

Dueling directors



Director-class switches have been used to maintain performance in large storage networks for years, but soon they'll act as the hubs that unify data center networks.

By Jacob Gsoedl

FIBRE CHANNEL (FC) as a technology has been relatively static over the past 10 years, and FC switch innovation has been incremental—from bandwidth support and additional features to increased resilience and availability—culminating in the high-end FC director platform. Aside from a few failed incursions by vendors like QLogic Corp. and others, Brocade Communications Systems Inc. and Cisco Systems Inc. now almost exclusively own the high-end Fibre Channel switch and director market. With their top-of-the-line platforms (the Brocade DCX Backbone and Cisco's MDS 9500 Multilayer Director Series), they duel for customers who require a combination of high FC performance and high availability. Features and suitability for existing infrastructure have typically been the primary director selection criteria. Enhancements related to Fibre Channel over Ethernet (FCoE), as well as the convergence of FC and Ethernet into a unified data center protocol, have added roadmaps and vendor strategies as relevant purchasing considerations. Protecting the new investment and ensuring its future are of paramount importance considering how profoundly storage-area networks (SANs) and data centers will be transformed.

The Brocade DCX Backbone and Cisco MDS 9500 Series have much in common. They're both chassis based and can be scaled by adding or changing hot-swappable line cards. With all components redundant and hot-swappable (blades, fans, power supplies), they present no single point of failure. From 1 Gbps/2 Gbps /4 Gbps /8 Gbps Fibre Channel to FICON, FC over Internet Protocol (FCIP) and Internet Protocol over FC (IPFC), and connectivity options for iSCSI (DCX via an iSCSI gateway and the MDS natively), they support all relevant storage networking protocols. With Brocade's Fabric Application Platform option and the Cisco MDS 9000 Storage Services Module (SSM), both vendors can turn their switches into intelligent app platforms for third-party services, such as EMC Corp.'s Invista for fabric-based storage virtualization and EMC RecoverPoint for fabric-based data protection. By way of the Brocade Data Center Fabric Manager (DCFM) and Cisco Data Center Network Manager (DCNM), both provide commensurate management applications to provision and proactively manage devices across multiple fabrics. Both vendors also support N_Port ID Virtualization (NPIV) to simplify connectivity, management and monitoring of proliferating virtual server environments. In addition, Cisco's ability to create completely isolated fabric topologies via its virtual SAN (VSAN) technology has been matched by Brocade's new Virtual Fabrics feature, which enables organizations to create logical groups of separately managed devices, ports and switches within a physical SAN.

"With the support of Virtual Fabrics, Brocade eliminated one of the competitive advantages Cisco had with VSANs."

—Bob Passmore,
research vice president, Gartner Inc.

"With the support of Virtual Fabrics, Brocade eliminated one of the competitive advantages Cisco had with VSANs," said Bob Passmore, research vice president at Stamford, Conn.-based Gartner Inc. Notwithstanding a long list of commonalities, Brocade and Cisco differ in some key areas and features, as well as in product strategy.

PRODUCT OVERVIEW

In an attempt to establish a product category that resides above traditional directors, Brocade doesn't categorize its DCX Backbone as a director-class product. However, with four times the chassis bandwidth, significant energy efficiency and several feature enhancements, the DCX Backbone can be viewed as a "green" Brocade 48000 Director on steroids. Among the DCX Backbone enhancements are inter-chassis links (ICL) that allow cascading two DCX Backbone switches via a 512 Gbps pipe, integrated Fibre Channel routing, fabric-based encryption, and the separation of control processor and core switching blades. Unlike the 48000 Director and the Cisco MDS 9500 Series, the DCX Backbone doesn't support

iSCSI natively and depends on external iSCSI gateways to interface with iSCSI SANs.

The DCX Backbone is available in two modular form factors. Built for large enterprise networks, the 14U Brocade DCX Backbone has eight vertical blade slots to provide up to 384 FC ports using Brocade-branded 4 Gbps or 8 Gbps small form-factor pluggables (SFPs). Built for midsized networks, the 8U Brocade DCX-4S Backbone has four horizontal blade slots to provide up to 192 Fibre Channel ports. Additionally, each chassis has two slots for redundant control processor blades and two more slots for redundant core switching blades that move traffic between blades.

With the Cisco MDS 9100 Series rack-mountable fabric switches at the low end and the midrange MDS 9200 Series—which provides MDS 9500-equivalent features at a smaller form factor and lower cost by sacrificing redundancy and scalability—the MDS 9500 director family tops Cisco's coherent FC switch product line. Cisco offers three MDS 9500 models: The MDS 9513, supporting a record 512 8 Gbps FC ports, with a total of nine slots for line cards and two slots for redundant supervisor modules, is targeted at enterprise networks with high port count requirements. For smaller networks and as edge devices, Cisco offers the MDS 9509 featuring nine slots, and the MDS 9506 with six slots.

Introduced by Cisco in 2002, the MDS 9500 is on track to follow the Catalyst 6500 Ethernet switch family in prolonging its life toward a decade, displaying longevity that's rare in the fast moving high-tech arena where assets usually depreciate in three to five years. "MDS 9506 and 9509 switches that shipped in 2002 still support 8 Gbps FC and will support FCoE," confirmed Paolo Perazzo, senior product line manager, Cisco's Data Center Switching Technology Group. In contrast, during the same period, Brocade has brought five different director platforms (12000, SilkWorm

FEATURE COMPARISON

	BROCADE DCX BACKBONE	CISCO SYSTEMS INC. MDS 9500
Models	*DCX Backbone (eight port blades) *DCX-4S Backbone (four port blades)	*MDS 9513 (11 port blades) *MDS 9509 (seven port blades) *MDS 9506 (four port blades)
Maximum port count	384	512
Bandwidth per slot	256 Gbps	96 Gbps
Biggest benefit	Fastest Fibre Channel (FC) director on the market. Significant lead in the total number of FC ports sold.	High degree of investment protection, long product lifecycle and coherent product line that uses the same NX-OS operating system
Fibre Channel over Ethernet (FCoE) support	Will be available once the standard is ratified	Available now via the Nexus 5000 switch; for the MDS 9500, Nexus 2000 and Nexus 7000, it will be released on standard ratification

24000, 48000, DCX Backbone and McData i10K) to market.

“Brocade is a smaller but very innovative company with shorter product lifecycles, partnering with best-of-breed vendors. Cisco, on the other hand, has a record of designing architectures that last for many years,” noted Bob Laliberte, an analyst at Enterprise Strategy Group (ESG) in Milford, Mass.

A case in point is the DCX Backbone. “The main difference between the 48000 and DCX is the separation and rearchitecture of the core switching and control processor blades into separate blades, which required a new chassis design,” said Bill Dunmire, senior product marketing manager at Brocade.

THE HIGH COST OF HIGH-END SWITCHES

PLATFORM	PRICE
Brocade DCX Backbone with eight 48-port 8 Gbps blades and full redundant configuration	\$628,000
Cisco MDS 9513 with eight 48-port 8 Gbps blades and full redundant configuration	\$593,000
Brocade DCX-4S with four 48-port 8 Gbps blades and full redundant configuration	\$328,000
Cisco MDS 9506 with four 48-port 8 Gbps blades and full redundant configuration	\$320,000

(List prices provided by Brocade/Cisco reseller)

COMPARING ARCHITECTURES

The Brocade DCX Backbone is based on a shared memory architecture where data moves from switching ASIC to switching ASIC along multiple internal ISLs that make up the path from an ingress port to an egress port. To load balance between these inter-ASIC links within the switch, the DCX Backbone relies on either exchange- or port-based routing. “Besides fewer components on blades, which reduces the likelihood of failure, in a shared memory architecture ASICs on the core switching blades talk to ASICs on port blades using the same protocol, minimizing protocol overhead,” Dunmire explained.

In comparison, the Cisco MDS 9500 leverages a crossbar architecture where frame forwarding is directly performed in ASICs on the line cards. The crossbar manages forwarding of packages, and a central arbiter ensures fairness and prioritization. While the MDS 9506 and MDS 9509 provide the fabric switching module and central arbiter on the supervisor blade, the MDS 9513 uses a separate pair of switching modules located in the back of the MDS 9513 chassis. “Unlike a shared memory architecture where traffic moves across internal switching ASICs along varying paths, resulting in varying latencies, in a crossbar architecture the latency between ports is consistent across all ports within the switch,” said Omar Sultan, solution manager, data center switching, data center solutions at Cisco.

Even though each vendor claims its architecture is superior, they each have their pros and cons. With the exception of a few vendor specific peculiarities, both platforms can be used to power the most mission-critical and largest SANs with comparable results and user experience; this is substantiated by Brocade and Cisco splitting the director market almost evenly. “The two products work very well and

by having competition, they're pushing new features and, in the long term, bring down the price," ESG's Laliberte said.

Nevertheless, there are noticeable differences between the two platforms. The DCX Backbone supports local switching, which allows traffic between ports on the same blade to be directly switched instead of having to go through the core switching module; this means lower latency for devices connected to the same blade and improved scalability by reducing the amount of traffic that has to pass through the core switching blades. Although Cisco rebuffs the local switching benefit, emphasizing bigger latency variances as a result of local switching, support for local switching in its latest Nexus platform suggests that the lack of local switching support in the MDS 9500 is a disadvantage.

In addition to reliability, performance and throughput are the most relevant attributes of a director platform. The Brocade DCX Backbone currently wins the raw throughput comparison with 256 Gbps throughput per slot vs. 96 Gbps for the Cisco MDS 9500. When combined with local switching, it can concurrently operate more ports at full 8 Gbps utilization than the MDS 9500, as verified by a February 2009 Miercom lab test (Report 090115B). As a result, the MDS 9500 depends to a greater degree on oversubscription than the DCX Backbone. In practical SAN reality, however, not all ports will operate at full 8 Gbps

In addition to reliability, performance and throughput are the most relevant attributes of a director platform.

rate, and the use of oversubscription combined with traffic prioritization and QoS makes the throughput difference less significant. In the past, increases in port and chassis throughput benefited mostly ISLs and, to a lesser degree, servers; but now the proliferation of virtual server environments definitely makes bandwidth capacity more relevant. "Server virtualization is a game changer, making oversubscription more problematic because physical servers running many virtual machines are more likely to fully utilize a SAN link," Gartner's Passmore said. Cisco confirmed that it's working on a next-generation switch fabric module that will match the DCX's 256 Gbps slot throughput; existing customers will be able to upgrade by simply replacing the existing switch fabric module. "Replacing the switch fabric module costs an order of magnitude less than a forklift upgrade," noted Bill Marozas, business development manager, Cisco Data Center Solutions.

Despite each vendor's claim that its platform requires less SAN architecting, each director platform has idiosyncrasies a SAN designer needs to take into consideration to ensure optimal performance. In the case of the MDS 9500, the SAN design effort will likely be related to managing oversubscription and traffic prioritization. Correspondingly, the DCX

Backbone requires SAN architects to take latency variances between different ports within the same chassis into account, as well as its use of port- and exchange-based routing to load balance inter-ASIC links. While both Brocade and Cisco support port- and exchange-based routing over external ISL links, Brocade's use of these protocols inside the switch has been somewhat controversial. Customers need to make a choice between one of two routing modes; despite repudiation by Brocade, benchmarks like the December 2008 Miercom report (Report 081215B) have shown slower performance if the switch is used with port-based routing instead of the default exchange-based routing; and some array vendors advise their customers to stay away from the DCX's default exchange-based routing for some of their arrays.

"HP does not typically make specific recommendations regarding switch routing, but we recommend using port-based routing with the StorageWorks Continuous Access EVA solution since exchange-based routing doesn't guarantee in-order frame delivery all the time across exchanges," said Kyle Fitze, marketing director for the StorageWorks Storage Platforms Division at Hewlett-Packard (HP) Co. Conversely, EMC and NetApp confirmed that all of their arrays work flawlessly using the DCX default exchange-based routing mode.

FCoE AND CEE/DCE

Compelled by the prevalence of Ethernet and its enhancements, and the success and simplicity of iSCSI, Brocade and Cisco have embarked on bringing Ethernet into the well-guarded FC domain via Fibre Channel over Ethernet. FCoE is a T11 standard expected to be ratified later in the year. It uses Converged Enhanced Ethernet (CEE)—Cisco calls it Data Center Ethernet (DCE)—as the physical network transport to deliver Fibre Channel payloads. However, unlike its Ethernet brethren, it's lossless and appears as native Fibre Channel to the operating system and apps. Unlike iSCSI, it's not routable and is designed as a low-latency, high-performance Layer 2 data center protocol.

The value proposition of FCoE and converged Ethernet is lower infrastructure cost realized by simplifying cabling and reducing the number of adapters from two host bus adapters (HBAs) and two network interface cards (NICs) to two redundant Converged Network Adapters (CNAs) through which the converged local-area network (LAN) and FC traffic traverses. Instead of connecting two NICs to an Ethernet switch and two HBAs to a Fibre Channel switch, the two CNAs terminate into a CEE/DCE-capable switch that delivers Ethernet traffic to the LAN and FC traffic to the SAN. Although FCoE and CEE/DCE are expected to eventually be used from core to edge, its initial use is primarily at the access layer to connect servers to CEE/DCE-capable switches.

Both Brocade and Cisco are committed to FCoE, but with different strategies. Brocade won't ship Converged Enhanced Ethernet products until the standard is ratified; at that point, Brocade will support FCoE and

CEE in its DCX Backbone via new blades. Older Brocade Fibre Channel products, such as the 48000 Director, will connect through the DCX Backbone or a new top-of-rack switch to interface with CEE components.

With the Nexus 5000 Series top-of-rack switch, Cisco is the first vendor to offer a pre-standard FCoE product. For the MDS 9500 director family, as well as the Nexus 2000 Series Fabric Extenders and Nexus 7000 Series switches, DCE and FCoE support won't be available until standard ratification, similar to Brocade's plans.

Overall, Cisco has a more coherent product strategy to support a unified data center protocol and has been working on it for the past five years. With the Nexus 7000 core switch and its formidable 15 Tbps planned throughput designed to support 100 Gbps Ethernet, its unified NX-OS operating system that's used by both the MDS 9000 family and the new Nexus platform, and its Data Center 3.0 initiative to unify computing systems, Cisco is amazingly ready for the battle for the unified data center. On the other hand, with the recent acquisition of Foundry Networks Inc., Brocade isn't standing still and, despite integration challenges, a legion of loyal Brocade and Foundry customers are likely to side with Brocade. Even though the battle has begun, broad adoption of CEE isn't expected until late 2010. "With the exception of some early adopters, broader adoption of CEE won't happen until 18 to 24 months from now," said Greg Schulz, founder and senior analyst at StorageIO Group in Stillwater, Minn.

MAKING A CHOICE

With much in common, including pricing (see "[The high cost of high-end switches](#)," p. 22), the most important director selection criteria are which platform fits best into your existing infrastructure and which vendor's strategy aligns more closely with your roadmap.

"We chose Brocade directors mainly because we already had McData and Brocade departmental switches and our staff was familiar with managing them. The DCX performance was icing on the cake," said Michael Kilmer, primary storage administrator at Daktronics Inc., a leading manufacturer of large format display systems in Brookings, S.D.

"The fact that we are an all-Cisco shop as well as the VSAN feature were the main reasons we went with the MDS 9500 platform," explained Colleen Rhode, systems analyst at East Carolina University in Greenville, N.C.

For Keith E. Moore, director of technology services at Bellevue, Wash.-based Applied Discovery Inc., a division of LexisNexis, scalability and features were the decisive factors for selecting Brocade. "One of the

The most important director selection criteria: which platform fits best into your existing infrastructure and which vendor's strategy aligns more closely with your roadmap.

reasons Applied Discovery chose Brocade is that Brocade is known to handle heavy traffic while sustaining wire speed. Our Cisco FC network consistently had performance issues under heavy load,” Moore said.

David Turner, vice president of technical operations at MobiTV Inc., a leading provider of video to mobile devices in Emeryville, Calif., is evaluating the Nexus 5000 to complement the firm’s MDS 9509 directors to lower costs and take advantage of copper cabling. “I decided to move off Brocade to Cisco for the sole purpose of taking advantage of the modularity of the MDS switches. It’s much more cost-effective to get to the next version with the Cisco platform. Unlike Brocade, it doesn’t require expensive forklift upgrades,” Turner said.

Fernando Mejia, senior manager of IT infrastructure at the Independent Purchasing Cooperative (IPCoop) Inc., the purchasing arm of the Subway franchise in Miami, acquired a Cisco Nexus 7000 instead of a Catalyst 6500 because of its high performance, scalability and the ability to replace his stackable Brocade FC switches once FCoE becomes available.

MATURING DIRECTORS

Regardless of whose product you choose, both platforms will reliably power your SAN, which is confirmed by the myriad storage-area networks currently powered by Brocade and Cisco. Both vendors are embracing the converged Ethernet paradigm in their product roadmaps, but unless you’re willing to debug the initial CEE/DCE flaws as an early adopter, you’re well advised to wait for at least another year until the standard and products have matured. ☉

Jacob Gsoedl is a freelance writer and a corporate director for business systems. He can be reached at jgsoedl@yahoo.com.