

Managing the information that drives the enterprise

STORAGE

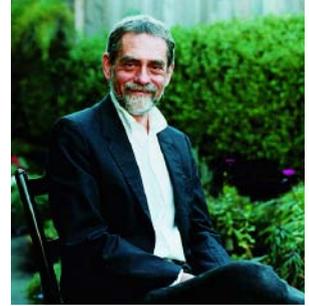
ESSENTIAL GUIDE TO

Storage Networking

Storage networking experts and technologists spotlight the advancements that will have the most significant impact in storage networking to help you prepare for the future.

INSIDE

- 2 New era of networking
- 4 The SAN of the future
- 13 Storage networking outlook
- 18 Timetable for 10 GbE
- 22 Fibre Channel switch best practices
- 28 Optimizing iSCSI SANs
- 31 Vendor resources



New era of networking

HERE WAS A TIME when storage networks were little more than after-thoughts. You started at the array and worked back toward the servers; if you needed high performance, you got a Fibre Channel (FC) array or if you were serving files, NAS was the choice. So the type of storage dictated the type of network you'd use to string it all together. And if that connective tissue you put in place didn't quite do the job, you got more of it or speedier parts.

Times have changed—in fact, they're on the verge of being turned inside out. The “inside” is the network, and pretty soon it's likely that it will be starting point when designing shared storage systems. This might seem like a revolutionary change, but it's really more a case of evolution. Today, a good-sized shop may have several FC fabrics, a few IP networks dedicated to NAS filers and, of course, a LAN to link hosts and clients to all that storage. It's a lot to manage, and managing it often crosses the boundaries of two or more IT disciplines.

Wouldn't a single network make more sense? It not only makes more sense, it's here now. Fibre Channel over Ethernet (FCoE) makes it possible to link these previously disparate networks, promising simpler administration, less complexity and lower costs. Any one of those would be irresistible to most IT shops.

Of course the success of FCoE hinges on a number of factors, including widely available—and affordable—10 GigE components, and the finishing touches on the new Ethernet standard that will make FCoE possible, Data Center Ethernet (DCE)—or Converged Enhanced Ethernet (CEE), as it's also known. Still, FCoE switches are here today, along with Converged Network Adapters (CNAs) that combine the functions of common NICs and HBAs.

But this doesn't suggest that you'll necessarily want to link all of your networks. In some cases the costs may outweigh the conveniences, or it simply may not be necessary. Still, it's prudent to take a fresh look at your installed networks, with so many protocol upgrades and new products available now or coming soon. ☉

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

Fibre Channel over Ethernet (FCoE) makes it possible to link previously disparate networks, promising simpler administration, less complexity and lower costs.

Nothing like a **Happy** IT Management Team



Unveiling Four Outstanding **NEW** products for:



No Nonsense Backups



High Availability



Seamless Migrations



Rapid Provisioning

Double-TakeBackup

Double-TakeMove

Double-TakeAvailability

Double-TakeFlex

SAN OF THE FUTURE

Storage technologists and industry analysts predict how the SAN will evolve and what you need to do to prepare for the changes.

By Alan Radding

IN FIVE YEARS, the enterprise SAN might be a service running in the cloud. Or a huge collection of DAS, like a giant mainframe DASD farm of old. It might be cableless, contained in a pre-wired cabinet or wireless. Object-based storage could make blocks and files irrelevant. The SAN might even be diskless if solid-state disk (SSD) economies of scale and adoption accelerate in a hockey stick curve. Whatever form it takes, the SAN of the future will be more consolidated, virtualized, automated and greener.

Or, as many predict, the changes will be evolutionary and not revolutionary; in five years, the SAN will be a lot like the enterprise SAN of today: just faster, packed with more disk capacity, cheaper on a cost/GB basis, a little easier to manage and less energy hungry.

Storage magazine recently asked storage vendors, industry analysts and technologists serving on storage industry associations about where they see the SAN heading. There may not be sweeping architectural changes in five years, but there will be changes in the basic building blocks of the SAN infrastructure: networks and protocols; switches; storage arrays, disks and controllers; and SAN management.

NETWORKS AND PROTOCOLS

Today, only about half of the storage deployed is networked, said Jackie Ross, VP, business development at Cisco Systems Inc. In five years, the amount of networked storage will increase to 70%, she suggested.

Among networked storage, Fibre Channel (FC) is the dominant storage networking protocol in the enterprise data center with more than an 80% market share, according to Skip Jones, chairman of the Fibre Channel Industry Association. Roger Cox, a research VP at Stamford, CT-based Gartner Inc., projects a 66% share for FC by 2012.

By then, 8 Gb/sec FC will be heading toward 16 Gb/sec, while 10 Gb/sec Ethernet will be aiming for 40 Gb/sec or even 100Gb/sec, keeping in line with the Ethernet's full order of magnitude increases. At that point, FC will face being left behind in terms of sheer network performance.

But before then, the game will shift. "We see the industry moving to a unified fabric," said Ross. That means combining FC and iSCSI on Ethernet. "The construct for FC storage won't change. You manage the SAN, provision LUNs and do masking the same way," she explained. What will change is the number of components the organization needs. There will be only one type of switch and one type of adapter. "Cabling, which represents 25% to 30% of the data center cost, is reduced, too," said Ross.

"In five years, the network infrastructure will have to be a unified platform that speaks multiple protocols," said Jason Schaffer, director of storage product management at Sun Microsystems Inc. "It will spit out whatever protocol the server or storage dictates." But Cox warned that "there are a number of issues that will keep convergence from happening." The biggest ones are organizational. "You have issues between networking and storage people that aren't easy to resolve," he said.

On the technical side, FC over Ethernet (FCoE), for example, isn't a slam dunk. For 10Gb Ethernet to provide the basis of FCoE, "you need a special form of 10Gb Ethernet," said Cox. Called Converged Enhanced Ethernet, (CEE) it will address such things as flow control, which is necessary to deliver the lossless networking that makes FC storage so popular. "The standards aren't yet in place," said Cox, who doubts they'll be ready for widespread deployment in five years, noting that "FCoE will achieve about 2% market penetration by 2012."

As for InfiniBand, forget it. "Maybe we'll see InfiniBand as an alternative for the converged network," said Greg Schulz, senior analyst at StorageIO Group, Stillwater, MN.

SWITCHES

Switches will be more flexible and intelligent. "By then, plumbing will be less important than intelligence," said Jon Toigo, CEO at Toigo Partners International, Dunedin, FL.

Cisco expects switches to be capable of providing networking services,

such as firewalls, load balancing and other quality of service (QoS) functionality. The switch will also play a central role in network management automation. “To get [end-to-end] automation, you’ll need intelligence at multiple places: in the converged network, adapters, HBAs, array controllers,” said Ross (see “[Where to put storage intelligence](#),” below).

No protocols will go away anytime soon. Instead, switches will handle multiple protocols, including FC, FCoE, Ethernet, enhanced or converged Ethernet, iSCSI and possibly InfiniBand. By 2013, multiprotocol SAN

WHERE TO PUT STORAGE INTELLIGENCE

LOCATION	ADVANTAGES	DISADVANTAGE	PROBABILITY
Server/Host bus adapter (source)	<ul style="list-style-type: none"> • Direct access to CPU for faster performance • Access to policies • Network and storage system agnostic 	<ul style="list-style-type: none"> • Diverts CPU resources from other tasks • Need to manage multiple adapters across servers 	60%; host bus adapter vendors will be adding intelligence
Application (source)	<ul style="list-style-type: none"> • Knows data requirements • Access to policies • Storage and network agnostic • More awareness of what’s happening with the data 	<ul style="list-style-type: none"> • Diverts CPU resources from other tasks • Slows performance • Licensing costs 	75%; will need software on the server to help it
Switch (network)	<ul style="list-style-type: none"> • Offloads processing from server • CPU, server and application agnostic • No need to manage software on server 	<ul style="list-style-type: none"> • Distance from application and storage increases latency • Potential for lock-in at the network • Possible bottleneck (if everything flows through switch) 	85%; switch vendors are there, but whose software will run the switch?
Network appliance (network)	<ul style="list-style-type: none"> • Offloads processing from server • CPU, server and application agnostic • No need to manage software on server 	<ul style="list-style-type: none"> • Distance from application and storage increases latency • Another device to manage • Potential for lock-in at the network 	95%; already being done extensively (IBM, NetApp); high user acceptance
Controller/Array (target)	<ul style="list-style-type: none"> • Offloads processing from server • Puts management adjacent to the storage to reduce latency 	<ul style="list-style-type: none"> • Vendor lock-in 	100%; shipping now

Source: Greg Schulz, senior analyst, StorageIO Group

switches should be commonplace, although the particular combination of protocols may vary. Switches will also be bigger, encompassing hundreds of ports and enabling thousands of ports on the network. Intelligence will reside in the core switches, and edge switches will connect to the core.

STORAGE ARRAYS, DISKS AND CONTROLLERS

Storage arrays will continue riding Moore's Law. 10Gb/sec Ethernet and 8Gb/sec FC will be standard interfaces for enterprise arrays. "You'll see the expected increases in performance and capacity from all of the major vendors," said Kyle Fitze, director of marketing in the SAN Division of Hewlett-Packard (HP) Co.'s StorageWorks group.

Storage arrays will continue to consist primarily of hard disk drives (HDDs) in 2013, although the size and form factor may vary. "In five years, most of the storage will be ultra-high-density arrays packing large numbers of drives into small footprints," said Schulz. These arrays will become the norm, not just something for firms facing energy or space constraints.

One technology that's not likely to replace HDD in the array is SSD or flash drives. Vendors currently incorporate SSD in arrays and will continue to do so, but SSD will be reserved for critical applications requiring very high IOPS. HP distinguished technologist Jieming Zhu said two main issues deter rapid adoption of SSD: price and SSD's inherent wear-out factor. Zhu added that work needs to be done on software that prolongs the life of SSDs and better integrates them with RAID and database applications. "It's a work in progress," he said.

HDD capacity will keep getting bigger, delivering

THE BENEFITS OF object storage

SAN storage in five years will be increasingly object based. Object-based storage resembles file-based storage except it makes greater use of metadata. But object-based storage isn't a total win-win proposition. It trades the efficiency and performance of block-based storage for easier management and more automation.

Object metadata will let you manage the storage more effectively and apply policies based on the data content, regulatory requirements, ownership of the data and so on. The metadata can also be used to dynamically store data at the most appropriate service levels.

price/performance increases of approximately 40% a year. Low-cost 1.5 TB SATA drives will be surpassed by even larger disk drives of 4 TB or more. For organizations needing performance greater than 15K rpm, "there's no reason why there can't be 20K or even 22K drives," said Ed Grochowski, conference committee chairman of the International Disk Drive Equipment and Materials Association (IDEMA).

What you're more likely to see are drives supporting 4K (4,096 bytes) sectors for error correction. This is a completed IDEMA standard and compatibility testing is under way this year. By 2013, the 4K sector will be in all new SATA drives (iSCSI drives aren't impacted by sector size) and possibly adopted by the SSD industry.

You should also expect to see more file-oriented, NAS-like storage in

FUTURE DIRECTIONS: SERVER VIRTUALIZATION

TODAY, THE BASIC CHALLENGES for storage posed by server virtualization are being addressed. VMware Consolidated Backup (VCB) offloaded the processing cycles required for backup to a centralized backup proxy. And, in its vSphere 4 release, VMware introduced vStorage APIs for data protection that enable third-party backup tools to integrate more easily with VCB. "In the future, administrators will be able to just click a box in the backup tool for the kind of backup and restore they want," when using tools that are integrated with the APIs, said Jon Bock, group product marketing manager at VMware. Administrators will have a choice of VM-level or file-level restores from a single backup pass.

It has been difficult to provision storage for moveable virtual machines (VMs). VMware's Virtual Machine File System (VMFS) abstracts details of the underlying physical storage and limits the number of times storage administrators have to provision storage for VMs. "The APIs in vSphere 4 enable storage and management vendors to see what storage is being used by VMs. Future APIs will enable even deeper visibility and integration," said Bock. This will enable storage administrators, for example, to see and resolve LUN bottlenecks resulting from unexpectedly heavy VM activity.

VMware vSphere 4, which was announced in April, provides features and interfaces to storage technology to support a range of storage activities, according to Bock. These include thin-storage provisioning of the virtual disks used by VMs, features for online expansion of VMFS volumes and virtual disks and APIs that enable third-party multipathing software to be used with vSphere environments. VMware vSphere 4 also has an integrated backup tool, called VMware Data Recovery, that uses deduplication to deliver VM backup for small environments.

the data center. “This will simplify provisioning; it’s not nearly as complex to manage as block-based storage,” said StorageIO Group’s Schulz. He expects file-oriented storage to be widely accepted even for database applications.

At about the same time, the data center will begin to see the early implementations of object-based storage, noted Schulz. Object-based storage contains richer metadata than block storage. “It becomes a question of which is the better level of abstraction: the richness of the object-based system or the efficiency of block storage,” said Rick Gillett, VP of data systems architecture at F5 Networks Inc. (see [“The benefits of object storage,” p. 7](#)). By relying on in-depth metadata, object-based systems will know more about the data and enable intelligence in the storage system to better manage the data.

Faced with surging volumes of data, more intelligence will be needed in storage systems. Where that intelligence should reside is an open question. “The SAN is taking over much of the intelligence that used to be in the server,” said ReiJane Huai, chairman and CEO at FalconStor Software. SAN-based intelligence already provides services like snapshots and replication.

By 2013, storage controllers will have sufficient processing power to run, for example, database apps. “Just think about running Oracle on a controller right next to the storage array. Just imagine what that could do for database performance,” noted Huai.

STORAGE MANAGEMENT

Storage management will get harder before it starts to get easier. Storage virtualization embedded into the SAN can simplify some aspects of storage management while server virtualization complicates it.

Server virtualization will continue to complicate storage management. “This is a new dimension for storage management,” said Joseph Zhou, senior analyst, storage research at Ideas International Inc., Rye Brook, NY. Virtualization requires dynamic reprovisioning to accommodate changes to virtual servers. In five years, dynamic reprovisioning should be supported for leading hypervisors (see [“Future directions: Server virtualization,” p. 8](#)).

Convergence of protocols over a unified fabric promises simplified management. “You will be able to manage across FC and iSCSI,” said Mike Karp, senior analyst at Enterprise Management Associates, Boulder, CO. Unresolved is who will manage the FCoE network: network admins or storage admins.

Faced with surging volumes of data, more intelligence will be needed in storage systems. Where intelligence should reside is an open question.

“Intelligent storage is the management solution,” said Steve Luning, VP, office of the CTO at Dell Inc. Storage intelligence could reside in the app, server, data, array, off in the cloud or some middle layer. “Maybe the hypervisor handles the management,” suggested Luning.

But some storage management tasks aren’t practical to automate. “You can automate the most common tasks, like backup, but these aren’t what cause problems,” said StorageIO Group’s Schulz. Problems caused by increased complexity and products that comply with standards at a high level but break the standard deeper down will continue to make storage difficult to manage.

“Where vendors provide management tools, they’re all stovepiped. Cisco or EMC can add management capabilities, but most often they only work in their environments. As soon as you go beyond the vendor, you lose the management benefits,” noted Schulz, adding that “this is unlikely to change.”

What’s needed is a common storage management platform that’s transparent from top to bottom. SMI-S doesn’t do the trick, according to Toigo, at Toigo Partners International. Instead, he envisions the SAN as a set of managed Web services.

STORAGE SKILLS

“Storage managers will have to get comfortable with server virtualization and moveable workloads,” said Dell’s Luning. “They’ll also need to know about the data, data classification, and better understand each app’s storage and performance requirements.”

The skills storage admins have today—setting up RAID, provisioning LUNs, zoning and masking—will be relegated to a few specialists or automation.

“The low-level skills will get folded into automation,” said Sun’s Schaffer. “The storage administrator’s expertise will lie in knowing what the data needs and what the requirements are.”

For example, a storage admin setting up storage for Microsoft Exchange “will need to know not only the number of mailboxes and their size, but the performance needs and protection requirements, the RPO and RTO,” said HP’s Fitze. Ideally, the admin can specify this at a high level and automation will set it up correctly.

In addition, storage admins may have to rethink their approach to RAID for extremely large (1 TB-plus) disk drives due to impossibly slow rebuild times.

“Extremely large drives raise questions about RAID. Administrators may have to do RAID across files or objects so they would have to rebuild only a small part of a disk,” said Enterprise Management Associates’ Karp.

“You can automate the most common tasks, like backup, but these aren't what cause problems.”

—Greg Schulz,
senior analyst, StorageIO Group

SAN TRENDS, 2013

DISK DRIVES

Likelihood: 90% Hard disk drives remain the dominant storage in 2.5-inch and 3.5-inch form factors; 4K sectors will emerge for enhanced error correction; expect capacities to reach 4TB disk capacity, but 15K rpm will remain the top choice for performance.

MANAGEMENT

Likelihood: 60% With the widespread adoption of VMware and other hypervisors, APIs by management tool vendors will simplify the backup of virtual servers and enable dynamic provisioning of mobile virtual machines. Intelligence embedded in the SAN and switch will enable more automated, policy-based data management. Object-based storage with rich metadata will allow more intelligence-driven data automation.

STORAGE ARRAYS

Likelihood: 85% Ultra-high-density storage arrays will pack more storage into a smaller, greener footprint; arrays will have multiple interfaces (IP, enhanced IP, Fibre Channel over Ethernet, Fibre Channel) to connect with converged fabrics; some solid-state disk will be incorporated for high IOPS data.

SWITCHES

Likelihood: 70% Multiprotocol switches will be common, and switches will have greater intelligence, which will be used for management.

NO CLOUD IN THE FORECAST

What the enterprise SAN won't look like in five years is a SAN in the cloud, although some storage operations may use the cloud. Similarly, the SAN is unlikely to exist as a set of Web services despite the widespread acceptance of Web services. A wireless SAN could eliminate cabling hassles and expenses, but the volume of data and security concerns make this unlikely. Large DAS farms are a possibility for special situations, but they're unlikely to replace the enterprise SAN despite the simplicity of DAS.

The SAN in five years may look surprisingly similar to the enterprise SAN of today. Protocol convergence, unified fabrics and server virtualization will simplify and complicate the SAN. Storage administrators will need new skills—a better understanding of virtualization, data and apps—while keeping their traditional storage skills sharp. It's not that SAN technology isn't advancing fast. Rather, organizations deploying enterprise SANs adopt change at a more measured pace. ☉

Alan Radding is a frequent contributor to *Storage*.

Storage networking outlook

Storage pros will need to learn more about the network than ever before.

By Bob Laliberte

ATA CENTERS are being transformed. Companies are consolidating geographically dispersed data centers into centralized ones to reduce footprints and costs, and to improve performance. One of the most visible technologies enabling this change is virtualization, particularly server virtualization. But despite all of the attention virtualization has received, probably less than 20% of available servers have been virtualized, leaving a lot of room for future growth.

Another significant part of this transformation is the expanding role of the network. To support all of the features and functionality of server virtualization, a networked storage environment is required. Research from Enterprise Strategy Group indicates that more than 85% of server virtualization shops leverage a networked storage environment. While vendors will argue the merits of various types of networks, the most common one is still Fibre Channel (FC), chosen for performance reasons. However, it's not used exclusively and many firms will deploy multiple storage networks based on performance needs, internal skills and budgets.

FCoE'S ROLE IN THE NETWORK

Just as data centers are transforming, the most popular storage networking technology is also evolving. While many companies were content to follow the FC roadmap—upgrading from 1 Gb to 2 Gb, then to 4 Gb and now 8 Gb—new technologies like Fibre Channel over Ethernet (FCoE) have given users something to think about before blindly progressing to 16 Gb FC. Why is that? And why should the storage team pay attention?

- ***There's the potential for much higher throughput/performance.***
 - FCoE leverages 10 Gb Ethernet (10 GbE). To be more specific, it leverages an enhanced version of the Ethernet standard referred to as Converged Enhanced Ethernet (CEE). The changes are mostly related to eliminating dropped packets and relieving congestion.
 - The roadmap for FCoE mirrors Ethernet. This means the next leap is four times the throughput (up to 40 Gb), which will quickly surpass the FC roadmap.
- ***Additional savings can be realized through convergence.***
 - Every IT organization is under constant pressure to reduce costs. FCoE provides the opportunity to reduce the number of cards and cables required at least at the rack level. This could also have an impact on power and cooling requirements.
 - List prices for 10 GbE ports are already less than \$500 per port and will continue to decline as sales volumes increase.
- ***Major vendors have made significant investments in this space.***
 - They understand the benefits of convergence and are building hardware and software portfolios to provide solutions to enable this transition. Some of the more notable acquisitions include Cisco Systems Inc. bringing in Nuova Systems Inc. and Brocade's acquisition of Foundry Networks Inc. Other firms like Emulex Corp. and QLogic Corp. have developed their own technology to deliver converged network adapters to replace host bus adapters and NIC cards.

Why is this important to the storage team? As data centers and the networks that power them continue to change, the line between data networks and storage networks will blur. Server virtualization and data mobility are forcing IT to rethink the traditional, siloed approach to data center technologies. For example, before Cisco announced its Nexus 1000 virtual switches at VMworld 2008, server admins controlled VMware virtual switches embedded in the ESX hypervisor through a VMware interface. Now, if users choose to deploy the Cisco Nexus 1000 in VMware environments, network admins can regain control of the switching environ-

ment and use Cisco's NX-OS to manage the virtual, as well as the physical, Ethernet switches.

WHAT THE TRANSFORMATION MEANS TO STORAGE TEAMS

It's important to understand where your company is when considering this transformation process. Has it implemented server virtualization? Is it in production? How is it connecting the virtualized server environment to the storage? What technologies are being considered? Take this opportunity to become more relevant to the business. Think in terms of how changes in the IT environment can positively impact the company's bottom line, not just enhance the visibility of your particular domain. More specifically, the following must be considered:

BOTTLENECKS. Once server virtualization technologies have been deployed correctly, bringing on a new application can require only a few minutes. But how long will it take to provision the storage to support it? If the answer is measured in days or weeks, the process needs to be reviewed and new storage technologies may be required.

Who controls the newly deployed network? Typically, storage teams dictate the type of network supporting the storage environment. Looking ahead, it's easy to imagine that changing. If FCoE takes off, will the deployment of Ethernet switches be controlled by storage or networking companies? Will the storage team or the networking group have the responsibility and budget? Will companies need hybrid IT groups with members from both of these teams?

Do FCoE products need certification from storage vendors? Traditionally, FC switch sales are influenced by the specific vendor but controlled by storage companies. Ethernet switches, sold into the network groups, haven't needed any approvals for NAS and iSCSI implementations. With FCoE, however, the game has changed. For now, all FCoE products are undergoing testing and so-called certification by major vendors. But will this trend continue? If you've been purchasing Ethernet switches for years without requiring storage vendor approval, why start now? And questions remain as to whether or not storage vendors will even have the time or desire to test all of the solutions. Will storage vendors retain final approval or will network vendors convince customers that it's an unnecessary step? It will be important to keep an eye on this potential shift in power.

FC won't disappear overnight. Remember when open systems were going to eliminate mainframes? There are still plenty of mainframes around and you could argue that the concept of virtualization is simply open systems recognizing a great benefit of logical partitions in the

Ethernet switches, sold into the network groups, haven't needed any approvals for NAS and iSCSI implementations. With FCoE, however, the game has changed.

mainframe world and adopting it. FC will be around for a while, but like ESCON and FICON, it may not be the fastest growing or most exciting segment to work in five to 10 years down the road.

CHART YOUR CAREER FOR CHANGE

Athletes cross train to break up the monotony of their routine and to increase their overall strength and endurance. IT shops should do the same. Begin to explore and educate yourself, but don't limit your studies to just your current responsibilities. Think about adjacent domains and always consider how a new technology will drive higher levels of service to the business.

Vendors, especially those with a convergence message, offer classes and certification programs to become better educated on these new technologies and virtualization products. Many are offered online and don't require travel. Take advantage of any company-sponsored training to expand your knowledge base and position yourself for future growth.

In addition, go to the certifying bodies themselves, like the IEEE and T11, to learn more about the status of CEE and FCoE. Other helpful sites include the FCoE home page.

LOOK FOR INTEGRATION POINTS

Top-of-rack switches are one of those areas where FCoE makes sense. They reduce the number of cables and cards deployed, but don't require a full rip and replace because they can direct traffic into an FC SAN or Ethernet LAN. You'll need to be ready to implement these new technologies when the opportunity arises.

Cisco is driving convergence with products and software, like the Nexus product line and NX-OS operating system, which combines storage networking SAN-OS and IOS into a single interface. The company's Nexus 1000 virtual switch resides in the hypervisor and replaces the VMware virtual switch. It's also controlled by NX-OS. You can expect Brocade, with its acquisition of Foundry Networks and its FCoE product announcements (converged network adapter and switch), to follow suit with a combined OS and single console to manage the storage and data network.

THE BIG PICTURE

Convergence is coming, so be supportive of these efforts and try to become part of the planning and testing teams. Remember, the two largest FC vendors made some major investments to solidify their portfolios last year. Become a proponent of mixed silos. Many companies have already begun to embrace the shift on a project-by-project basis. Network convergence in virtualized environments is still a relatively new model. Yes, there are products available and there's some testing taking place, but it's certainly not too late to get up to speed. This year, dedicate time to

learning about FCoE and meeting with vendors that supply this technology. Think of 2009 as the year to kick the tires. More advanced companies may start limited production deployments and by 2010 to 2011 most data centers should be on board with converged networks. The question, however, is: Will you be? ☉

Bob Laliberte is an analyst at the Enterprise Strategy Group.

Timetable for 10 Gige

The next generation of Ethernet is likely to have a profound effect on storage—pumped-up iSCSI performance may challenge Fibre Channel's tier 1 dominance.

By Christine Cignoli

IT'S BARELY EVEN HERE YET, but 10 Gigabit Ethernet (10 GbE) is going to have a hard time living up to its hype. Hailed as a “game changing” technology by some, it carries the burden of being a cure-all for storage (and network) managers’ problems. But when you look beyond the hyperbole surrounding 10 GbE, you’ll see the technology is, in many ways, still just emerging.

That’s not to suggest that 10 GbE won’t deliver on its promise—perhaps not a cure-all, but certainly destined to give iSCSI storage and Fibre Channel over Ethernet (FCoE) topologies a big boost. Still, real products are few and far between at this time, and per-port prices are still at a very un-Ethernet premium. While there have been some early adopters, its first few inroads into the market have been in higher-end implementations such as super-high-performance computing.

Switch vendors are starting down the 10 GbE road with product releases and certifications. But it may still take some time for 10 GbE to hit the mainstream. Brad Booth, chairman of the board of the Ethernet Alliance, thinks 2012 is when costs will be low enough for the market to see widespread adoption. “Generally when a standard is written, it takes about 10 years before it really hits what we consider the big volume adoption,” he said. Booth calls those 2012 adopters “tier 3,” and said that quicker-adopting tier 2 data centers are probably looking at 10 Gb now and considering

which cabling they'll choose. He cites Google as a tier 1 data center; they're now running a brand-new data center with 10 GbE throughout, according to Booth.

One research institute has already moved to 10 GbE. Bruce Allen, director of the Hannover, Germany-based Max Planck Institute for Gravitational Physics, chose a Woven Systems Inc. switch as part of their early 2008 systems upgrade. "I think 10Gb is now a proven technology," he said. And being an early adopter is standard for Max Planck. "We're a research institute," he said. "The one guaranteed way to fail is to wait. The wrong

10 Gig rides on fiber or copper

Optical fiber: The earliest 10 Gb Ethernet components were built around optical fiber, but copper cable-based products soon followed. In most data centers, copper is the standard transport for data and storage networks because it's relatively cheap and easy to install. But copper cabling typically can't carry a signal over long distances, so it's used primarily within buildings, while optical cabling is usually used to link facilities over longer distances.

There are a number of IEEE standards for optical cabling with 10 Gb Ethernet, each designated by a unique identifier. For example, 10GBASE-SR (short distance) is the standard for the cheapest optical implementation that uses standard optical cabling; it can cover distances of approximately 30 meters to 90 meters. The next step up in price and distance is 10GBASE-LR (long range), which can carry signals up to 10 km. There are several additional standards for 10 Gig over optics that offer greater ranges.

Copper connections: For copper cabling, 10GBASE-CX4 is an IEEE standard approved approximately five years ago that uses familiar twin-axial cable. Also referred to as 802.3ak, this standard is based on cabling and connectors used for InfiniBand, although some reengineering was done so they're not identical. 10GBASE-CX4 is effective at distances up to 15 meters.

10GBASE-T, approved in 2006, is likely to gain widespread popularity. It uses unshielded (or shielded) twisted-pair cables and will work at up to 100 meters; companies may opt to use already installed Cat 6 cabling, but the distance will be effectively halved. Connectors for 10GBASE-T are familiar, too, RJ-45-style connectors rated at 650 MHz.

decision is waiting.” The institute stores and analyzes data. Allen says storage is secondary for them, with compute cycles the primary goal. “I wanted wirespeed nonblocking to all nodes, and that meant Ethernet,” he said.

Woven was one of the earliest vendors to offer 10 Gb switches, along with others including, Cisco Systems Inc., Force10 Networks Inc., Foundry Networks Inc. (now part of Brocade) and Mellanox Technologies Inc. Joseph Ammirato, vice president of marketing at Woven, said 10GbE will be more affordable by 2010. And, for more widespread adoption, “we need 10 GbE as a default feature on server motherboards,” he said. “Users still have to buy it as a separate add-on.” He targets mid-2010 for 10 Gb to be built in.

Ammirato names several customer requirements he thinks will drive 10 Gb adoption: consolidation, server virtualization and what he calls desktop virtualization—a need for organizations to better control mission-critical applications, running them from the data center and not user desktops.

Mellanox is trying to move into the 10 GbE market from its base of InfiniBand products. T.A. Ramanujam, Mellanox’s senior product marketing manager, said that in the past few years, Ethernet has come into the spotlight as a consolidation or unifying data center fabric. “To accomplish that,” he said, “they’ve taken features from technologies like Fibre Channel and InfiniBand and added that to [a] new Ethernet specification, what’s going to be called data center Ethernet or converged enhanced Ethernet.”

Ramanujam says per-port prices are approximately \$300 to \$400 for a 10 GbE switch, and predicts that the market will start adopting 10 GbE when its price per-port cost drops closer to Gigabit Ethernet’s cost. “It’s

A SAMPLER OF 10GBASE-T NICs

Early on, networking vendors were challenged by issues such as power consumption and overheating with their designs for copper-based interface cards for 10 Gig Ethernet. Many of these issues have been resolved. Today, there are a number of vendors offering 10 Gig Ethernet network interface cards (NICs) that use the 10GBASE-T that supports twisted-pair cables and RJ-45 jacks, including the following:

Manufacturer	Product
Chelsio Communications Inc.	S310E-BT Storage Accelerator
Intel Corp.	Intel 10 Gigabit AT Server Adapter
Mellanox Technologies Inc.	ConnectX EN MNTH18-XTC
SMC Networks Inc.	SMC10GPCIe-10BT TigerCard
Tehuti Networks Ltd.	TN7588-S and TN7588-D (dual port)

not very far-fetched,” he said, estimating that the second half of 2010 will bring 10 GbE on the motherboard along with more affordable switches. “By then, the price of 10 gig switches will be in the region of \$100, \$150 per port.”

It also remains to be seen how FCoE will affect 10 GbE’s adoption and pricing. “I tend to think that Fibre Channel is one of those technologies that doesn’t really have a good place in the future,” said Allen, at the Max Planck Institute. “You’re better off sticking to commodity stuff.” But storage users managing the deeply ingrained FC infrastructures of many businesses might disagree.

In the meantime, the Ethernet Alliance’s Booth said the group is busy putting on interoperability demonstrations and educating possible 10 GbE users on what they’ll need to do to be ready for the transition—namely, consider their current server utilization and decide on cabling. “Looking at a server running 1gig, do they want to upgrade that server by putting in a new NIC and upgrade it to run 10 gig?” said Booth. “Or do they wish to put a better server in there and put this 10 gig card in it?”

Users must also think about whether they can reuse their current optical data center cable, or whether they’ll need or want to replace copper cabling. Allen said the one mistake the Max Planck Institute made in setting up 10 GbE was its cabling choice. They’re using copper 10 Gb cables, which Allen calls “bulky and inflexible.” For \$50 more apiece, he said, they could have had optical cables.

10 Gb also signals a shift in thinking toward consolidation and centralization. “Unlike previous Ethernet technologies, where primarily volume went to desktop,” said Booth, “now most of the volume of 10 gig is in the data center, into the core of the network.”

Christine Cignoli is a contributor to various TechTarget sites.

FC Switch Advice

Storage professionals can keep their Fibre Channel SANs up and running by following these best practices for switch environments.

By Carol Sliwa

ONCE A COMPANY DECIDES to go with a Fibre Channel storage area network (SAN), the next thing to do is settle on a switch vendor because it's not a good idea to mix and match brands.

Fibre Channel switches need to communicate and cooperate with each other to manage the overall fabric. The best way to ensure that that happens reliably is to select a switch from one of the top three switch vendors: Brocade Communications Systems Inc., Cisco Systems Inc. or QLogic Corp.

“There's a standard for this communication [between the switches], but the standard is kind of a weak, least common denominator of the functions required to build a SAN,” said Robert Passmore, an analyst at Gartner Inc. “All of the switch vendors have a much more robust overall set of management functions that are proprietary to [each of them].”

Some of the best practices that are common to all of the Fibre Channel switching environments are as follows:



Planning considerations

Plan your SAN for what you expect to need over the next three years.

Project your future needs based on the number of apps, physical servers and storage in use during the past two years. Take into account new technologies that may be deployed, such as virtual servers. Think about the impact different components will have on the overall environment.

“Whatever you think you’re going to need over the next three years, double it and build it for that,” said Marc Staimer, president of Dragon Slayer Consulting. “It costs you more, but it will make your life infinitely easier.”

Not planning your SAN upfront is “a nightmare of immense proportions,” according to Staimer. “The more you plan, the less rework you will have.”

Determine application throughput and I/O to size and design the SAN environment most cost effectively.

Most users don’t know their application throughput and I/O, said Dan Iacono, a SAN engineer at Philadelphia Insurance Companies, who was formerly senior SAN systems engineer in Hewlett-Packard (HP) Co.’s multivendor systems engineering group. Performance tools built into Unix and Windows or offered by network, SAN or storage vendors can help them calculate those metrics.

Understanding the applications and knowing their throughput will determine what type of ports (oversubscribed or full throughput) will work best and how to build out the SAN design to most cost effectively use bandwidth. Many users opt for a core/edge design, often with 16-port or 32-port switches at the edge going into a bigger director switch, connected via an interswitch link (ISL).

“You need to know your throughput on all your edge switches to connect the appropriate amount of ISL to your director,” Iacono said. Minimizing the ISL count can free up ports on the switch and give more money back to the SAN.

Companies with mature SANs may discover they need to shift an especially high-throughput application from an edge switch directly into the director to reduce hops and move it closer to the storage.

Don’t be afraid of oversubscription.

Most servers don’t require the full bandwidth of a Fibre Channel switch, so it’s common practice to oversubscribe or allocate more potential demand than the switch can handle because statistically it’s unlikely to need it all at the same time.

Still, Howard Goldstein, president of Howard Goldstein Associates, finds that administrators “tend to be conservative when they don’t need to be.” He said that, in most SAN environments, “you’re using one-tenth of the capacity of the switch port.”

Assess power consumption and cooling requirements in advance.

Technology vendors often consolidate their offerings into the smallest possible packages, but customer sites can't always handle them. Most requests for proposals come with questions about power consumption, according to Mario Blandini, director of product marketing in Brocade's data center infrastructure division.

"You'd be surprised at how many IT environments literally have no more additional electrical capacity," Blandini said. "Most [hospital or university] buildings were built 75 years if not 100-200 years ago. And when they put the electricity in, no one ever fathomed you would be consuming in a 19-inch square space 10,000 W of electricity."

Build two independent Fibre Channel fabrics for redundancy.

A SAN needs to be up 24/7. The more servers the SAN supports, the higher the consequences of failure. To make sure the SAN never goes down, there need to be two paths from the servers to the storage.

If there's a failure along one of the paths—with an HBA, switch, cable, port or anything—the other path allows the application and its storage to continue to communicate. Another benefit is that upgrades can be done while the SAN is operating.

"Fibre Channel is a disruptive technology," Staimer said. "Anything you change, anything you add, whatever you do to your system, will disrupt the application using it at that time. So what you do is force them onto one fabric while you make your change on the other one. You're the least disruptive when you have dual fabrics."

Management: The technical side



Deploy path management software to automatically switch the I/O request from one path to another in the event one path fails.

Some operating system environments provide basic capability. Some storage vendors have their own path management software that may cost more, but it offers additional features that may make it worthwhile, Passmore said.

Set up, tune and monitor hardware and performance alerts.

Iacono remembers a large consulting company that got 6,000 alerts per day and didn't do anything with them. One switch vendor used to have a default alert set to go off whenever the SAN hit 0 MBps. That could trigger a thousand emails per day.

"You simply had to turn that off," he said.

But even just a few hardware bit-level errors are cause for concern, since that could signal an impending failure. “About 95% of failure rate in SANs, we’re seeing [alerts] beforehand, but the alerting was not addressed,” Iacono said. “If you’re getting too many alerts, maybe you need to tune your alerting environment to get rid of the erroneous errors, or maybe there’s a real issue that you need to address.”

Back up the SAN configuration information to a hard drive not on the SAN.

SANs don’t go down much, but when they do, they go down hard. If the SAN documentation is backed up on a server connected to a network drive, and that drive is over the SAN, the storage team will lose the information it needs to restore the systems.

“I could tell you Fortune 50 companies that do this,” Iacono said. “It’s amazing.”

Many companies don’t even have updated documentation. They often start with an Excel spreadsheet and the best intentions, and then rarely update it because they have more pressing responsibilities.

“If they have to troubleshoot something, they have no idea what’s connected to what port,” Iacono said. “I’d say everyone has some sort of documentation. Probably 50% to 70% [of it] isn’t up to date.”



Management: The personnel side

Employ a dedicated storage team and rigid change management procedures.

When a SAN goes down, it’s usually because of human error. Strict change management policies reduce the chances that will happen. So does a dedicated storage team that manages the systems proactively.

Server administrators need to communicate and coordinate their needs with the storage group, which handles the storage design. Storage pros write down the process steps and setup instructions, including the actions on the storage array and the switches. Ideally, another storage specialist reviews the change design and quality assurance is done.

“Organizations that follow these kinds of processes are the ones that, in essence, go year in and year out without ever having a failure in the SAN,” Gartner’s Passmore said.

Set separate user accounts and passwords for each administrator and third-party consultant with access to the SAN.

It’s not uncommon for an administrator with a new SAN switch to tweak parameters and not tell colleagues, Iacono says. When he finds a switch configured differently and asks what happened, he usually

hears that “Joe was doing this and Steve was doing that.”

According to Iacono, “Once you create accountability, all that disappears. We want to be able to audit who’s doing what.”

Not only will the IT group be able to determine the source of any problems, it won’t need to reset the universal password when a SAN administrator leaves the company.



Security

Create zones at the same time LUN masking and binding is done.

When storage is created for a new server, tools are used to carve out a storage volume and give it an address, or an SCSI LUN. LUN masking hides the LUN from entities that don’t own it; LUN binding attaches the LUN only to the worldwide ID of the HBA in the server.

At the same time that LUN masking and binding is done, a storage specialist should go into the switch and create a zone that will allow only specified adapters to talk to certain storage ports.

“In essence, the switches, through zoning, reinforce the LUN masking and binding,” Passmore said. “And in one more step, switches at the port that talk to the server can be programmed to check the worldwide ID and therefore reinforce the LUN masking and binding that’s been done in the storage arrays.”

Use Secure Shell (SSH) protocol to access the SAN.

If an administrator logs into a SAN switch using the Telnet protocol, the password isn’t encrypted, leaving it at risk of interception. SSH provides a secure channel.

“With SSH, everything is encrypted,” Iacono said. “This is a standard if you’re managing your Windows or Unix environment, but for some reason, no one does this for SAN environments.”



Virtual servers

Make sure the bandwidth in and out of the servers into the switches and the targets is adequate to accommodate the environment.

When once underutilized servers run multiple application workloads on virtual machines, the bandwidth requirements escalate. Users need to design their SANs with that in mind.

“A typical x86 server last year would be hard-pressed to do more than a gigabit per second of throughput,” Staimer said. “The current generation of x86, [which is] typically dual-quad core, can easily push 10 Gb, if the applications can. If you’re running 20 applications concurrent-

ly, you're going to push that 10 Gb. It's pushing the I/O that in the past the server really didn't push, because one application was rarely going to do it."

Make sure every physical server with virtual machines is in the same zone.

Using virtual server technology, an administrator can move an application from one physical server to another without any downtime, but those physical servers need to be in the same Fibre Channel zone to be able to access the storage.

"What happens when an application can't see its storage? It crashes!" Staimer said.

Use switches and HBAs that support N_Port ID Virtualization (NPIV).

If one physical server has five virtual machines running on it, NPIV will permit each of those virtual machines to get a unique identifier on a single HBA, and an NPIV-capable switch will recognize each distinct ID. That, in turn, means each virtual machine can have access to a different LUN.

Without NPIV-capable devices, the physical server would get one port ID. NPIV is supported in new switches and HBAs, but anyone using legacy hardware might need to check with the vendor about a firmware update. NPIV works with blade servers similarly to the way it does with virtual machines. ☉

Carol Sliwa is the Features Editor for SearchStorage.com.

iSCSI switch choice

Solution providers can optimize a customer's iSCSI SAN by guiding them toward the right Ethernet switch.

By Jerome Wendt

Storage magazine's Spring 2009 purchasing intentions survey indicated that about 44% of respondents who work with value-added resellers (VARs) either have deployed or will deploy iSCSI storage in 2009. This means solution providers need

to identify and put in place best practices to optimize their customers' Ethernet networks to support these iSCSI SANs. Toward the goal of delivering problem-free, scalable iSCSI SAN configurations, solution providers will need to choose and deploy the right type of Ethernet switches and properly configure Ethernet switch features like flow control and jumbo frames for optimal iSCSI performance.

While choosing an Ethernet switch may seem like a no-brainer, since any Ethernet switch that supports TCP/IP traffic will support iSCSI traffic, it really is a strategic decision. The right Ethernet switch can better position a customer's iSCSI SANs for future network and storage growth, and it can improve iSCSI SAN performance. Factors to consider in this choice are the storage system and the SAN's growth rate.

Here's why the storage system in use is an important factor in your customer's choice of Ethernet switch: Some storage systems, such as Dell Corp.'s EqualLogic PS Series iSCSI storage system, are stackable systems that can create one logical virtual pool, or group, of storage. To create this storage group, every member in the group needs to communicate with the others through the network. Also, every server host needs access to every member in the Dell EqualLogic storage system group.

In that scenario, if your customer expects minimal or no growth, 24- or 48-port switches with support for interswitch links (ISLs) will be fine. ISLs connect different Ethernet switches and route TCP/IP and iSCSI traffic between them. Using switches that only support ISLs, storage systems can communicate with one another, and any server can access any storage system connected to the iSCSI SAN.

If your customer expects rapid growth of its iSCSI SAN, ISL-supported switches are not the best choice, because ISLs' use of Ethernet switch ports increases as members are added to the storage group. Instead, your customer should use stackable Ethernet switches for fast-growing iSCSI SANs. These switches have dedicated ports for routing TCP/IP and iSCSI traffic, keeping more ports free for server and storage connections. Stackable switches are also available from some vendors with 10 Gigabit Ethernet (10 GigE) ports, which can further improve performance on an iSCSI SAN.

Another consideration in Ethernet switches is OSI layer. Layer 3 routing is preferable for switches that have dedicated ports for TCP/IP and iSCSI traffic; Layer 3-supported switches can route traffic between switches based on the server or storage system's IP address. Layer 2 Ethernet switches route network traffic only based on MAC addresses, so if the switch cannot resolve the MAC address, it needs to send the packet to a router. Depending on the size of the iSCSI SAN and the amount of network traffic, this can become a performance bottleneck, one that Layer 3 Ethernet switches avoid.

KEY ETHERNET SWITCH FEATURES

Customers that have a lot of sequential traffic, such as video streaming or backup, can benefit from the correct implementation of the jumbo frame feature on Ethernet switches. Jumbo frames optimize end-to-end data transfers by putting more data in each TCP/IP packet, which minimizes transmit and receive interrupt processing on server and storage hosts. Jumbo frames are by default about 9 KB in size, so there needs to be enough data to fill the jumbo frame, and both the initiator's and target's network cards need to support jumbo frames.

Flow control, the process of managing data transmission rates between two nodes, is also important. Most switches handle flow control in a similar manner, but you should check the flow control settings on the host's Ethernet network cards to ensure that buffer credit allocation is sufficient for each application. With correct buffer credit allocation, the card can tell the switch to slow down traffic to avoid dropped packets. Without flow control, packets can be dropped and need retransmission.

Customers that have a lot of sequential traffic, such as video streaming or backup, can benefit from the correct implementation of the jumbo frame feature on Ethernet switches.

A word of caution when your customers use both jumbo frames and flow control: Using both features can affect storage system stability, according to Jay Kramer, iStor Networks' vice president of worldwide marketing. iStor has seen some situations where the Ethernet switch did not have enough horsepower to drive jumbo frames when flow control was enabled; this affects system stability, Kramer said. Before turning on both jumbo frames and flow control for your customers, you should verify with the storage system vendor that using these features at the same time will not negatively affect the storage system. ☉

Jerome M. Wendt is the founder and lead analyst of The Datacenter Infrastructure Group.

Check out the following resources from our sponsor:



[Whitepaper: More Functionality, Less Infrastructure: Managed IT with Double-Take Flex](#)

[Double-Take Software Goes Green: 24/7 Access to all Datasheets, Whitepapers and More.](#)

[Watch this Video to See Double-Take Software's Failover and High Availability Capabilities in Action.](#)