

*Early reports from users praise IP storage for its performance and management of primary storage and backup and restore. Some initial problems with implementation seem to be the only sore point*

# IP storage delivers

The rumors about IP storage really are true, early adopters say. The technology is fairly easy to set up and manage, and performance is fine for many types of applications—assuming that the IP-based networks or storage components are designed and/or implemented properly.

BY JOHANNA AMBROSIO

# Users' own definitions of what constitutes IP storage or an IP SAN can vary widely

IP storage brings many benefits to the data center, including lower costs compared to Fibre Channel (FC) storage area network (SAN) products. We talked to eight organizations, some of which have been using IP SANs—or IP-based storage in general—for a year or more. Three are using IP as their primary storage mechanism for everything from database to e-mail and customer billing applications; one is using IP for backup and restore only; two are using IP for primary storage and backup; another is using its IP setup as a file server to replace a failing network-attached storage (NAS) box. One organization, Sandia National Laboratories, is still in testing mode and hasn't begun using its IP SAN in production.

In the user world, some installations have mixed FC and IP protocols—as opposed to pure FCIP or iSCSI. This is often due to older storage boxes—or even fabrics—that may use FC to link to other IP-based components of a SAN. This is the situation at the Cancer Therapy Research Center (CTRC) in San Antonio, TX, which is using older, FC-based disks with an IP-based storage router to get its planned ROI out of the storage boxes.

Users' own definitions of what constitutes IP storage or an IP SAN can vary widely. Some are using IP storage over WANs or are using iSCSI over IP or Ethernet networks to connect servers and arrays. Some are doing IP within one data center; others are using the protocol to connect far-flung storage boxes and servers.

Most of the IP implementations are still fairly small at this point, with an average of 3TB to 6TB of actual storage. In most cases, this is a fraction of the storage still in direct-attached mode or even on an FC SAN.

Five of the eight organizations are using an FC SAN in addition to IP, or at least started with FC. All have different ideas about connecting—or not connecting—the two worlds. The global law firm Clifford Chance is hoping to connect its Fibre and IP SANs in the next year or so. The Public Broadcasting Service (PBS) plans to use both for different applications for the foreseeable future—Fibre for its high-performance, transactional needs and IP for pretty much everything else—with hopes that it can someday manage the two from one management suite. Mortgage provider HomeBanc and water-purification supplier Zenon Environmental will continue to use their FC gear for one dedicated application that's already on the system, but will use IP for almost everything else.

## iSCSI for primary storage

There are many reasons to get involved with IP storage. The CTRC implemented an IP network approximately two years ago and hasn't looked back. The CTRC of-

fers radiation and oncology services to major hospitals in four locations, and serves roughly 200 patients each day. Its two primary facilities are 22 miles apart.

The center has seen a huge growth in its storage needs, driven by digital imaging requirements. Many types of scans—CT, MRI and others—move around the network on a regular basis. In 2001, the organization built a new main campus, which was an opportunity to redesign its major IT infrastructure, says Mike Luter, CTO. The major criteria was availability and reliability, because “to have someone miss or have to reschedule a radiation treatment is simply unacceptable.”

Luter says it was clear from the get-go they'd need Gigabit Ethernet to provide adequate bandwidth for its images and other files. So while they were at it, CTRC decided to kill many birds—iSCSI storage, voice over IP telephony and an application server at both data centers—with one stone. And with this setup, one IT staffer can now handle all of a user's needs.

Of its 45 major servers—mostly Compaq—some two-thirds are on the IP SAN, with approximately 7.5TB of total storage. There's a spare Compaq DL360 g2 at each location, and there's synchronous mirroring of data for business continuity.

The setup includes two EMC Clariion 4700s, with Cisco SN 5428 storage routers. The 4700s are actually Fibre-based, so they link to the 5428s via Fibre, and then the IP-based 5428s link to the backbone network.

## Storage for Space Odyssey

The Denver Museum of Nature & Science started planning a major new project called Space Odyssey and knew that its current technology wouldn't be up to the task. “Our storage is almost all direct-attached and we're in transition,” says Vince Wolfe, program technology manager. The museum is in the process of changing into a Microsoft shop from an old Novell environment.

Space Odyssey, though, needed a new kind of storage. The hands-on exhibit and all-digital planetarium required guaranteed uptime and scalability, in part to help deliver information-on-demand to human guides that are a critical part of the program.

The guides carry laptops to look up a range of information to help answer questions—everything from the history of space exploration to recent information gleaned from the Mars rover.

With this in mind, Wolfe says, it was clear they needed a SAN for shared storage, and the organization started looking at alternatives. “We did a pretty thorough evaluation and looked at Fibre Channel” and IP suppli-

Most of the IP implementations are still small, averaging 3TB to 6TB of actual storage

ers, Wolfe says. "We decided that an IP SAN was more flexible than Fibre." It was also much more affordable.

Ultimately, the museum went with LeftHand Networks as its IP SAN supplier, and now has 11 Network Storage Modules (NSMs). One reason was the architecture's "easy expandability," Wolfe says.

INTEC Engineering, a Houston-based company serving the oil and gas industry, had another equally clear-cut reason for looking into IP—its NAS was failing. "It was a cheap box, one of the first IDE-based ones that had come out," says Chris Warlick, IT director. "We were beating the heck out of it, using it beyond anything it was originally intended for." After it crashed several times, the company started looking around for something else. As Warlick recalls, "When your data is down, it's an easy sell."

They, too, looked at Fibre, but concluded they "wanted

to save money and to be leading-edge, but not necessarily bleeding-edge." When its existing networking supplier, Cisco, started selling an IP switch, INTEC bought one.

At this point, the SAN functions as a file server, with the firm's SQL Server and Exchange application data housed in DAS. As Warlick explains, "If we get to the point where we feel it would benefit us to move the actual e-mail and database data, we'll do that."

#### Backup and disaster recovery

Clifford Chance, a New York-based law firm, wanted a better way to perform disaster recovery and backup "without interfering with our existing production systems," says Milton Morgan, senior IT manager. And its existing EMA 12000 SAN from Hewlett-Packard Co. (HP) was too expensive to add capacity for this task, at least from a main-

## Top ten IP storage tips

If you're thinking about going down the IP path, here's some advice from early adopters to consider.

- 1 **DESIGN CAREFULLY.** Sure, this is true for any storage area network (SAN), but it's especially important because IP performance might be a bottleneck if there are large distances or many remote locations involved.
- 2 **GET YOUR STORAGE and networking experts together, and make them play nice.**
- 3 **MAKE SURE YOU'RE starting with at least a Gigabit Ethernet backbone to ensure adequate performance, especially if you're planning on using IP to handle your primary storage.**
- 4 **OUT-OF-THE-BOX Gigabit Ethernet NICs work just fine for most applications; you'll want to consider TOE or other kinds of interface adapters only for the applications with really specialized or high-performance needs.**
- 5 **SEPARATE THE IP SAN TRAFFIC on its own VLAN, for both performance and security reasons.**
- 6 **UNDERSTAND WHAT YOU'LL need to do to your applications to make them happy on iSCSI—making registry changes, pointing SQL Server files to the correct locations, etc. While Fibre Channel customers have long had to deal with this, these issues generally aren't part of the direct-attached world, and there might be a learning curve involved for IT staffers.**
- 7 **TAKE ADVANTAGE OF THE FACT that many vendors in the IP space are small and hungry and really want your business. Some customers have demanded—and received—benefits, including extra consulting and support and even money-back guarantees.**
- 8 **MAKE SURE YOUR IP SOLUTION of choice supports all the platforms you need—Linux and Unix iSCSI drivers aren't as popular as those for Microsoft.**
- 9 **START SMALL, with a couple of applications and/or servers, and grow from there.**
- 10 **DON'T BE AFRAID.** Despite some vendor-driven FUD out there, every single person we talked to says they're happy they went down the IP road and that things are working out just fine.

tenance perspective. So, they're running FalconStor's IPStor software on an ATA-based StorageTek box connected to an IP network via Cisco 6509 switches. Under this scenario, the firm mirrors its DAS to the IPStor/StorageTek setup in New York, then replicates the mirror to

the office in Washington, D.C. All of the firm's legal documents and home directories are handled this way, to the tune of about 4TB.

This approach saves money, but perhaps not for the most obvious reason. "There's not a huge difference in the price of the drives," Morgan says. There used to be, but HP and IBM Corp. have been "more willing to bring their storage prices down to ATA levels." The bigger problem is the cost of maintenance. The HP SAN—with a little more than 3TB of storage—costs \$12,000/year for maintenance. And the StorageTek box—with more than 4TB of storage—costs approximately \$6,000 in annually.

### The performance question

The motivations for going with IP storage are as different as the organizations using the gear—but there are some constants. Performance was an issue for some companies. Most of the companies interviewed are using Gigabit Ethernet switches, network interface cards (NICs) and associated gear, and have put their IP traffic onto its own VLAN for both performance and security considerations.

The Denver Museum of Nature & Science's Wolfe says flat-out that "there's a performance hit if the network's really busy."

Other adopters, though, say they have no complaints on the performance side. Helen Chen, network researcher at Sandia National Laboratories, says she has seen peak performance of 180Mb/s over its IP network for a write function. "I believe iSCSI has the potential to deliver performance for scientific computing," she says. The lab has already proven that IP performs "as well as Fibre" in a local connection "if we tune it correctly."

The lab is also more interested in open-source solutions. On the hardware side, the problem is that "you can't tune hardware very easily" to do things like increase the distance between initiator and target, Chen says. "So, we used the software stack in the Linux host to modify parameters," including iSCSI command size, TCP/IP window size, etc. "We want to open up the parameters so we can fill up a big, fat pipe between two remote locations," she says. "We're working with Adaptec to get the interface to be able to modify the TOE [TCP/IP offload engine] hardware."

PBS, Alexandria, VA, uses a StoneFly IP concentrator to network the storage for the applications that don't require the highest bandwidth provided by its other SAN, based on FC. They're "quite pleased" with IP's performance, but Ken Walters, senior director of enterprise platforms, emphasizes that not all applications are natural fits for iSCSI. "I use it for equipment used for development

work, and on machines that aren't critical or that don't have massive I/O requirements." These are applications that "I can't cost-justify connecting to my Fibre SAN," he says.

And so, he says, PBS is starting to provide iSCSI storage where they really need it. Some of the applications connected via the StoneFly device are a Linux box that processes the logs from the PBS Web site and a SQL Server cluster that's a development environment for an internal PBS user. Walters' group recently received an iSCSI cluster router from StoneFly. "I wanted to get clustering set up before I handed out a lot more storage because I didn't want a single point of failure," he says.

"So now, we'll start to provide iSCSI storage where appropriate." The organization uses the StoneFly Concentrators to stream all Real and Windows media from the Web site.

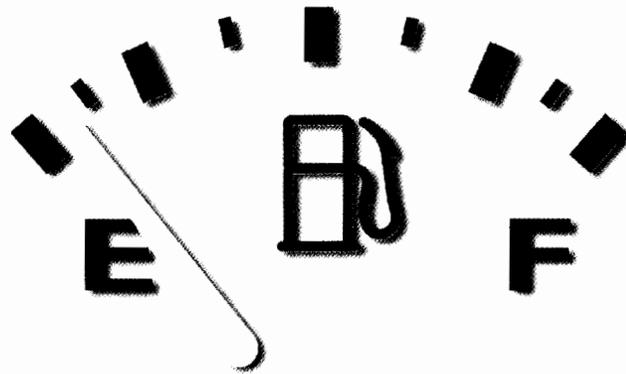
At Zenon Environmental in Ontario, Canada, Shawn Eveleigh, senior systems administrator, says he's happy with his PeerStorage setup from EqualLogic. He used utilities in Exchange that test how fast a disk subsystem performs. "Our total disk I/O speed was around 40 to 45Mb/s, and most of our systems will never get to that kind of load," he says. "That was forcing the server to go as fast as it possibly can. And the disk I/O kept up with that."

### Designing for performance

Almost all the organizations we talked to are using standard-issue, out-of-the-box Gigabit Ethernet NICs, but PBS is an exception. The group is using a combination of regular NICs and cards that of-

flood part of the IP and iSCSI processing load from the server. These include Intel Storage Network Interface Cards (SNICs) and TOE cards from Alacritech and other suppliers. "The offload does work as advertised," PBS' Walters says.

Which type of card he uses to hook up each server to



## Are your gauges really telling you the truth?



by our booth at  
Storage Decisions and  
enter to win a  
MERCEDES SLK

Sometimes when you think you are running on fumes, there's still enough to get you where you want to go. The same may be true for your storage environment. Global Storage Manager from Storability Software enables you to measure and manage precise levels of your valuable storage capacity. Join the growing list of Storability® customers who now can see what they've been missing.

Visit us at [www.storability.com](http://www.storability.com) or call us at 508.229.1700 and find out how you can get extra mileage out of your storage investment.

**STORABILITY**  
SOFTWARE

Storage Management. Made Simple.

© copyright 2004 Storability and the Storability logo are trademarks of Storability, Inc.

# Most of the customers who bought IP for primary storage plan on expanding it to handle backup and restore

the SAN depends on how much processing power the server has to deal with iSCSI's additional workload. "Much of our hardware is fairly new," he says, "with plenty of processing power. So, I use mostly regular NICs, for about \$130, where SNICs are about \$550 and TOEs are around \$850."

There are some nitty-gritty things to consider for performance, too. Akil Woolfolk, network operations team lead at Atlanta-based mortgage lender HomeBanc, says they had to make registry changes to make sure SQL Server works well with iSCSI. "We had to do some configuration to create a dependency" to make sure the SQL Server wouldn't start before the iSCSI initiator was up and running," he says. Otherwise, the SQL Server will crash because some of the files it needs to see aren't present.

Also, Woolfolk and his team at HomeBanc figured it was wise to plan where to put each application's database associated log files. "That makes a difference in your implementation," he says. "We put the log files on the local disk, with the OS to help boot up storage. Then, we run the iSCSI initiator to connect the servers to the IP SAN."

HomeBanc is using EqualLogic's PeerStorage to host file services for more than 500 users, with more than 80% of its production SQL data on the IP SAN.

## Redundancy and security

INTEC bought two Cisco Catalyst 3550 switches and two Cisco 5428 storage routers for redundancy in its IP network. "We did some absolutely rigorous tests from the standpoint of user and server access," Warlick says. "What does a user see if the router fails, if the drives fail? We tried every failure you can imagine, and the test results were phenomenal. If we turned off one of our switches, that would delay the server connection for a few seconds. But you'd still see the data."

Some customers say they're not too concerned about redundancy from their IP SANs. "I'm not trying to recreate the high availability of a FC SAN," PBS' Walters says. "If I need the ultimate in load-balancing and failover, the application belongs on Fibre," and not on IP. Indeed, if you try to design the same level of redundancy and security into IP, you start to erode some of the cost benefits of the technology, he says.

With PBS' clustered StoneFly gear, however, that level of redundancy is already built in. "I pulled the power cable, NIC cables—I powered them off, and the machine never lost a beat," Walters says. Unlike Fibre, where "I'd be lying if I said we didn't have problems. Fabrics can collapse due to hardware problems; sometimes multipathing doesn't work and you lose your connec-

tion to the [Fibre] SAN," Walters says. "If you ask people and they're honest about it, they'll tell you that in practice there can be as much or more downtime with Fibre as with DAS if you're not very careful."

Another suggestion Walters has for redundancy in an IP environment is to put in multiple NICs and run multiple drops to different IP switches. It doubles the NIC costs, but it's still less expensive than going with all Fibre, he says.

On the security front, "eventually we'll be able to use IP-based authentication, but it's not quite there yet," says Sandia's Chen. "Password-based security isn't good enough for us."

But for most users, putting the IP SAN on its own VLAN—and using whatever built-in security comes with their associated products—seems to be good enough. HomeBanc is using the security model that already comes with its PeerStorage SAN. With that, "you have to enter passwords, not just set up a workstation and run the iSCSI initiator and connect," says Woolfolk. "And because it's on a separate VLAN, you can't just hit the PeerStorage box from anywhere on the network."

At Zenon Environmental, Eveleigh says, "Because it's a separate network, it's not exposed to anything. It sits behind the servers, so if someone plugged into our network, they wouldn't see the storage side."

"Most people don't worry much about security," Walters says. "You'd have to break into my DC [data center] to get anywhere near the Fibre network. IP is much closer to the hackers, so maybe we do need to start worrying about it."

## Driver and other problems

Most of the issues with drivers seem to be resolved with Microsoft's iSCSI Initiator. But that only works with newer versions of Windows—after Windows 2000—and the Linux and Unix worlds don't yet have this same level of standardization, customers say.

The museum's Wolfe says he's looking forward to implementing LeftHand's new Linux drivers as a way of driving even more IP storage in his shop. "But I doubt they'll ever have direct support for Irix," also used at the museum. In general, the more homogenous the environment, and the more up-to-date the operating system, the fewer problems encountered on the driver front. But there are other kinds of things that bite customers, too.

Stan Rehffuss, senior system admin at the museum, says that with the Cisco 4507 switch, the 48-port cards are set up in a way that you need to distribute your gigabit ports or else you don't get the full gig for performance. "It's just

# IP SAN more Fibre”

Vince Wolfe, program technology manager at The Denver Museum of Nature & Science.

a limitation on the 4507,” he says.

Clifford Chance ran into problems, all of which “are behind us now,” Morgan says. The firm wanted to mirror its DAS to a cheaper ATA system, so they bought a StorageTek ATA box. “We mirrored the data, and found that after 10 days the server abended,” he says. “So we went to Novell and found there was an issue with the way that NetWare 5.1 worked with IP—it didn’t. Version 6.0 worked, but we were in the process of moving to NT file storage.” Once they moved to NT, it worked fine. But “the issues on the Novell side forced us to increase the pace of a planned migration that was already in progress to Windows NT,” says Morgan. Another problem area was implementing jumbo frames on its Cisco switches, to allow more data to be moved at one time.

“When we did that, we had to enable jumbo frames on both switches,” says Morgan. “But then that affect-

ed other network traffic. It may be the hardware we’re using, or it may be the version of Cisco IOS that we’re on. But the only solution we came up with was to disable jumbo frames on our main production switches, and then move the IP SAN to its own, separate, much smaller switches.” And not everyone’s run into problems on that score, either. Zenon Environmental enabled jumbo frames with no problems on its Dell Gigabit switches, Eveleigh says.

Overall, Eveleigh and the others say they’re happy they invested in IP SANs. Most of the customers who bought IP for primary storage plan on expanding it to handle backup and restore—and vice versa. As Eveleigh says, “Once you set it up, it’s not that fancy a technology. It does its job. I don’t worry—and that’s a good thing.” ☉

*JOHANNA AMBROSIO is a freelance writer in the Boston area who focuses on computer technology.*