

Managing the information that drives the enterprise

STORAGE

new directions for storage networking

Developments taking shape now will change the way your SAN infrastructure looks in a few years.

what's inside:

3

Data center building blocks

6

8Gig Fibre Channel

11

FCoE

15

InfiniBand longshot

Storage networking is getting ... **sexy**

By Rich Castagna

The “N” in SAN and NAS has long been the Rodney Dangerfield of storage shops. Networking gear has simply lacked the cachet of sophisticated and complex storage systems. After all, it’s just the cables and cards used to connect the truly important things, right?

More intelligence is moving from servers and storage arrays into the network, with sophisticated switches and special-purpose appliances providing a central point of administration and a logical place to add new storage management functionality.

There’s hardly a shop out there that doesn’t have a mix of storage protocols on the floor—Fibre Channel (FC) for primary applications, NAS for user shares and iSCSI for less critical apps. And each one comes with a separate network and set of tools to manage them, a situation that can easily unravel into a networking nightmare. But newer protocols, like Fibre Channel over Ethernet (FCoE) promise to converge those networks into a single, more manageable network without having to alter the storage environment. Converged network adapters (CNA) will replace host bus adapters (HBAs) and NICs with a single interface that can connect to both FC and Ethernet.

Few companies have yet to virtualize at least some of their servers. But the connective tissue between those two environments can also be virtualized in the form of “big backbone” concepts forwarded by switch vendors or with innovative devices that let you create virtual interfaces. So if you need a new HBA or NIC for a connection, or need to change one to the other, you can just do it in the software. The potential to create truly agile networks is within reach.

With these developments (and more to come) in storage networking, the network isn’t only finally getting some respect, it’s getting downright exciting. ☺

Rich Castagna (rcastagna@storagemagazine.com) is Editorial Director of the Storage Media Group.

**More
intelligence
is moving
from
servers and
storage
arrays into
the network.**



data center building blocks

By Dave Raffo

Storage networking vendors have laid out their vision of the data center of the near future. It consists of Fibre Channel (FC) and Ethernet working in perfect harmony, aided by new developments, such as Data Center Ethernet and FC over Ethernet (FCoE) feeding into massive backbone switches.

The main architects of the vision are Brocade Communications Systems Inc. and Cisco Systems Inc., but their visions differ about the specific building blocks and timing. Meanwhile, all the talk about FCoE, Data Center (or Enhanced) Ethernet and converged networks could make a storage manager's head spin.

"It's probably overkill now," says Ed Delgado, storage administrator at financial services organization RiskMetrics Group, of the choices next-generation data centers will bring.

Delgado is considering upgrading from Brocade fabric switches to a director because his firm's EMC Corp. SAN keeps growing. But his more immediate concern is whether to go to 4Gb/sec or 8Gb/sec FC instead of worrying about FCoE or converged networks.

"Most of our architecture is 2Gb/sec," he says. "Four gig came out already, and now 8Gb/sec is out. So, maybe let's just skip 4Gb/sec and jump right to 8Gb/sec."

FCoE probably won't become mainstream before 2010, but that hasn't stopped Brocade and Cisco from laying out their convergence visions.

From Brocade's perspective, the future is their DCX Backbone director. When two DCXes are trunked, you get an 896-port behemoth. The DCX supports Brocade's bread-and-butter FC (including 8Gb/sec), as well as iSCSI, Gigabit Ethernet, Ficon and FC over IP. Brocade also has 8Gb/sec FC switches and host bus adapters (HBAs), and in July acquired Ethernet switch vendor Foudry Networks for \$3 billion.

Cisco counters with the Nexus platform, a larger version of its Ethernet Catalyst family. The first of these switches, the Nexus 7000, holds 512 ports for 10Gb/sec Ethernet and 768 ports for 1Gb/sec Ethernet, but has no FC capabilities yet. Cisco added a smaller Nexus 5000 line, beginning with a 40-port switch that supports FC with an expansion pack. Cisco's FCoE technology comes from Nuova Systems, a startup it funded and then acquired in 2008. Cisco's 8Gb/sec MDS FC switches aren't due out until late 2008.

The two major ingredients of both products are waiting on standards to be worked out. Data Center Ethernet is a spec that's designed to avoid the dropped packets that prevent Ethernet from being used in most storage networks today. When FCoE is ready, Brocade's DCX and Cisco's Nexus switches will be able to integrate FC and Ethernet connectivity.

Brocade and Cisco have different expectations on the timing of FCoE and Data Center Ethernet. Cisco maintains there will be a good amount of FCoE adoption in 2009, while Brocade predicts FCoE won't show up in

The main architects of how your data center will look in the future are Brocade and Cisco, but their visions differ concerning the specific building blocks and the timing of new products and protocols.

servers until 2010 and won't become mainstream in storage before 2011.

But Jackie Ross, VP of business development at Cisco, is much more optimistic, predicting that Cisco customers will begin using FCoE in production on servers in late summer, with volume shipments coming in November. She predicts that SAN arrays with native FCoE support will appear by late 2009.

Analysts say the next-generation data center visions of Brocade and Cisco won't seem so different when FCoE arrives but, today, each centers on its traditional strength.

"Fundamentally, they're going to look similar," says Greg Schulz, senior analyst and founder at StorageIO Group, Stillwater, MN. "Cisco will have more of a network look and feel. Brocade will have more of a storage look and feel. Brocade has more Fibre Channel out there, but less Ethernet. Eventually, they'll square off with their functionality."

Bob Laliberte, an analyst at Milford, MA-based Enterprise Strategy Group (ESG), agrees that DCX and Nexus won't cross paths soon: "A lot of people are pitting them together, and they will be head-to-head, but Cisco will be favored in networking and Brocade in storage."

That means DCX Backbone and Nexus will be used by different teams in the interim. "The networking team will be interested in Nexus and storage teams will be interested in DCX Backbone," says Laliberte. "If the networking side is buying Nexus and the storage side is buying DCX, the interesting thing will be how well they work together. The intersection point is when FCoE comes out."

There will be other vendors' wares in play, as well. HBA vendors Emulex Corp. and QLogic Corp. are hurrying out 8Gb/sec devices, as well as cards that support FCoE.

Xsigo Systems (which uses InfiniBand to virtualize I/O), as well as Ethernet switch vendors Extreme Networks, Force10 Networks, Juniper Networks and Woven Systems, hope 10Gb/sec and Data Center Ethernet will make their products more attractive to users. But in the long run, the backbone for the data center will likely come down to spine provided by Brocade and Cisco.

That backbone will be costly. IBM sells the Brocade DCX Backbone for an entry-level price of \$390,000. A Nexus 7000 is expected to cost approximately \$200,000 for a typical data center configuration, and that's without FC switches. The Nexus 5020 is approximately \$900 a port. ☺

Dave Raffo (draffo@techtarget.com) is Senior News Director of the Storage Media Group.

8Gig Fibre Channel

Switches and host bus adapters are available, but 8Gb/sec arrays won't be released until next year.

By Alan Radding

UNLIKE THE PREVIOUS FIBRE CHANNEL (FC) specification transition from 2Gb/sec to 4Gb/sec in 2004, there may be a legitimate enterprise need for 8Gb/sec FC.

“There are some select applications that require 8Gb/sec FC right now, like high-definition video,” says Tim Lustig, solutions architect at QLogic Corp. However, “the key driver for 8Gb/sec is virtualization and server consolidation,” says Kyle Fitze, director of SAN marketing at Hewlett-Packard (HP) Co. Large companies are adopting virtualization to consolidate thousands of servers. Bob Gill, managing director, servers at TheInfoPro, a New York City-based research organization, says “95% of respondents [to our latest survey] state that virtualization is critical to achieving their business objectives.”

NOT LIKE 4Gb/SEC

When the transition from 2Gb/sec to 4Gb/sec FC began, the big question was what organizations would do with the extra performance. At that time, almost no organization, regardless of size, was saturating its 2Gb/sec pipes, said Tony Asaro, then a senior analyst at Enterprise Strategy Group (ESG), Milford, MA.

The transition promised to be transparent and painless, and there was no price premium attached to 4Gb/sec components. At the time, Framingham, MA-based IDC predicted that after the industry began 4Gb/sec FC general shipments in 2005, it would achieve 90% market penetration within two years.

Now, 8Gb/sec FC products are available. By virtualizing three, four or more servers within a single physical server, organizations will generate considerably more I/O. "In that case, you'll need more ports and more bandwidth," says ESG analyst Bob Laliberte.

Virtualization can benefit from bigger network pipes, just as it benefits from increased memory. However, "even with virtualization, you're not automatically going to saturate your pipes," says Greg Schulz, founder and senior analyst at StorageIO Group in Stillwater, MN. Most of the apps virtualized aren't high I/O database transaction apps that would benefit from the performance kick of 8Gb/sec FC.

"We tested four applications on one processor and we hit an HBA [host bus adapter] bottleneck at 50,000 IOPS," says QLogic's Lustig. With 8Gb/sec, the HBA bottleneck could be pushed off until 200,000 IOPS.

END TO END

At this point, 8Gb/sec FC is available for the server and switch components of the SAN. Here's the rub for the storage architect: "Behind the switch, you still have 4Gb/sec storage," notes Lustig. Even with HP, which has been actively cheerleading 8Gb/sec FC, the storage arrays remain at 4Gb/sec. "For complete 8Gb/sec throughput end to end, you'll need 8Gb/sec to the disk array," says the firm's Fitze.

FIBRE CHANNEL OVER ETHERNET ROADMAP

<i>Product</i>	<i>Throughput (MB/sec)</i>	<i>Line rate (Gbaud)</i>	<i>Year T11 spec completed</i>	<i>Market availability (year)</i>
1GbE	240	1.25	2008	Market demand
10GbE	2,400	10.52	2008	2009
40GbE	9,600	42.08	TBD	Market demand
100GbE	24,000	105.2	TBD	Market demand

SOURCE: FIBRE CHANNEL INDUSTRY ASSOCIATION

FIBRE CHANNEL ROADMAP				
<i>Product</i>	<i>Throughput (MB/sec)</i>	<i>Line rate (Gbaud)</i>	<i>Year T11 spec completed</i>	<i>Market availability (year)</i>
1Gb/sec	200	1.065	1996	1997
2Gb/sec	400	2.125	2000	2001
4Gb/sec	800	4.25	2003	2005
8Gb/sec	1,600	8.5	2006	2008
16Gb/sec	3,200	17	2009	2011
32Gb/sec	6,400	34	2012	Market demand
64Gb/sec	12,800	68	2016	Market demand
128Gb/sec	25,600	136	2020	Market demand

SOURCE: FIBRE CHANNEL INDUSTRY ASSOCIATION

There's usually a lag between when new technology components are introduced and when they're built into general-purpose storage arrays. HP expects to release an 8Gb/sec EVA array in 2009. Other vendors will probably follow in the same timeframe.

The industry will follow what it did with the previous FC speed transition; the 8Gb/sec components will automatically sense slower components and ratchet down the speed. As a result, the best an organization can do, end to end, today is 4Gb/sec—unless it wants to aggregate pairs of 4Gb/sec array ports.

NEW SAN OFFERINGS

Brocade Communications Systems Inc. has announced the availability of 8Gb/sec FC switches and HBAs. Beyond speed, Brocade is promising to build in advanced capabilities for QoS and support for virtual server mobility through dynamic state migration of link and port personality profiles.

IBM Corp. has announced three new 8Gb/sec switches: IBM System Storage SAN24B-4, SAN40B-4 and SAN80B-4. Pricing starts at \$4,420, and the switches have higher port counts (up to 24, 40 and 80 ports, respectively) and, according to IBM, consume less power than previous models. IBM is also releasing three new 8Gb/sec switch blades for the IBM TotalStorage SAN256B director. And Cisco Systems Inc. says it will offer a "transparent" upgrade to 8Gb/sec for its MDS 9500 Series directors in Q4.

InfiniBand update

- ▶ InfiniBand is an industry standard, channel-based, switched-fabric interconnect protocol for servers and storage. Today, it offers three levels of performance: 2.5Gb/sec, 10Gb/sec and 30Gb/sec for low-latency, high-aggregate throughput.
- ▶ InfiniBand is deployed primarily between servers in clusters for the purposes of high performance or failover. It's positioned to be complementary to Fibre Channel and 10GbE. The InfiniBand Trade Association envisions storage networks connecting into the edge of the InfiniBand fabric to communicate with the high-performance compute resources found there.
- ▶ "InfiniBand is mainly for high-performance computing. It has limited opportunities in enterprise storage," says Mike Karp, senior analyst at Enterprise Management Associates. "How many protocols do you really want to support in your data center?"

PRICE PREMIUM

The move to 8Gb/sec FC will entail a price hike, at least initially. "The 8Gb/sec components require new optics. The price won't drop to 4Gb/sec levels until the volumes ramp up," says Scott McIntyre, VP of software and customer marketing at Emulex Corp. Adds Mario Blandini, Brocade's director of product marketing: "There will be a 20% to 30% price premium over 4Gb/sec components initially."

For the extra money, 8Gb/sec components will bring some new capabilities. "These involve how data corruption is handled and how you authenticate the host to the FC fabric," says McIntyre. For example, there will be more Cyclic Redundancy Checks (usually a mathematical checksum), which detect data alteration during transmission or when stored by comparing the data stream going in and coming out.

One upshot of the transition to 8Gb/sec FC may be improved energy conservation. "You may be able to turn off lanes on the PCI bus and power off lanes in the HBA and ports," says QLogic's Lustig. These green capabilities, however, will come from new firmware and software, not from the 8Gb/sec HBA alone.

FCoE

Although 8Gb/sec FC doubles FC performance, it still comes up shy of 10Gigabit Ethernet (10GbE). However, FC enterprises can achieve 10GbE speed through Fibre Channel over Ethernet (FCoE). With FCoE, the FC protocol and FICON run native on a premium lossless, low-latency and QoS-enabled Ethernet. FC becomes just another network protocol running at 10Gb/sec alongside IP on the 10GbE network.

FCoE also allows convergence at the cabling level through a single Ethernet cable handling both FC and IP. FCoE will appeal to enterprises with FC SANs but with no plans to migrate to iSCSI and Ethernet. "It lets organizations with FC and FICON preserve their FC skills and tools," explains StorageIO Group's Schulz. They can converge their network to a fast 10GbE backbone and reap savings by running and maintaining only one set of cabling for all traffic.

Because FCoE doesn't use TCP/IP, it isn't routable. In addition, at this point FCoE is considered local technology only. "FC shops that want to follow strategies requiring long distance, such as remote mirroring, should look to FCIP [Fibre Channel over IP] or other protocols," says Schulz.

To use FCoE, organizations will have to deploy a Converged Network Adapter (CNA) that will look to the server as both an FC HBA and an Ethernet NIC. A number of vendors have declared support for FCoE, including switch makers like Brocade and Cisco, and component makers like Emulex and QLogic. Despite industry support, "this technology is still in its infancy," adds Schulz. "It will be 2010 before FCoE is ready for use by anyone except early adopters."

BEYOND 8Gb/SEC FC

The Fibre Channel Industry Association, which manages the FC roadmap, is signaling the next performance jump to 16Gb/sec. Work on the 16Gb/sec FC specification is underway with completion slated for 2009. Actual products should be delivered by 2011. After that, the future gets murky with development of 32Gb/sec, 64Gb/sec and even 128Gb/sec FC dependent on market interest.

If an organization needs the performance of 8Gb/sec FC, it can begin upgrading to available switches and HBAs, or it can ponder a future jump to FCoE. When the time comes to finally make the move to end-to-end 8Gb/sec FC, prices should have dropped. Otherwise, sit tight. ☉

Alan Radding is a frequent contributor to *Storage* magazine.

FCoE

Fibre Channel over Ethernet is speeding along the certification path, and now is the time to determine what it can do for you.

By Bob Laliberte

Y NOW, YOU'VE PROBABLY heard the hype surrounding Fibre Channel over Ethernet (FCoE). Why should you care? Because FCoE has the potential to reduce data center complexity and make the world a little greener by decreasing the number of cards, cabling and network devices in the data center. In some large organizations, the ability to reduce cable bundles could have a positive impact on air flow and reduce cooling costs.

The Ethernet part of the protocol isn't just any Ethernet, but a special, still-to-be-ratified Data Center Ethernet (DCE). To make Ethernet suitable for Fibre Channel (FC) transport, the 802.1Q IEEE standard is being modified to accommodate data center traffic to improve its priority traffic flow and allow it to operate in a lossless manner (no dropped packets).

The goal is to deliver FC via a different network protocol leveraging an Ethernet fabric, while maintaining the same or better performance that FC-only networks have enjoyed. Based on the IEEE modifications currently being considered, our best guess puts the timeframe for the protocol's delivery around Q4 2008. The protocol must be ratified by the FC standards body and the INCITS T11 Technical Committee.

IS FCoE NECESSARY?

Now that we understand FCoE, the real question is whether or not we need it. Why not jump to iSCSI? While iSCSI utilizes TCP layered on top of IP, neither FCoE nor iSCSI is routable. InfiniBand also looked promising in this area, but has so far only found a home in high-performance computing environments and hasn't seen widespread adoption otherwise.

So why do we want to converge FC and Ethernet into a single fabric? By combining FC and Ethernet, a single cable and a single card can replace current network interface and host bus adapter (HBA) cards. And because every new technology needs a three-letter acronym, the resulting interface card will be known as a Converged Network Adapter (CNA). These will feed switches or directors, and accommodate FC or FCoE and Ethernet over 10Gig

links. In addition, the same switches and directors will handle both protocols while simultaneously allowing storage and networking domains to control their traffic independently. This preserves the separation of the storage and networking management domains while consolidating hardware.

Who's involved in this effort?

From a storage perspective, we should look at the usual suspects: Brocade, Cisco Systems Inc. (Nuova Systems), EMC Corp., Emulex, Finisar, Hewlett-Packard Co., IBM Corp., NetApp, QLogic and Sun Microsystems Inc. Other players interested in participating in the FCoE market include Blade Network Technologies, Broadcom, Intel and Mellanox. The most obvious and most vocal proponents of this technology have been Cisco and Nuova (the former Andiamo team), which Cisco funded and then acquired outright.

Beneath the marketing gloss, FCoE looks very promising. It's certainly moving down the certification route faster than any other protocol in recent history, on track to go from inception to production in less than two years. This has been accomplished through universal backing by all participants, as well as by Nuova relinquishing its patents to FCoE with the understanding that proprietary protocols don't create a big market. Does all of this unprecedented cooperation guarantee success? Probably not, but it should help to relieve any interoperability concerns.

Most new technologies aren't implemented in production environ-

FCoE provides an elegant way to migrate Fibre Channel to Ethernet, while protecting existing Fibre Channel investments and skill sets.

ments until they're fully tested and proven. For users, the key to FCoE will be understanding its impact on data center operations. Once test and development cycles have been completed, early deployments will most likely be in the form of server fan-in environments—especially high-density blade server environments—and then move into the core from there.

ISSUES TO CONSIDER

But there are some FCoE issues that go beyond technology and certification. Most of them are related to culture and domain segmentation. While we do see some progressive IT shops working to integrate technology silos, convergence is still a challenge in many places. Here are some critical questions you'll want answered.

- ▶ Will the storage team or the networking team own the infrastructure? If co-managed, who has the deciding vote?
- ▶ Which department will pay for it? How will chargeback be calculated and future growth determined?
- ▶ Will the teams be integrated? Typically, the networking team is responsible for IP switches, while the storage team is responsible for FC.
- ▶ Who will own day-to-day operational issues? If a decision needs to be made regarding whether more bandwidth is given to LAN or SAN traffic, who makes the call? Will companies have to create a single, integrated connectivity group?

Aside from potential cultural issues that must be faced, another question is whether FCoE is compelling enough to merit a forklift upgrade. At this time, I don't believe we'll see FCoE conversion projects. Rather, FCoE will probably be implemented as part of a bigger IT project, like server virtualization or a storage array technology refresh. Because some new hardware (read: capital expense) will be required, justifying the investment may be difficult, especially considering the macro environment. However, if the math becomes compelling enough, say below \$500 per port, adoption may accelerate.

So why should you be thinking about FCoE now? Products are becoming available and some vendors have made claims that FCoE will be in production environments before the end of this year. Users currently testing FCoE environments using alpha/beta equipment seem quite satisfied with it. With some products available now and more due this fall, I say companies should consider creating a test and de-

velopment platform for FCoE to become comfortable with the technology. Companies should definitely be planning to include FCoE in their 2009 budget if it isn't already in this year's.

THE FUTURE OF FCoE


FCoE appears to have all the technology bases covered. It enables companies to retain an existing FC infrastructure, keeps existing FC management tools in place, provides the same level of performance (with DCE) and lowers costs. The biggest variable in all this is the economics of FCoE solutions. Given the present economy and the pressure IT is under to reduce costs, compelling FCoE pricing may accelerate adoption faster than any marketing pitch or certification.

What will be the role of FC going forward? Saying that FCoE will be the end of FC makes for good headlines, but the reality is that FC is here to stay, at least for a while. The Fibre Channel Industry Association continues to drive toward 16Gig FC. FC will co-exist with FCoE for a number of reasons. FCoE still needs to be tested and proven, and FC will continue to deliver services until then. Given that data centers are very slow to change—with plenty of people taking the “If it ain't broke, don't fix it” position—adoption could take a while. Lastly, there will always be protocol zealots resisting any new technology. It should be noted, however, that FCoE was designed to co-exist with FC, and these technologies will work with existing management tools.

The green impact of FCoE is a compelling factor. The combination of lower power and cooling requirements (not from the protocol, but from the reduction in equipment) with reduced cabling will certainly be attractive as greater emphasis is placed on green initiatives.

You could also consider FCoE the beginning of the end for the cultural barriers that exist between technology domains in large-scale data centers. At the very least, it provides the storage team with an opportunity to get to know the networking team so they can work together to provide higher levels of service to the business and reduce costs. Ultimately, it could provide an opportunity to embrace the sort of change that delivers significant benefit to the business. ◉

Bob Laliberte is an analyst at Enterprise Strategy Group, Milford, MA.



InfiniBand longshot

*Converged networks may
keep InfiniBand on the
storage sidelines.*

By Dave Raffo

IBRE CHANNEL (FC) STORAGE VENDORS see a need to play nice with Ethernet. That's why their vision of the next-generation data center includes a nascent FC over Ethernet (FCoE) standard as the key piece of what they consider a unified fabric.

Such unification leaves little room for InfiniBand. The low-latency, high-bandwidth protocol isn't used much for storage, but has replaced Ethernet in many high-performance computing (HPC) environments. Proponents say InfiniBand is a better fit for FC convergence than Ethernet.

FC-centric storage networking vendors such as Brocade, Emulex Corp. and QLogic Corp. consider Ethernet compatibility the key to extending FC's life. So does Cisco Systems Inc., which is primarily an Ethernet company, but competes with Brocade in the FC switching area.

But InfiniBand vendors are working up their new standard, FC over InfiniBand (FCoIB). They point out that InfiniBand already supports 20Gb/sec with 40Gb/sec expected by next year, while Ethernet is at 10Gb/sec. And Ethernet won't be ready to run high-performance storage apps until Enhanced Ethernet is adopted to prevent dropped packets.

With FCoIB, InfiniBand adapters will encapsulate FC packets in hardware and move them over InfiniBand. Thad Omura, VP of product marketing at Mellanox Technologies, says InfiniBand adapters support the standard, but software development won't be complete until late this year or early 2009.

But FCoIB doesn't have backing like FCoE. Cisco and QLogic sell InfiniBand switches and adapters, but Mellanox and Sun are the main voices behind FCoIB. "There hasn't been a lot of noise about it," admits Omura.

Arun Taneja, founder and consulting analyst at Hopkinton, MA-based Taneja Group, predicts InfiniBand's low latency and high bandwidth will give it a seat at the table, but as a separate fabric. He says there's no reason to use it in a network with Ethernet.

"If I take a 40gig-per-second protocol and put that on 10gigE converged Ethernet, I should have my head examined," says Taneja. "If I need 40gigs per second, I'm going to stay with InfiniBand, and more than likely I'm looking for InfiniBand at the front end." Only LSI and DataDirect Networks have that now.

QLogic expects InfiniBand to remain mostly an HPC fabric. Sash Sunkara, marketing VP for QLogic's network solutions group, thinks InfiniBand will play a large role in server I/O virtualization but not in storage.

"Fibre Channel is the SAN fabric today," says Sunkara. "It could move to FCoE, but I don't see InfiniBand in that [SAN] space any time soon." ☉

"If I take a 40gig-per-second protocol and put that on 10gigE converged Ethernet, I should have my head examined."

Dave Raffo (draffo@techtarget.com) is Senior News Director of the Storage Media Group.