, which reducing infrastructure costs and freeing more dollars for companies to investment in business growth.

The Brocade DCX “backbone-class” performance is mandatory for deploying virtual servers:

- Up to 384 ports running at 8 Gbit/sec
- At least 4x increase in switching performance compared to current SAN directors
- An innovative 1 Tbit/sec Inter-Chassis Link (ICL)

Increased server utilization means increased bandwidth, particularly at the fabric core. Growing demand for server virtualization, combined with new blade servers designed for 8-core Control Processor Units (CPUs) (and greater) mean that the fabric core must be able to scale non-disruptively to keep up with ever higher ratios of virtual machines per server.

BROCADE DCF TECHNOLOGY

The Brocade DCF architecture integrates products and technologies that improve performance, scalability, and manageability for virtual server environments. The products and technologies below are essential components for building a cost-efficient data center fabric for server virtualization.

NPIV

Virtual servers require secure access to storage in the same way as physical servers do. However, a single physical server connection is unable to provide independent storage access to individual virtual servers. Instead, all storage ports and Logical Unit Numbers (LUNs, or storage) are exposed to all virtual machines, reducing security and manageability.

N_Port ID Virtualization (NPIV) is an ANSI standard designed to solve this problem. With NPIV, a Host Bus Adapter (HBA) installed in one physical server can provide up to 255 unique World Wide Port Names (WWPNs) for use by virtual servers. Fabric switches with NPIV support can then assign unique fabric IDs to each virtual server as they log in to the fabric. With NPIV support, standard fabric zoning and storage LUN masking can be used with Virtual Machines (VMs) to isolate storage ports and LUNs to the appropriate virtual server just as they are with physical servers.

Access Gateway

The Brocade Access Gateway option aggregates Fibre Channel (FC) traffic from all the blade servers in a blade server chassis without consuming scarce switch addresses (domain IDs). It greatly improves scalability of blade server environments by eliminating a large number of 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 impact on the fabric. F_Port Trunking, high bandwidth utilization, and automatic link failover simplify management and reduce costs. With NPIV support, storage management is extended past the physical server HBA to the virtual machine.

Adaptive Networking

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. Brocade Adaptive Networking services create multiple logical connections over a single physical link, which efficiently isolates traffic by Quality of Service (QoS) priority (High, Medium, or Low) and efficiently allocates fabric resources when connecting virtual servers to virtual storage. Adaptive Networking features include Ingress Rate Limiting, Congestion Detection and Management, Traffic Isolation, and Fabric Dynamic Profiling, such as "Top Talkers."

These services simplify management and ensure high utilization of fabric resources. Multiple virtual channels per physical link can be used with server SAN boot, storage I/O, and server cluster traffic to isolate their dynamic behavior 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.

The Brocade DCX Backbone

At the core of the Brocade DCF is the Brocade DCX Backbone. The Brocade DCX delivers the demanding performance required by server consolidation and provides "no-compromise" virtual server deployment. The Brocade DCX Backbone can scale to meet large increases in bandwidth demand created by the ever-growing number of virtual machine per server. In future, the Brocade DCX Backbone architecture will non-disruptively integrate FCoE and CEE for VMware cluster connectivity.

USE CASES SCENARIOS

The following use cases demonstrate how to leverage the Brocade DCX Backbone with other Brocade products to meet three challenges of virtual server environments: scalability, performance, and manageability.

The Brocade DCX Backbone can support up to 384 full-performance, non-blocking ports of 8 Gbit/sec Fibre Channel at the core; and each Brocade DCX port can support up to 255 NPIV addresses (connections). Clearly, this simple-to-manage, "single-domain" fabric can consolidate impressive numbers of physical servers. And making it a simple-to-manage building block for consolidating virtual servers.

Use Case #1: Physical Server Consolidation at the Edge

Physical server consolidation has great economic benefits. For example, a blade server chassis can consolidate as many as 14 servers into 9U of rack space, including space for Ethernet and Fibre Channel switches. These provide cluster communications and SAN storage connections to the blade servers.

Looking back to our early adoption of blade servers, it is easy to see that Brocade appreciated their impact on the fabric. The problem was that many small-port-count switches quickly exhausted the Fibre Channel switch, or domain ID, addresses, which are limited to 239. Brocade introduced the Access Gateway feature in Fabric OS® (FOS) to improve performance, reduce management complexity, and boost fabric scalability at the blade server edge. As shown in Figure 1, the Access Gateway can be used with existing 4 Gbit/sec embedded Fibre Channel switches in blade servers—converting them from a switch to an "N_Port cluster." On the back end of the Access Gateway, an F_Port connects to each blade server in the chassis. On the

front end, the Access Gateway presents multiple N_Ports (from 4 to 6 depending on the model), which can be aggregated into a load-balanced F_Port trunk attached to the core. NPIV support in the Brocade DCX Backbone allows multiple blade server N_Ports to log in on a single Brocade DCX switch port.




Figure 1. Server consolidation at the edge

There are several advantages to using Access Gateway mode to collapse the fabric edge. First, the Fibre Channel address space is used more efficiently, greatly improving scalability. Rather than consume an address from the smaller domain address space for every 8 to 14 servers, the much larger array and port address spaces (approximately 65,000 addresses) can be used, greatly increasing scalability. Second, server administrators can configure and manage the Access Gateway independent of the storage administrator's management of the SAN. And third, the Access Gateway provides load balancing between back-end and front-end ports using F_Port trunking, which delivers high bandwidth utilization without overprovisioning the number of physical links. If a link in the F_Port trunk fails, traffic is automatically rebalanced over the remaining links in the trunk.

NOTE: The Brocade 200E switch also supports Access Gateway mode, so that standard 1U servers can be consolidated at the edge with the same benefits as seen with blade servers. The Access Gateway feature will be added to other Brocade platforms in future.

Use Case #2: Virtual Server Consolidation at the Edge

A natural next step to increase the payoff from server consolidation is to deploy server virtualization. Server virtualization increases server utilization by creating multiple virtual servers in a single physical server. This combines multiple application workloads onto a single server, which improves server utilization and reduces the number of physical servers required. To simplify management, pools of physical servers can be created and workload management software can automatically move a virtual server to a different physical server to accommodate changes in application workload. Although virtual servers greatly improve server utilization and decrease costs, they create new management challenges for administrators.

Since a physical server now supports multiple applications, the fabric bandwidth per physical server increases, which in turn affects network performance, particularly at the core. There can be issues with zoning that require all storage LUNs connected to the physical server to be exposed to all the virtual machines, which increases risks and can decrease fabric performance. When creating physical server clusters, VMs can migrate to any physical server in the pool. And, all storage LUNs available to all physical servers become visible to all the virtual machines in a fabric—which further reduces access control and security.




Figure 2. Server consolidation at the edge

The solution to this problem is N_Port ID Virtualization (NPIV). With NPIV, each Virtual Machine is assigned its own virtual N_Port (or V_Port), allowing fabric zoning to once again ensure that VMs see only the storage port assigned to their own storage LUNs. Note that both the server HBA and the fabric switch ports must support NPIV.

Virtual server management software frequently supports migration of a virtual machine to a different physical server. The migration process must also move the assigned WWPN with the virtual machine. Many virtual servers use configuration files, called “personality profiles,” which store virtual machine configuration information. The configuration files are always accessible by the Virtual Machine no matter what physical server is currently hosting it. In this case, the VM rereads its original configuration file, retaining its original WWPN and so no changes to fabric zoning are needed.

Use Case #3: Scalable Performance at the Edge

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

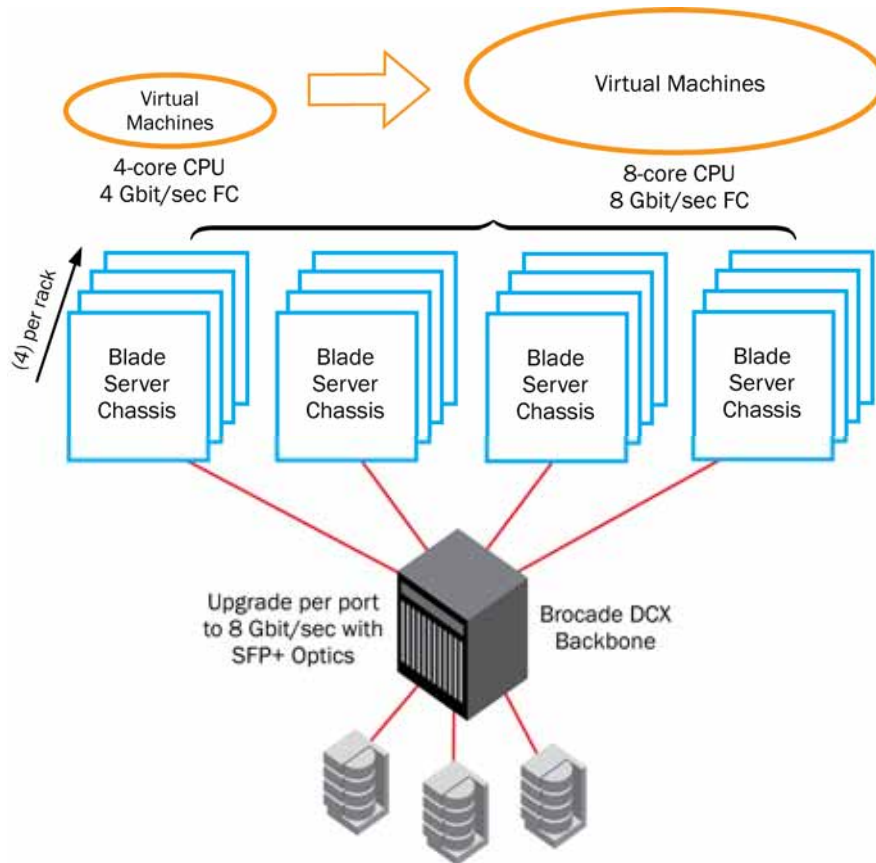


Figure 3. Scaling virtual servers with the Brocade DCX 8 Gbit/sec technology

The next-generation blade server chassis will increase VM density by a factor of two, three, or greater beyond the current levels. With its integrated 8 Gbit/sec technology, the Brocade DCX Backbone running Fabric OS 6.0 or higher, can easily scale its performance to meet this demand. Figure 3 shows a Brocade DCX initially connecting to today's 4 Gbit/sec, 4-core CPU blade servers. Since the Brocade DCX supports up to 255 NPIV addresses per physical switch port, the unique 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 Gbit/sec connectivity.

The upgrade to 8 Gbit/sec service can be accomplished at the port level by replacing the 4 Gbit/sec SFP optics with 8 Gbit/sec SFP+ optics. The upgrade is at the port level (not the card level), easy to implement, and requires minimal disruption. The procedure is simple: unplug the cable and 4 Gbit/sec optics and plug in the new cable and 8 Gbit/sec SFP+ optics.

Use Case #4: Adaptive Networking Services

Virtual Machine workloads can have performance requirements that differ from their application service level, because performance requirements are dictated by business needs. For example, a high-workload application may have a medium business importance, while another application creates a lower workload, but has higher business importance. Adaptive Networking services can be used to separate resource requirements (workload) from service levels (QoS).

For example, in Figure 4 virtual servers are deployed for various workloads on blade servers and 1U servers. The blade servers support 5 virtual servers per physical server, while the 1U servers support 10 virtual servers per physical server due to lower application workloads. However, the applications running on the blade servers have a lower service level (Medium) than the 1U server applications (High). Adaptive Networking can be used to assign a medium QoS to the blade server applications and a high QoS to the 1U server applications.

Traffic for each QoS priority is assigned to specific virtual channels across all switches, maintaining the QoS priority across the fabric. QoS ensures that the data flow from each virtual server has the same priority throughout the fabric, preventing unexpected congestion in the fabric when virtual servers migrate between physical servers.




Figure 4. Adaptive Networking services simplify fabric management

Another Adaptive Networking service, Traffic Isolation, can be used to dedicate Inter-Switch Link (ISL) paths to applications that may require large amounts of sustained bandwidth. The workload from the DBMS servers can cause high bandwidth requirements with large sustained peaks. To ensure that this traffic does not interfere with virtual server application traffic, Traffic Isolation is deployed on certain ISL connections. Traffic Isolation constrains the DBMS traffic to specific ISLs. Bursts in DBMS traffic now flow smoothly without the risk of an unexpected bottleneck. A bottleneck caused by a DBMS data flow could trigger a needless virtual machine migration if automatic workload migration is used to overcome virtual server resource bottlenecks. Traffic Isolation can avoid needless "thrashing" of virtual servers.

Using Adaptive Networking services with the Brocade DCX Backbone improves resource utilization, simplifies management as virtual servers migrate between physical servers, and avoids unnecessary virtual server migrations.

© 2008 Brocade Communications Systems, Inc. All Rights Reserved. 08/08 GA-TB-079-01

Brocade, the B-wing symbol, DCX, Fabric OS, File Lifecycle Manager, MyView, and StorageX are registered trademarks, and DCFM 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.