



DATA CENTER

Technical Brief: Adaptive Networking—Advanced Data Center Fabric Technology

The Brocade data center fabric is a foundation and evolutionary strategy for designing, building, and managing enterprise data centers. A critical advanced technology integrated into the data center fabric is Brocade Adaptive Networking services. As application consolidation accelerates, service level management gains in importance, Adaptive Networking provides Quality of Service (QoS) to ensure fabric support for critical application service levels.

BROCADE

CONTENTS

Executive Summary	3
The Evolving Data Center	4
The Brocade Data Center Fabric	5
Converged Protocols	7
Brocade's Virtual Channel Technology	8
Brocade Adaptive Networking Services Overview	8
Application-Aware Fabrics	8
Adaptive Networking Overview	10
Device Virtualization and Dynamic Workloads	10
Virtual Channels	11
Adaptive Networking Architecture	13
ASIC Functionality.....	13
Adaptive Networking Services	14
Quality of Service.....	14
Traffic Management Services.....	15
Fabric Dynamic Profiling	17
Resource Recovery Services (Future)	17
Management of Adaptive Networking Services	18
Adaptive Networking in Fabric OS 6.0 and Beyond	19

EXECUTIVE SUMMARY

Data centers are being transformed from static stovepipes of dedicated servers and storage hosting fixed applications to virtual environments, in which applications run on virtual servers sharing compute resources and storing their data in virtual storage pools. A transformation brought about by the relentless growth in digital data, Web content, and the demands of a global supply chain. Just to keep up, data centers have to dramatically increase the utilization of physical server and storage assets, while at the same time simplifying infrastructure management. A compelling benefit of increased device utilization in a data center virtual environment is reduction in power, cooling and floor space, which are becoming scarcer and more expensive. Much higher utilization of high-end storage is enabled by much better connectivity to high-capacity arrays.

Two significant technologies, virtual servers and virtual storage, are essential for this transformation but they introduce new complexity into the data center. As the data center infrastructure becomes "virtual," applications are dynamically provisioned on virtual machines and can be quickly moved from one virtual machine to another as their workloads change over time. Virtual storage applications automatically move data from one storage tier to another, actions dictated by access patterns and data retention policies. This dynamic movement of application workloads and data can create unexpected bottlenecks, which in turn cause unpredictable congestion in the networks hosting server-to-storage traffic, server-to-server clusters, and storage-to-storage replication. Provisioning, configuration management, capacity planning, security, and fault isolation are all more complex in a virtual environment.

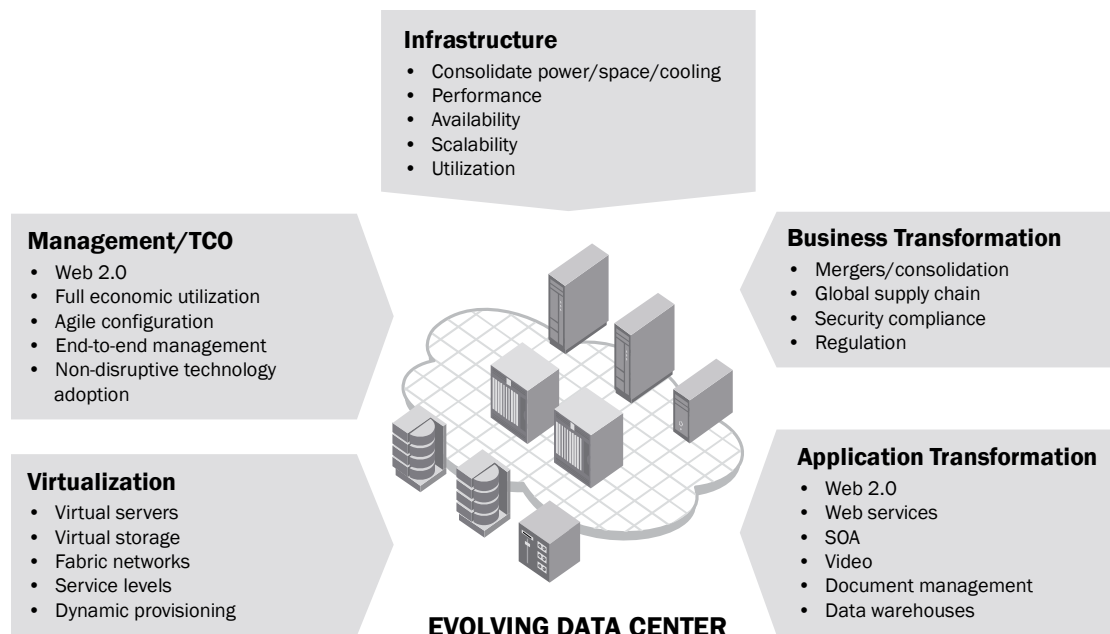


Figure 1. Forces driving data center transformation and their affect on the infrastructure

The Brocade® data center fabric strategy anticipates this transformation and its impact on the physical infrastructure. Adaptive Networking is part of the advanced fabric services in the Brocade data center fabric. It uses deep network intelligence to anticipate congestion and to dynamically make adjustments in the fabric so that application traffic continues to flow. Adaptive Networking will also be integrated with virtual server and storage management applications for provisioning, configuration management, capacity planning, and fault isolation to simplify data center operations and management.

Adaptive Networking introduces four new networking services: Quality of Service (QoS), Traffic Management, Fabric Dynamic Profiling, and in the future, Resource Recovery:

- **Quality of Service** defines high, medium, and low priorities of service; applications are assigned a QoS priority using zoning (Source ID/Destination ID pairs). Today, with Brocade Host Bus Adapter (HBA) and Converged Network Adapter (CNA) products, you can assign QoS priorities at the application level on the server.
- **Traffic Management** implements automatic rate limiting and advanced queuing priority algorithms to remove congestion and can dedicate bandwidth to specific applications so that overall bandwidth utilization is improved.
- **Fabric Dynamic Profiling** provides knowledge about end-to-end resource use across the converged fabric, supporting traffic management and capacity planning.
- **Resource Recovery** can detect depleted buffer credits as part of the flow control mechanism and “slow-drain” devices to prevent resource exhaustion and application disruption.

Whereas virtual servers and storage make fabrics more complex to configure and manage, Brocade Adaptive Networking services will simplify provisioning, configuration management, capacity planning, security management, and fault isolation for the evolving virtual data center.

THE EVOLVING DATA CENTER

With the emergence of Fibre Channel (FC) Storage Area Networks (SANs) in the late 1990s, Fibre Channel demonstrated its ability to improve performance, maintain mission-critical levels of availability, and enable enormous growth in affordable storage. The underlying technology, a fabric network, delivers deterministic latency, high performance, scalable bandwidth (doubling every two years), and very high availability (99.999 percent). The combination of these features is what differentiates a fabric network from a “best-effort” communications network, such as Ethernet.

A logical question to ask is: Where else can fabric networks with their proven mission-critical capabilities be used?

Brocade asked partners and customers to disclose their technology strategies and their requirements for transforming the data center; this information helped Brocade frame a technology strategy. Five fundamental business drivers transforming the data center were identified:

- Relentless growth of digital data
- Server and workload consolidation
- Service-oriented application architectures
- Rapid and flexible application deployment
- Agile infrastructure and scalable data management (to quickly and economically support ever faster business transformation)

Brocade recognized the singular advantage of fabric network technology in the future data center. Advanced fabric intelligence can simplify management and improve fabric utilization, while delivering massive performance, constant availability, and non-disruptive scalability—the essential qualities in virtual server and virtual storage environments.

One important advanced technology integrated into the data center fabric is Brocade Adaptive Networking. Adaptive Networking services will become more and more important as data center fabrics are used to collapse server-to-storage, server-to-server and storage-to-storage connectivity onto a common network connecting virtualized devices.

THE BROCADE DATA CENTER FABRIC

Brocade's strategy anticipates the transformation of the data center according to seven interlocking principles that apply to data center infrastructure:

- Converge protocols at an intelligent “access edge” into a high-performance core enabling a logical fabric to consolidate (pool) physical resources in the data center
- Extend physical connectivity to devices at the intelligent access edge for virtual connections to pools of shared resources.
- Integrate Virtual Fabrics to provide flexible management, traffic isolation, and switch partitioning.
- Integrate application service levels with the fabric for agile provisioning, high resource utilization, and Adaptive Networking services.
- Integrate fabric-hosted encryption services for disk and tape with Frame Redirection for simple configuration and no traffic disruption.
- Integrate Fibre Channel over Ethernet (FCoE) and Converged Enhanced Ethernet (CEE) at 10 Gigabits per second (Gbps) for physical consolidation at the server edge
- Extend data management policies into the fabric to unify, control, and scale data mobility applications (replication, migration, copy, and so on).

The Brocade data center fabric extends Brocade's technology leadership and product portfolio, consolidating fabric technologies onto a unified logical fabric. The management layer connects deeply integrated device and fabric intelligence with workload management, provisioning, and capacity planning tools provided by our storage and server virtualization partners. The benefits are simple management, lower operating cost, and a dynamic infrastructure that can evolve at a speed that matches the speed of business innovation.

The transformation to physical resource consolidation in the data center has a number of impacts. As more resources are concentrated, multiple workloads share the same physical resources, making it essential to have continuous availability, high performance, and non-disruptive scalability. These qualities are the hallmark of fabric networks in general and Brocade products in particular. Figure 2 illustrates how these components work together.

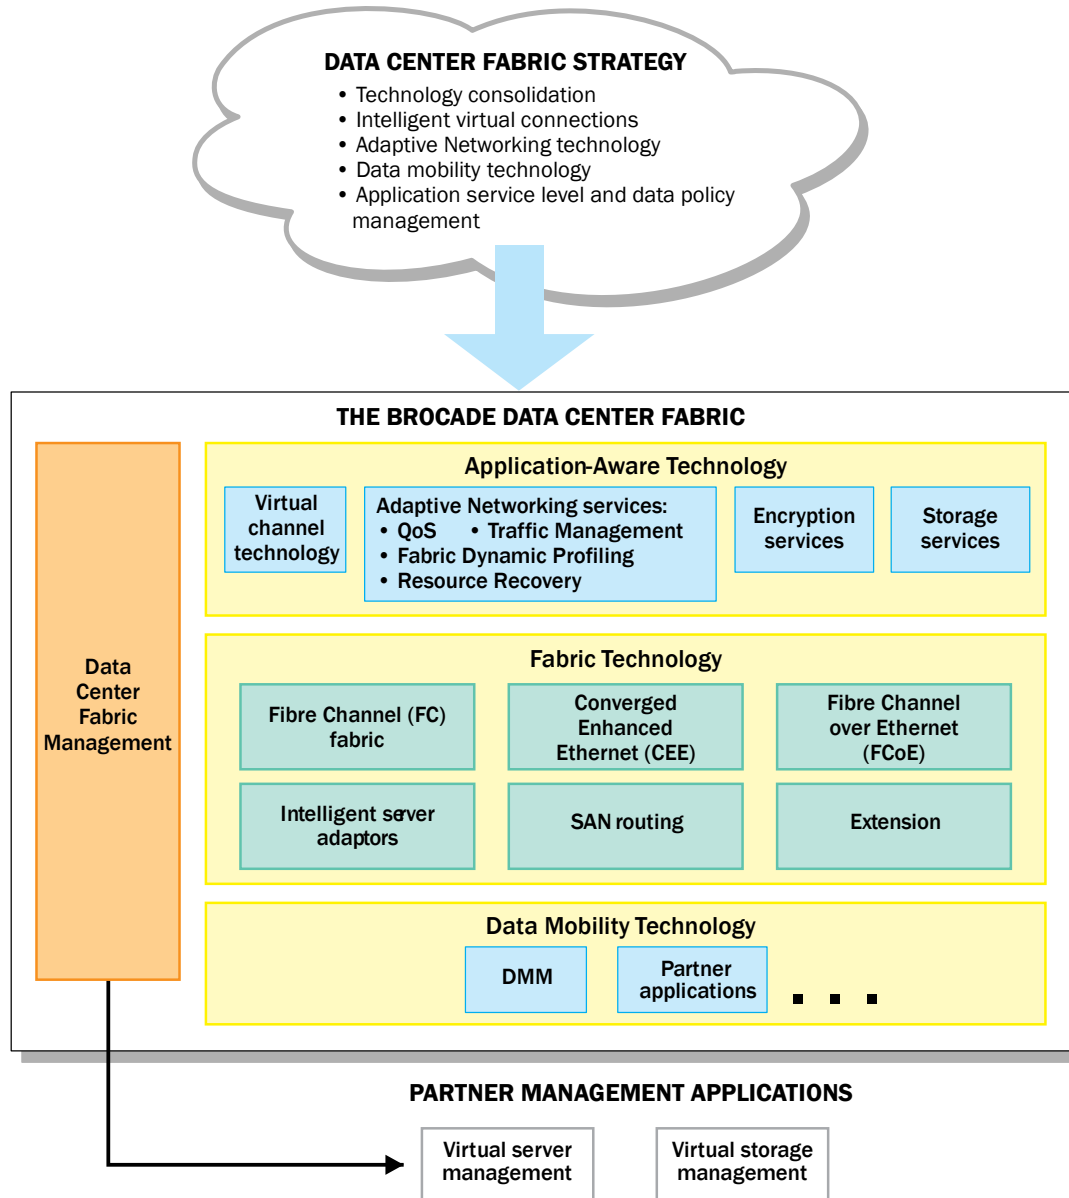


Figure 2. Data center fabric technology architecture

Converged Protocols

Applications, not the network, should control access to and management of their data. But it becomes more difficult for applications to access their data as new networking protocols and infrastructure virtualization are added to the data center. To simplify data access, the Brocade data center fabric extends a proven fabric technology to accomplish consolidation of networks and protocols, as illustrated in Figure 3.

Enterprises actually deploy and manage many physical networks: separate server-to-storage (SANs), server-to-server (clusters), and storage-to-storage replication networks are common within and between data centers. Brocade's data center fabric strategy consolidates these networks onto a single fabric that can non-disruptively integrate emerging technology, such as Fibre Channel over Ethernet (FCoE) and Converged Enhanced Ethernet (CEE).

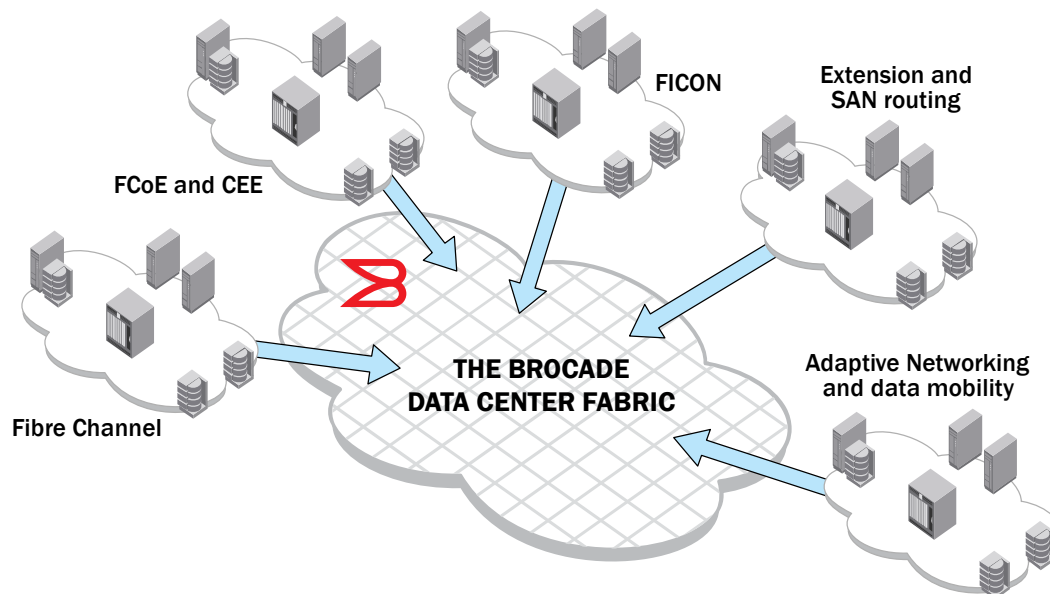


Figure 3. The Brocade data center fabric delivers fabric technology consolidation

Brocade's Virtual Channel Technology

Virtual channels in a physical fabric improve network utilization, simplify virtual server and storage provisioning, and reduce network configuration effort as application workloads are dynamically moved between virtual servers. Brocade's virtual channel technology provides intelligent virtual connections designed to logically connect virtual servers to virtual storage. A virtual channel allows physical resources in a physical link to be logically divided, providing fine-grained management and control of the fabric resources. An essential characteristic of Brocade's virtual channel technology is the isolation of data flow dynamics within the virtual channel. This allows multiple virtual server applications with different data flow dynamics to reside on the same physical link without disrupting each other. Figure 4 helps illustrate this concept.

Fabric intelligence about resource utilization must extend end-to-end across the data path and across heterogeneous protocols (Fibre Channel, iSCSI, FCoE) and not be isolated to a single network device. Adaptive Networking services help virtual channels span protocols, resulting in a uniform, logical data path within and between data centers.

Brocade Adaptive Networking Services Overview

Adaptive Networking services rely on Brocade's virtual channel technology. For example, QoS allows fine-grained allocation of fabric resources to applications, while categorizing the relative importance of the application (high, medium, low). Using Quality of Service is easy: adding a special prefix to a zone name classifies the zone's traffic as having a high, medium or low QoS priority. Specific virtual channels are assigned to each QoS priority. When applications become dynamic, the QoS priority must follow the application as it moves between physical server and fabric connections. It is very important for the fabric to monitor resource use and to detect (or predict) congestion anywhere along the data path. Virtual channel technology makes this possible. When congestion is detected or predicted, QoS priorities help traffic management applications dynamically adjust resource usage to avoid network congestion.

Traffic Management services provide congestion management (ingress rate limiting and traffic shaping) at network ingress points to ensure application service levels are enforced. They can also provide dedicated bandwidth to high-performance applications for more effective fabric utilization.

Fabric Dynamic Profiling services provide fine-grained, end-to-end analysis of individual data flows and their resource use (bandwidth, BB-credits, queue time), providing deep information about fabric performance, which is essential for traffic management in virtual server environments.

Application-Aware Fabrics

The Brocade data center fabric includes an intelligent edge with a point of presence in the server, with a Brocade server adapter (either a Brocade HBA or CNA). Applications can use the Brocade server adapters to communicate with fabric applications to request application service levels from the fabric. This extends the physical fabric connection at the server into an application-aware, logical connection—which is integrated with intelligent QoS, traffic management, and security services applications. The fabric uses its intelligence to monitor resource usage and to ensure it can deliver the service level requested by the application. Deeper knowledge about application requirements improves fabric resource utilization and capacity planning and simplifies provisioning, configuration, and fault isolation.

To summarize, the Brocade data center fabric creates a common, logical fabric that converges multiple networks and protocols. *It delivers a dynamic, cost-effective, scalable, and secure fabric connecting virtual servers and storage.* Data center fabric management can integrate with workload provisioning, configuration management, and provisioning applications, as shown in Figure 4.

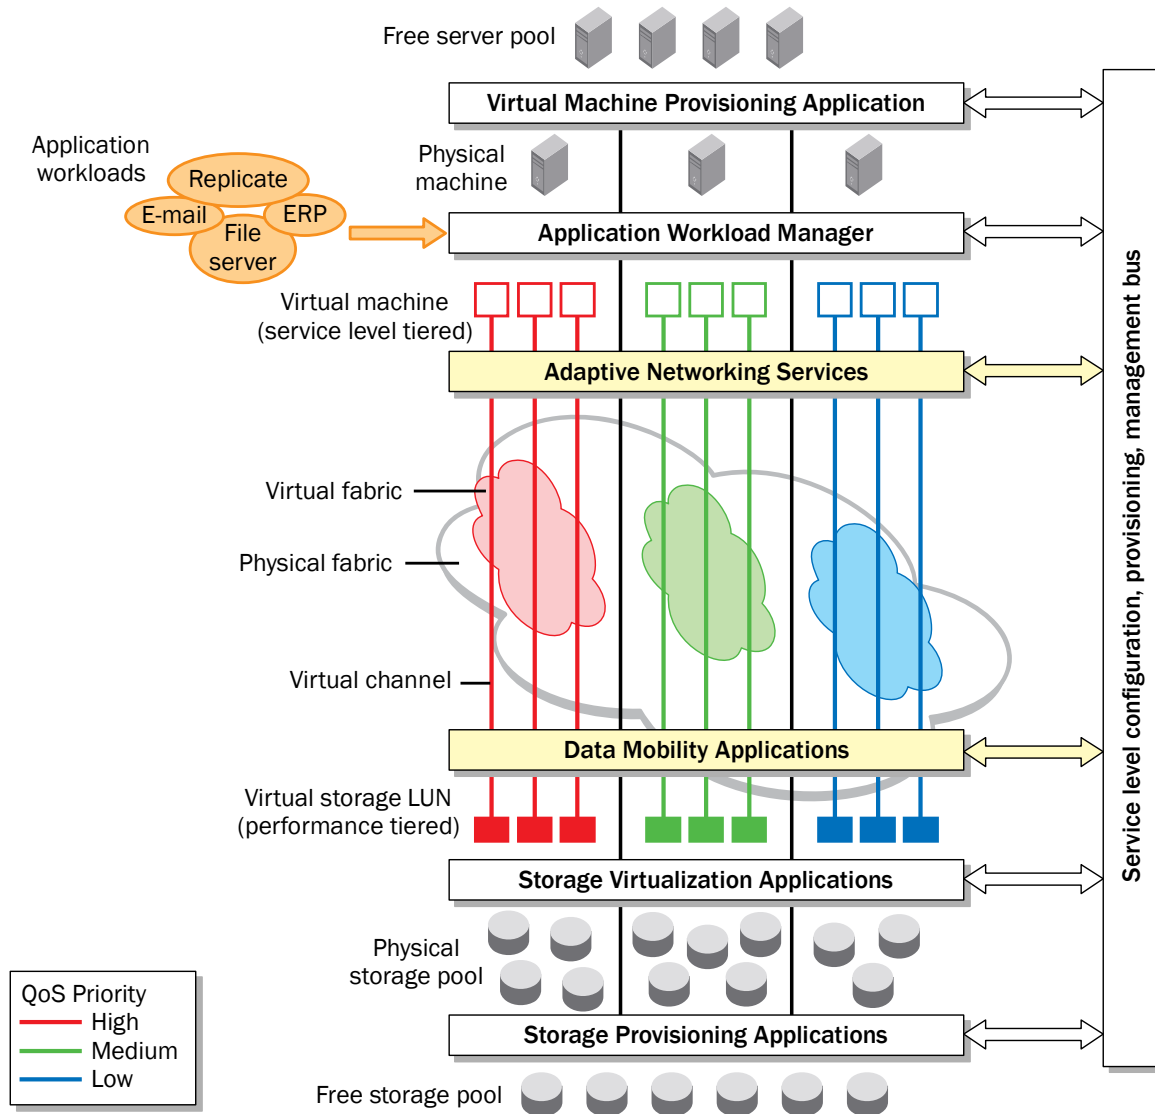


Figure 4. The Brocade data center fabric aligns with data center transformation

ADAPTIVE NETWORKING OVERVIEW

Some of the increase in fabric complexity comes from the natural growth of storage, as more business is conducted using digital information. But the use of server and storage virtualization to enable cost-effective sharing of pooled physical resources (computing, storage) also increases complexity. It is becoming more difficult to understand the dynamic behavior of the fabric as traffic patterns shift more quickly. This transformation defines important requirements for adaptive networking.

Device Virtualization and Dynamic Workloads

Until recently, servers and storage connected to the fabric were isolated physical resources. In most cases, a physical server ran one application workload, which resulted in well-defined network bandwidth requirements. But storage virtualization and server virtualization have replaced the static, physical connections between servers and storage with virtual machines and virtual storage Logical Unit Numbers (LUNs). These new technologies increase the complexity involved in designing and managing the fabric. For example, shared network resources (Inter-Switch Links (ISLs) and F_Ports using N_Port ID Virtualization (NPIV)) are now subject to variable workloads as applications are shifted from one virtual machine to another and as storage is moved from one virtual storage LUN to another. This can create unexpected congestion on shared fabric links if resource exhaustion occurs anywhere along the data path. The data center fabric has to be able to allocate its resources flexibly and to adapt resource use as workloads move between virtual hosts and virtual storage devices.

Virtual Channels

An essential fabric technology required to connect virtual servers to virtual storage is virtual channels. Virtual channels create multiple logical data paths across a single physical link or connection. They are allocated their own network resources (queues, flow control mechanisms such as BB-credits). *Virtual channel technology is the fundamental building block used to construct Adaptive Networking services.*

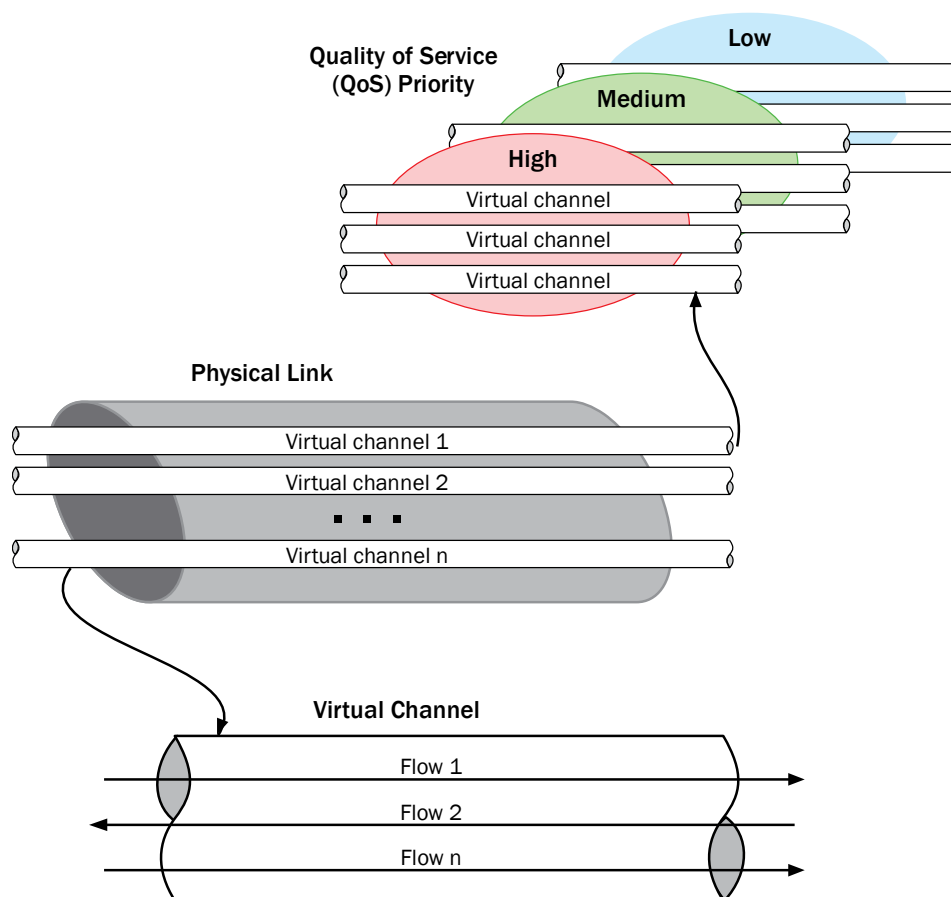


Figure 5. Virtual channels within a physical link and QoS priorities

Adaptive networking technology increases the fabric's intelligence by adding new services that automatically react to changes in virtual server workloads. This allows the fabric to increase the utilization of shared physical resources and smoothly adjust them as changes occur in the data flows between virtual servers and virtual storage LUNs. If congestion occurs, flows can be rerouted or throttled according to application service levels to reduce traffic until congestion is relieved. Fabric capacity planning is enhanced, as historical, fine-grained information about physical resource utilization is available for trend analysis. Bottlenecks are easier to identify and resources are added at the point of maximum need. With the introduction of the Brocade 8 Gbps HBA products, Virtual Channels now extend into the server to directly link application service levels to QoS priority.

Brocade Adaptive Networking integrates innovation in ASIC design (virtual channels and in-depth resource monitoring tools), new firmware functions, and fabric applications to create a fabric-wide resource monitoring and management service within the fabric. The Adaptive Networking layer can dynamically allocate and adjust fabric resources to application workloads, ensuring that higher-priority workloads continue to flow even when congestion occurs. Adaptive Networking is also architected to act proactively and avoid pending congestion conditions when resource utilization unexpectedly spikes. The main services required for Adaptive Networking include QoS, Traffic Management, Fabric Dynamic Profiling, and automatic Resource Recovery.

Figure 6 shows examples of how future releases of Adaptive Networking could support differing granularity of QoS assignment using virtual channels. For example, a QoS priority could be assigned to the physical Host Bus Adapter (HBA), to a virtual machine running a single application, or to execution threads on a virtual machine on which some threads support server-to-server cluster messages over Ethernet, some support storage I/O over Fibre Channel, and others are used by a backup application for data protection. The virtual channel architecture provides a flexible way to apply QoS to applications in virtual server environments.

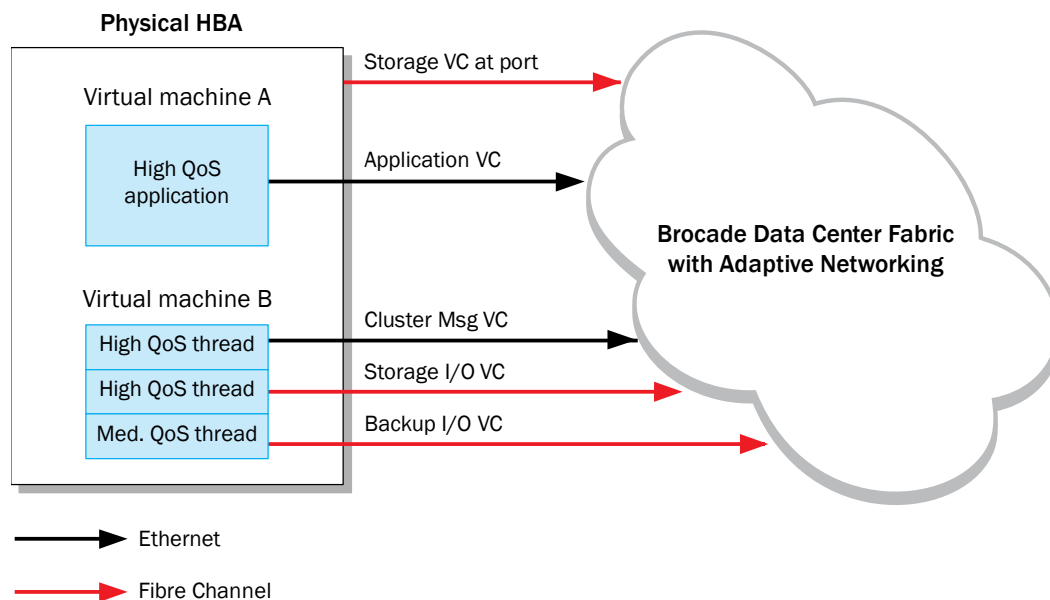


Figure 6. Converged fabric with virtual channel and QoS mappings

ADAPTIVE NETWORKING ARCHITECTURE

The Brocade architecture for Adaptive Networking includes deep integration of innovative ASIC technology with flexible firmware functions and new Adaptive Networking services—all managed from a unified management layer. Figure 7 shows how adaptive networking technology is implemented on the Brocade DCX® Backbone, a new class of network infrastructure that resides at the core of the Brocade Data Center Fabric. The Brocade DCX architecture supports Adaptive Networking services that exploit Brocade's deep technology integration from the bottom to the top.

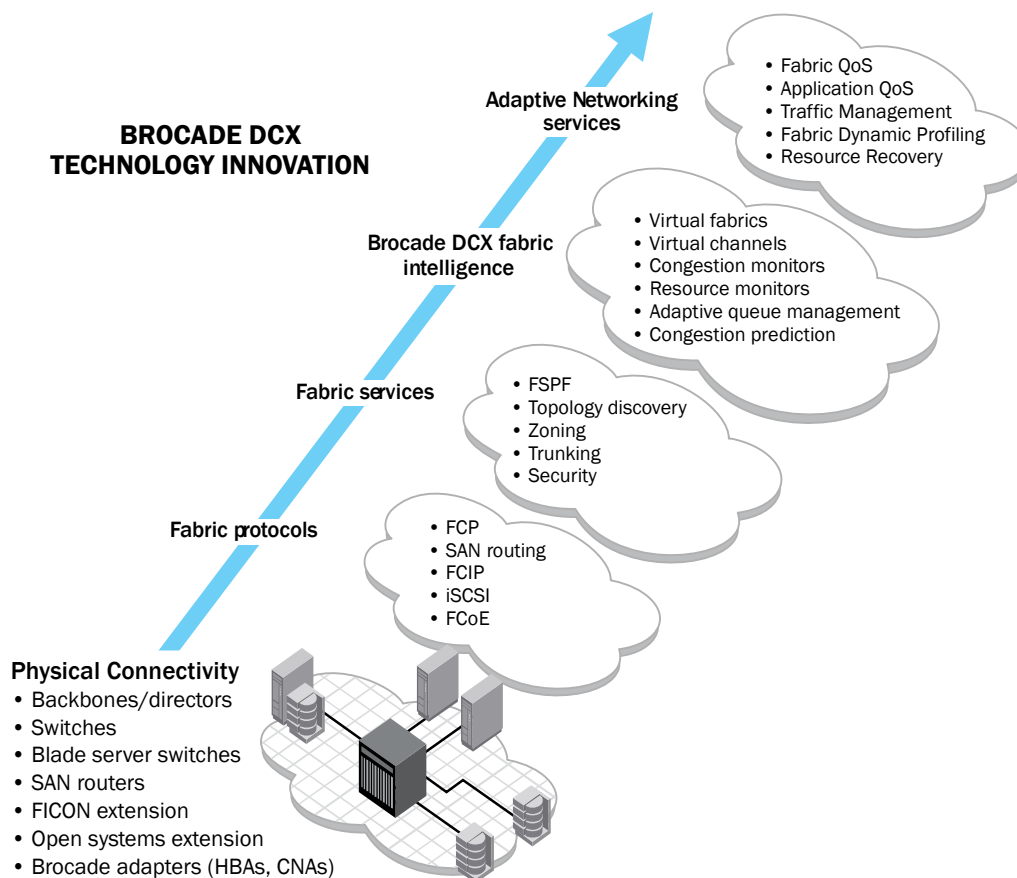


Figure 7. Brocade DCX technology integration for Adaptive Networking services.

NOTE: The terms “Brocade DCX Backbone” or “Brocade DCX” refer to both the Brocade DCX and the Brocade DCX-4S Backbone unless otherwise noted. In illustrations, the Brocade DCX is shown to represent both backbone platforms.

ASIC Functionality

Innovation in ASIC technology is a Brocade core competency. We introduced the core ASIC functions (virtual channels, destination-based QoS, and port and switch resource monitors) required by adaptive networking technology in the Brocade 4 Gbps Fibre Channel technology and the Condor ASIC. The Brocade DCX Backbone infrastructure with the Condor 2 ASIC extends virtual channel technology for more scalability and added resource monitoring functions, such as congestion detection, rate limiting, flow-based QoS, and additional port resource monitors.

Virtual channels implement the strategy of providing virtual access to physical resources. In simple terms, the resources of a physical link can now be allocated to a logical virtual channel connection. This enables pooling of the underlying physical link resources promoting higher utilization and lower cost infrastructure; and service levels are managed by monitoring resource allocation to virtual channels.

The Brocade 4 Gbps SAN product family supports 16 virtual channels per port, while Brocade DCX technology can support up to 32 virtual channels. In our SAN products, we dedicated one of the virtual channels for exclusive use by Class F frames, helping separate control frames (Class F) from data frames (Classes 2 and 3), which greatly improves fabric stability, scalability, and performance. In the Brocade DCX infrastructure, we reserve 16 virtual channels to support Adaptive Networking services. ISL, F_Port with NPIV, and Inter-Chassis Links (ICLs, not supported on the Brocade DCX-4S) all support virtual channel technology.

Adaptive Networking Services

As introduced earlier, Brocade's Adaptive Networking architecture includes several types of services:

- Quality of Service (QoS)
- Traffic Management
- Fabric Dynamic Profiling
- Resource Recovery

Quality of Service

The notion behind Quality of Service originated with TCP/IP communications networks. QoS was introduced for best-effort communication networks relying on TCP/IP to prioritize which frames are dropped by routers when a port becomes congested. Since these networks offered only "best effort" for data and message flows, congestion in the network randomly impacted all traffic via frame discards. However, application importance must be considered by the router in deciding which frames to discard. QoS allows network routers to decide which class of frames (tied to application workloads) to drop based on a tag in the frame header identifying the QoS priority of the frame. Lower-priority frames are dropped first, allowing more critical application workloads to continue through the network at points of congestion.

NOTE: TCP/IP QoS comes into play only when congestion is detected in a router's queues. When there is no congestion, frames flow freely as application workloads generate them and are not impacted by the QoS priority. In other words, QoS does not provide "privileged" access to network resources (pre-allocated bandwidth) for application workloads.

Fabric QoS

In a Brocade Fibre Channel environment, a fabric ingress port determines the QoS priority based on information contained in the frame header and administrator configuration of zoning services. For example, a fabric zone identifies a host initiator and the storage targets to which it needs a connection. A fabric zone establishes a logical association and thus represents a data flow through the fabric between the initiator and the target(s). To implement QoS for a data flow, the SAN administrator creates a special zone with a QoS priority prefix in the zone name. The QoS priority prefix defines the QoS priority, while the zone members (initiator and target(s)) identify the data flow associated with this QoS priority. In addition to specifying QoS at the zone level, QoS can be defined using the virtual fabric, the Logical SAN (LSAN), or even the contents of the frame payload.

Within the fabric, multiple virtual channels can be assigned to each of the three QoS priorities. Each virtual channel has its own input queues. The QoS queuing algorithm is the same for all switches in the fabric and provides weighted servicing of frames in the input buffers for all QoS priorities. The weighting depends on the QoS priority and provides more frame forwarding for the high-priority QoS class than for the low. Note that the low priority queue continues to be serviced (is not starved) when congestion occurs. This approach implements the same QoS level for a data flow across a fabric regardless of the number of switches, ISL, or IFL connections used. With multiple virtual channels per QoS priority, Resource Recovery services will in the future remove ISL congestion caused by a "slow drain" device.

Fabric QoS inspects the Source ID/Destination ID (SID/DID) of the frame at the fabric ingress ports (F_Port) and compares it to the list of special QoS zones. If it finds a match (that is, SID/DID are in a QoS zone), it uses the QoS "name tag" contained in the zone name to decide which QoS priority to assign to the flow (high, medium or low). If the zone name tag indicates a high QoS priority, the ingress port puts the frame into the queue for the virtual channels assigned to high priority traffic. At each switch in the data path, the same virtual channels are used to forward frames for a given QoS priority. And each switch schedules frame forwarding from the virtual channels using the same queuing algorithm. Hence, the frame receives the same frame forwarding priority across an arbitrary-sized fabric.

NOTE: Today extensions to zoning are used to define the QoS priority. A zone contains the World Wide Port Name (WWPN) of the host HBA and the WWPN of the storage port. The name server converts these addresses to the appropriate fabric IDs. When a frame enters the fabric through a port on the ingress switch, the name server looks up the Switch ID (SID) in the zoning WWPN/Fabric ID table. That tells the switch what zone this frame belongs to, and the zone name tells the switch what QoS priority to deploy.

Application QoS

The host can directly request a QoS priority if it has an interface to the fabric to do so. Brocade HBAs and CNAs can send a QoS priority request for a data flow (initiator, target) to the fabric QoS service interface:

- If the fabric has sufficient resources to meet the request, it returns an acceptance, along with information about which virtual channel(s) to use for forwarding the frame. The QoS request is made only once per data flow, not each time a frame is sent. The host adapter now knows to put the frame for this initiator/target into the adapter queue allocated to that QoS priority. In this case, the fabric ingress port (F_Port) does not assign the QoS priority, but is configured to accept the frame from the virtual channel it is received on and to forward it to the proper egress port on the same virtual channel.
- If the fabric is unable to meet the QoS service level requested by the host interface (insufficient physical resources), the host can accept a lower QoS. Adaptive network management applications can remember the requested QoS priority, and when physical resources permit can upgrade the QoS level by signaling the adapter to move the data flow to virtual channels assigned to the higher QoS priority.

Traffic Management Services

Traffic Management is a way to control and allocate network bandwidth, reducing latency along the data path and minimizing congestion at network ports. Traffic management provides several services, including ingress rate limiting, traffic isolation, and in the future, congestion management.

Understanding Ingress Rate Limiting

To understand Ingress Rate Limiting, think of traffic entering a congested freeway from an entrance ramp with a traffic light. Cars wait on the entrance ramp until signaled to enter, thereby throttling access to the freeway to reduce congestion. Traffic Management services can be used to regulate the rate of frame entry to a fabric when congestion has occurred or is about to occur. Application workloads are prioritized in advance so that the workloads of less-critical applications are throttled when congestion occurs. Lower-priority workloads are still serviced when shared network data paths are nearing physical resource exhaustion, but at a reduced rate to avoid congestion.

As physical data paths become less congested, the lower-priority applications are once again allowed to enter the network as frames become available. This is similar to the freeway entrance ramp keeping the green light on later at night, when traffic is light so cars do not have to wait to enter. Today, Ingress Rate Limiting can be set on an F_Port to restrict the received frames into the port to a preset bandwidth limit.

Traffic Isolation

A second Traffic Management service is traffic isolation. User-defined data flows can be allocated to specific ISLs without changing Fabric Shortest Path First (FSPF) routing. This ensures that the application's traffic always has network resources dedicated to it, which is valuable when sustained, high-bandwidth traffic is created by the workload. For example, it is useful in backup applications using physical or virtual tape devices, or for array-to-array replication applications providing synchronous or asynchronous block replication. Other application workloads avoid using the physical connections allocated for traffic isolation, in which case their performance is not impacted by the high-bandwidth workloads.

Congestion Management (Future)

A third category of Traffic Management service is congestion management, which is applied at the virtual channel level providing fine-grained management. Two types of congestion management are congestion control and congestion avoidance. These will be advanced services exploiting deep network intelligence about resource use at every port along the data path.

Congestion Control. Congestion control is reactive; and when congestion occurs, it sheds load from the network to eliminate the congestion. As mentioned above, today Adaptive Networking provides manual ingress rate limiting.

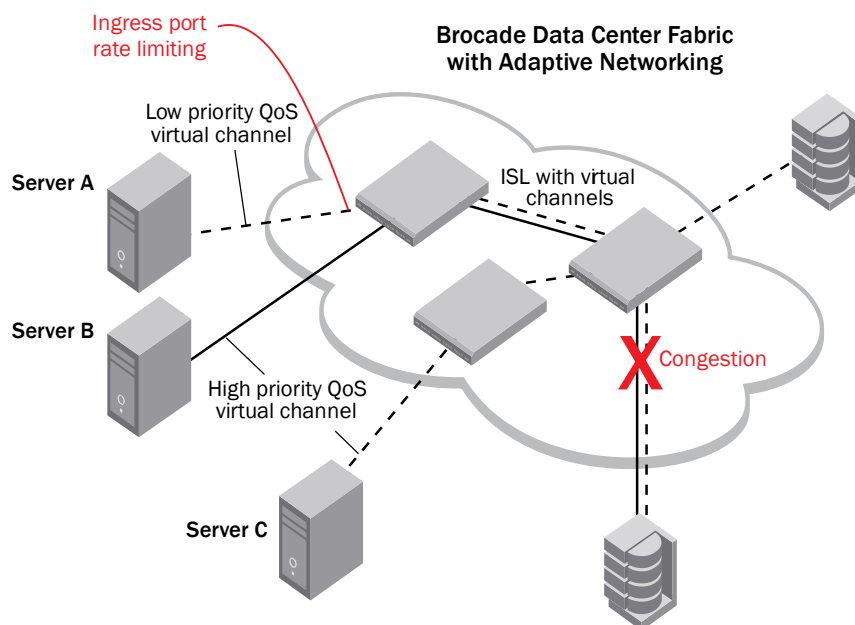


Figure 8. Virtual channels with QoS using ingress port rate limiting for congestion management

Automatic rate limiting will be a valuable improvement over manual methods. It will predict impending congestion and then apply a defined rate limit to the data flows moving through the congestion point until the congestion has cleared—at which point normal frame flow resumes. One algorithm for automatic rate limiting uses a decaying throttle rate for the ingress frames. As the rate of throttling increases, frame entry to the fabric is slowed, allowing the congestion to clear. Over a period of time, the throttling rate “decays,” allowing full frame entry rates to resume.

Congestion Avoidance. This is proactive and prevents congestion from occurring by closely monitoring physical resource use and comparing it against a high water mark. When resource utilization exceeds the high water mark, special congestion avoidance techniques are activated along the data path. Congestion avoidance requires deep network intelligence (such as Fabric Dynamic Profiling) to monitor physical resource utilization (queue depth, BB-credits, latency, and so on) at all ports in the fabric and heuristic models that can predict the onset of congestion at any point in the fabric. Congestion avoidance uses special queuing and rate limiting algorithms to reduce shared resource utilization until resource use falls below the high water mark. High priority application workloads are favored at the expense of lower priority ones allowing changes in resource allocation and application specific rate limiting to prevent congestion.

Fabric Dynamic Profiling

This Adaptive Networking service provides profiles about baseline resource utilization across the entire fabric. It identifies points of congestion, and monitors and reports on numerous statistics counters for monitoring physical resource utilization. This advanced reporting can be used to support provisioning, capacity planning, and end-to-end fault isolation tools that simplify fabric management. For example, the Top Talker profile is used by the Top Talker report in the Advanced Performance Monitoring feature.

The Top Talker report identifies which data flows (SID/DID) have the highest bandwidth utilization. Today, this profile can be used in two modes, F_Port Top Talkers and Fabric Top Talkers. F_Port Top Talkers ranks and then displays the highest-bandwidth data flows on a given F_Port. Fabric Top Talkers ranks and displays the highest-bandwidth data flows on all E_Ports for a selected switch. Top Talker information is generated on a port-by-port basis and can be processed to generate a fabric-wide tabulation of which application workloads consume the most fabric bandwidth. Knowing which applications are Top Talkers helps identify where to set manual ingress port rate limits to avoid congestion.

Resource Recovery Services (Future)

Buffer-to-Buffer (BB) credits is a flow control mechanism that represents the availability of resources at the receiving port. Fabrics are subject to conditions in which BB credit resources can be depleted, preventing traffic flow over shared data paths. Two conditions can cause this: damaged or missing R_RDY frames and "slow-drain" devices. The first condition can be overcome with BB credit recovery and the second by using dynamic virtual channel mapping of data flows.

BB Credit Recovery

In fabrics that rely on BB credits for flow control, missing or damaged credit recovery frames (R_RDY) can exhaust BB credits. The transmitting port will not send frames if there are zero BB credits for the receiving port. The receiving port signals resource availability by returning an R-RDY frame to the transmitting port after it has processed an incoming frame. Over time, BB credit loss can result in a network port being unable to send frames, since it no longer has BB credits for the receiving port. The Resource Recovery service uses in-band signaling (ANSI specifications define the BB_SCn primitives for this purpose) to allow a port to determine if lack of BB credits is due to lost or damaged R_RDY frames. If it is, then the ports can negotiate recovery of BB credits.

Dynamic Virtual Channel Mapping of Data Flows

Another Resource Recovery service is appropriate when a slow-drain device on a shared data path prevents other data flows from using that path. This condition can occur on ISL connections when a storage target is unable to process received frames quickly enough (often due to internal resource exhaustion, such as storage array cache misses, internal queue overflow, or bus saturation), resulting in BB credit exhaustion for the transmitting E_Port of the ISL. Without available BB credits, all other data flows using that ISL are halted. In a data path that traverses multiple switches over multiple ISLs, BB credit exhaustion propagates backward to the ingress ports for all data flows that use the congested ISLs.

Detection of a slow-drain condition before BB credit exhaustion occurs would allow that data flow to be shifted to a separate virtual channel with its own BB credit pool as long as the slow-drain condition exists. The remaining data flows sharing that ISL can continue through the network on the original virtual channel, without being held up. Dynamic virtual channel mapping prevents physical resource exhaustion in an end device (storage port) from disrupting other application workloads that do not rely on that resource.

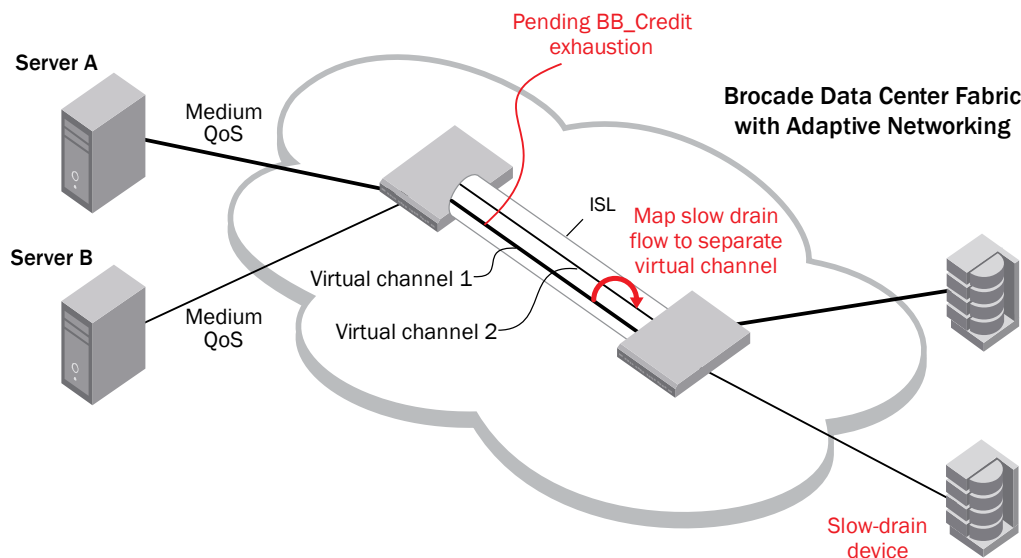


Figure 9. Dynamic virtual channel flow mapping applied to slow-drain devices

As shown in Figure 9, Servers A and B both have a medium QoS priority. One of the arrays is behaving as a slow drain and is about to exhaust BB credits on the ISL ports. Dynamic virtual channel flow mapping is invoked and moves the slow-drain flow between Server B and its array port from virtual channel 1 to virtual channel 2. Since each virtual channel has its own allocation of BB credits, virtual channel 1 quickly regains its BB credits, while virtual channel 2 isolates the slow-drain flow and prevents it from congesting the ISL traffic.

MANAGEMENT OF ADAPTIVE NETWORKING SERVICES

A primary goal of Adaptive Networking services is to simplify fabric management. As virtual servers and virtual storage are implemented in the data center, creating more dynamic application workloads, fabric intelligence has to provide more in-depth knowledge about the state of fabric-wide resource utilization. Then, the intelligence needs to associate resource use with application workloads (data flows) across all virtual devices. Consequently, device management and fabric management software— Brocade Data Center Fabric Manager (DCFM™), Brocade Fabric Manager, and Brocade Enterprise Fabric Connectivity Manager (EFCM)—have to be extended to provide simple configuration and management of Adaptive Networking services. As application QoS becomes available in Adaptive

Networking, the management layer will integrate with workload provisioning, capacity planning, and end-to-end fault isolation management tools available in higher-level management applications such as IBM Tivoli, HP OpenView, EMC ECC, Microsoft Management Console, and so on.

ADAPTIVE NETWORKING IN FABRIC OS 6.0 AND BEYOND

Fabric OS 6.0 and later releases offer the following Adaptive Networking services:

- Fabric QoS with high, medium, and low QoS priority for SID/DID pairs (FOS 6.0)
- Ingress Port Rate Limiting for F_Ports (FOS 6.0)
- Traffic Isolation for FC fabric (FOS 6.0)
- Traffic Isolation for FC Routing (FOS 6.1)

Brocade's plans to enhance the initial release of adaptive networking technology have the goal of extending application-aware service levels into the converged data center fabric. This requires Adaptive Networking services to work across multiple protocols and to provide application interfaces in the server.

Brocade will simplify how application workloads on virtual servers are provisioned and connected to advanced Adaptive Networking services, examples of which include automated bandwidth allocation, automatic ingress rate limiting, the ability to assign QoS priority at the HBA, application workload, and application or application thread level. *The overarching goal is to provide advanced Adaptive Networking services without adding complexity to provisioning and configuration.*

With the introduction of the innovative Access Gateway mode, now available for 8 Gbps embedded blade server switches and the Brocade 300 Switch, Brocade has partnered with virtual server vendors (VMware, Microsoft, and XenSource) to develop the required software interfaces that will simplify provisioning virtual servers, virtual channels, and virtual storage. Leveraging our investment in HBA technology, Brocade HBA and CAN server adapters allow applications to request fabric resources that are automatically provisioned by Adaptive Networking services. This evolution of adaptive networking technology is essential to our partners as they extend virtual server technology to more critical application workloads.

© 2009 Brocade Communications Systems, Inc. All Rights Reserved. 07/09 GA-TB-058-04

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 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.