

REPRINTED from

COMPUTERWORLD®

THE VOICE OF IT MANAGEMENT ■ WWW.COMPUTERWORLD.COM

APRIL 4, 2005 ■ VOL. 39 ■ NO. 14

INVASION OF THE iSCSI ARRAYS

As the technology has matured, IP-based storage arrays have established a beachhead as the preferred low-end SAN option.

BY ROBERT L. MITCHELL

WHEN JIM TARALA oversaw the rebuilding of his firm's network and IT infrastructure last year, he never dreamed he'd throw out his Fibre Channel SAN in favor of IP-based networked storage. The savings, however, were just too large to ignore.

Tarala, CIO and chief technology officer at Schenck Business Solutions, a 500-partner accounting firm in Milwaukee, was comfortable with his EMC Clariion storage system, but it was running out of space. Tarala also wanted to eliminate direct-attached storage on his mySQL and Microsoft Exchange Server systems in favor of networked storage. He decided to replace the entire system but initially dismissed the idea of using IP storage-area

networks (SAN) — systems that use the iSCSI protocol to allow servers to access stored data over an IP network — rather than direct-attached or Fibre Channel arrays. "I wasn't comfortable with the overall architecture," Tarala says.

Then he discovered that replacing the aging FC700 with a 1.5TB system would cost more than \$90,000, while a 2.5TB iSCSI-based PS Series system from EqualLogic Inc. in Nashua, N.H., would cost just \$47,000. Schenck gave Fibre Channel the boot. Tarala now has three PS Series systems with more than 7TB of capacity that support 12 servers.

"I spent half of what I budgeted, doubled my capacity, and it performed flawlessly," he says.

More than two years after the iSCSI protocol was ratified as an Internet Engineering Task Force standard, early adopters say IP SANs are not only ready for production deployments but also offer an alternative to Fibre Channel for low-end and midrange storage. Performance and reliability of iSCSI arrays have improved, and iSCSI SANs are significantly less expensive to set up and manage than Fibre Channel SANs, users say.

The high costs of traditional SANs have restricted the technology to mostly first- and second-tier data center applications. Now storage administrators are setting their sights on iSCSI as an alternative for some second-tier applications.

A Second Wave

And a second wave of storage consolidations is already under way: Administrators are replacing direct-attached storage used in departmental servers with local, iSCSI SANs. Robert Gray, an analyst at IDC, says the strongest growth is coming from large companies, which have huge numbers of departmental servers that can benefit from consolidating storage.

Early misgivings about iSCSI's capabilities have faded. Tarala was wary of the performance and reliability of the Serial ATA (SATA) drives used in EqualLogic's system but says the new SAN's performance has been comparable to the system he retired.

The new storage system is also more efficient to run. Because the SAN is IP-based, Tarala's Windows server administrators can manage it.

"You don't have to wait to bring in an outside technician to check that the [Clariion] tuned the array," he says. And moving from direct-attached to networked storage has improved staff productivity. "Everything boots from the SAN, and they can bring up a new Windows 2003 server in 20 minutes," Tarala says.

Until recently, the major SAN vendors were reluctant to release iSCSI products, citing maturity issues and server CPU performance bottlenecks that might arise in processing traffic associated with iSCSI and

the chatty TCP/IP protocol. But thanks in part to faster processors, that bottleneck never materialized for Tier 2 applications.

Smaller Vendors Lead

Smaller vendors, such as EqualLogic, StoneFly Networks Inc. in San Diego and Intransa Inc. in San Jose, have taken the lead in offering iSCSI target storage, while larger players such as EMC Corp. have offered iSCSI ports on Fibre Channel SANs and multiprotocol switches.

But native iSCSI storage arrays are less expensive, and early buyers are deploying them to support e-mail and database servers, backup and other departmental applications that don't require the high I/O that Fibre Channel delivers — and that won't support the cost of Fibre Channel SAN switches and host adapters.

Now Tier 1 vendors are jumping in. IBM, which withdrew an early iSCSI array, has returned to the market with the TotalStorage DS300. EMC recently announced new Clariion AX and CS Series models that offer native iSCSI connectivity to Fibre Channel, parallel ATA or SATA disk arrays. Every major vendor will have a native iSCSI SAN offering by year's end, says Gray.

That's important to users such as Robert Stevenson, a technology strategist at Nielsen Media Research Inc. in New York. "Initially, we were very cautious about moving into the iSCSI space [because] the larger players were dismissive of it," he says. But like Tarala, he found the cost and manageability benefits outweighed his initial concerns.

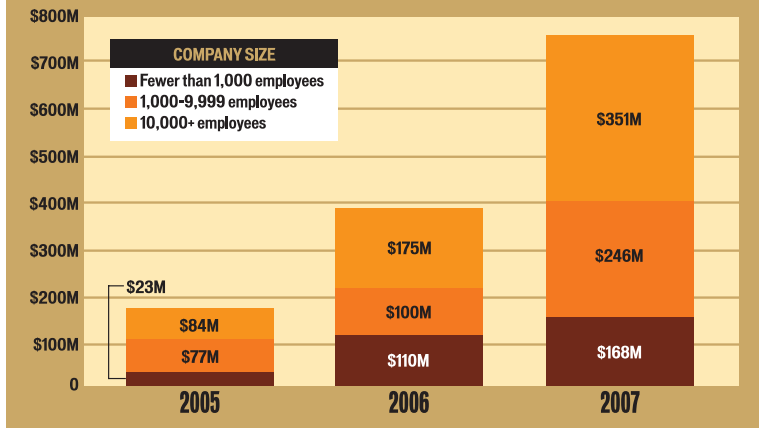
Stevenson started last year with an IP5000 iSCSI target array from Intransa to support a virtual tape library application, then he added another unit to house small Sybase Inc. databases that sit behind Nielsen's TV ratings system applications.

A third IP SAN from EqualLogic is in the lab as part of a project to host a larger data warehouse. At about \$3 per gigabyte, storage on the Intransa system is "very economically compelling," Stevenson says.

With more than 10TB on iSCSI-based storage, Stevenson's biggest concern now is rolling administration of those systems into the storage resource management tools that control the rest of his 1.2 petabytes of networked storage. Unfortunately, his current tools don't fully support iSCSI.

Management of iSCSI SAN

PROJECTED iSCSI ARRAY SPENDING



systems has trailed behind hardware and infrastructure development, and standards like the Storage Networking Industry Association's (SNIA) Storage Management Interface Specification have yet to catch up.

Support for mixed environments like the one Stevenson is considering are an even bigger challenge. "How do I manage an end-to-end environment when the iSCSI host may be several hops away on a Gigabit Ethernet switch or IP router and the proxy Fibre Channel target is on the other side of a multiswitch SAN?" he says.

But those concerns aren't stopping users from creating stand-alone IP SANs to target specific applications. For example, Siemens Corporate Research Inc. in Princeton, N.J., added an iSCSI interface to its Network Appliance Inc. filer to back up its IBM Rational ClearCase change management software. ClearCase wanted to issue block-writes to direct-attached disks; an iSCSI interface allowed the storage to be migrated to the filer, where it could be backed up using NetApp's

St. Paul, Minn. Expanding the direct-attached arrays was sometimes problematic and required taking the system down after hours.

Waslie moved the data onto three iSCSI-based Network Storage Module 150 appliances from LeftHand Networks Inc. in Boulder, Colo. "The last time I had to expand [storage], I did it in minutes — and I didn't have to come in on a Sunday," he says.

Still, experienced users have other reservations about IP-based storage — especially with regard to the IP network. Waslie isolated his IP SAN traffic on a physically separate network for security and to allow for out-of-band management.

Stevenson says project planners should make sure sufficient bandwidth is available on the existing network before adding iSCSI traffic. And adding IP SAN devices, which require static IP addresses, also increases complexity. "These static IP connection points make it very different to upgrade the storage frames in a heterogeneous environment," he says.

While iSCSI is gaining ground for backups and second-tier applications, Stevenson already envisions using IP SANs for a more mission-critical application at Nielsen. His group wants to create a copy of a 40TB data warehouse for the development team, but without spending \$4 million on a Fibre Channel SAN. "You can get cheap blade servers and get iSCSI to them and put it on SATA [disk arrays], and you've got low cost," he says.

As Nielsen eventually migrates to 10 Gigabit Ethernet, Stevenson expects IP SANs to move still higher into his tiered storage architecture. Even today, he says, "it tends to perform at a higher tier than you would think."

IP SAN SPECS AND STANDARDS

iSNS: The Internet Storage Name Service, currently an IETF draft standard, provides for both automated discovery and authentication support for iSCSI devices. The standard is expected to be finalized this year.

SMI-S: The Storage Management Interface Specification is a SNIA initiative to develop a common management interface for storage networks. SNIA's work is being standardized through the InterNational Committee for Information Technology Standards. Support for iSCSI within the SMI-S specification is still evolving.

MPIO: Microsoft Corp.'s multipath I/O technology enables multipathing for Win-

dows hosts attached to iSCSI or Fibre Channel SANs. MPIO can be used to facilitate fail-over or load balancing. A new feature in Version 2 of Microsoft's iSCSI initiator for Windows, MPIO is likely to become the de facto standard for providing multiple communication paths between Windows hosts and IP SANs.

VDS: Virtual Disk Services is Microsoft's tool for managing heterogeneous storage systems for Windows systems. The current version supports only Fibre Channel, but VDS 1.1 will also support iSCSI. It will be released in mid-2005, six weeks after Service Pack 1 for Windows Server 2003 ships, says Microsoft.

— Robert L. Mitchell

AT A GLANCE