

S STORAGE

PIPELINE *Managing Business Data*

» Cover Story

iSCSI STORAGE:

READY FOR LIFTOFF

By Steven Hill

Although iSCSI accounts for only 2 percent of the SAN market, its low cost and ease of use position it for growth

It used to be said you could get a Ford Model T in any color, as long as it was black. Until the ratification of the iSCSI standard in February 2003, SAN choices were much the same: You could get Fiber Channel or Fibre Channel. Although Fibre Channel storage still garners the lion's share of new large SAN installations, iSCSI is gaining acceptance. This technology now accounts for slightly more than 2 percent of total networked storage sales, but a recent report from storage analyst IDC shows a 22 percent increase in iSCSI sales revenue from Q4 2004 through Q1 2005.

The iSCSI SAN platform's lower cost and relative simplicity compared with those of FC have contributed to its growing popularity. An iSCSI SAN module can be plugged into any Gigabit Ethernet environment, then switched and routed like conventional network traffic to provide expandable, block-level storage to any host without a specialized HBA (host bus adapter). In most cases, setting up an iSCSI SAN takes only minutes and requires no extensive training or specialized personnel.



Furthermore, iSCSI initiators are available for Apple Macintosh OS X, Linux, Microsoft Windows, Novell NetWare and Unix.

Modular iSCSI SANs offer an affordable entry point for small and midsize businesses that need networked storage consolidation. On an enterprise scale, iSCSI SANs can provide primary storage for second-tier, departmental and remote-office servers, as well as secondary storage for disk-to-disk-to-tape backups or remote replication for disaster recovery. For high-density Windows environments, iSCSI offers convenient boot-from-SAN capabilities for bladed servers or clusters. iSCSI SANs are even developing a following at universities and other campus environments where distributed, Ethernet storage pools can be extremely effective.

We tested four dedicated iSCSI SANs in our Green Bay, Wis., Real-World Labs®: EMC's Clariion CX300i, EqualLogic's PS200E, LeftHand Networks' LeftHand SAN with Network Storage Module (NSM) 150 and MPC Computers' DataFrame 420. We also sent invitations to Aberdeen, which declined because its product is still in development, and Raid, which turned us down because of prior commitments. Relatively few vendors have committed to building dedicated iSCSI SANs. Instead, many are easing their way into this market by adding support for virtual, block-level iSCSI volumes to their NAS (network-attached storage) products. Although these solutions offer a level of convenience for environments that require a mix of file- and block-level storage, we chose to maintain our focus on SAN-specific products dedicated to providing block-level storage.

Three of those products—the EqualLogic, LeftHand

Networks and MPC—are independent SAN modules that can be virtualized into a dynamically scalable storage environment and targeted in the sub-100-TB range. This building-block premise takes advantage of the power of multiple storage controllers and linear bandwidth increases through the clustering of multiple modules and their active Gigabit Ethernet ports.

The EMC, by contrast, uses the conventional SAN model: A centralized storage controller manages multiple dependent disk arrays. Although this architecture reduces the cost of additional disk arrays, the CX300i supports a maximum of 60 drives and tops out at 19.2 TB. Additional drive cabinets are linked with dual FC loops, but because the number of available Gigabit Ethernet ports stays the same even when you add drive enclosures, beware of the potential for bandwidth bottlenecks as the system hits its maximum capacity.

All four systems have features we consider basic to an iSCSI SAN environment: dynamic volume expansion, provisions for snapshots and replication, support for CHAP (Challenge Handshake Authentication Protocol), iSNS (Internet Storage Name Service) discovery and LUN (logical unit number) masking. In all cases, additional disk arrays can be added to the storage pool, managed from a centralized console and assigned without the interruption of production, though array performance will be reduced during reallocation. Every SAN we tested came fully loaded with drives in every slot, though each one lets you minimally configure it with three to five initial drives while offering a range of drive-size options for future growth. For more information, see our features chart at www.nwc.com/go/0905f1.jhtml

» Executive Summary: iSCSI SANs

At last, there's a real choice in the SAN arena. Once the exclusive domain of Fibre Channel, SAN setups are now available with the less expensive, easier-to-use iSCSI interface. The result is storage that can connect easily in any Gigabit Ethernet environment.

Still, iSCSI SANs are facing the

standard catch-22: Few vendors want to jump in with both feet until the technology proves to be a hit, and that won't happen until there's enough choice. As a result, many vendors are adding block-level iSCSI support to their NAS (network-attached storage) lines, but few have introduced dedicated

iSCSI SANs.

We found just four vendors with iSCSI SAN modules to test: EMC, EqualLogic, LeftHand Networks and MPC Computers. The EqualLogic PS200E's combination of great features, winning performance and easy setup made it a shoe-in for Editor's Choice.

In our real-world performance tests, the EqualLogic dominated the competition, with MPC and LeftHand Networks following at a distant second and third. For reasons EMC and NETWORK COMPUTING couldn't resolve at deadline, the CX300i had severe problems with Iometer read tests larger than 64 KB and produced questionable results from the matching write tests, even though smaller transfers worked as expected. Because all the other systems performed these tests under identical conditions and with repeatable results, we had no choice but to give EMC a low score in performance.

We compared the features, performance, management, scalability and price (both list and cost per terabyte) of all four systems. It was a relatively easy call to present our Editor's Choice award to EqualLogic. The PS200E Storage Array has noteworthy ease of use, and its performance ran circles around the rest. LeftHand Networks' NSM and MPC's DataFrame share common software, but the DataFrame has the lowest cost per terabyte, while LeftHand's product offers growth through smaller modules. EMC's Clariion has the widest matrix of configuration options and would be a great fit for environments with an installed base of EMC gear, but it tops the list in price as it unexplainably bottomed out in performance.

EqualLogic PS200E Storage Array EqualLogic under-

A- stands iSCSI. It offers a killer combination of storage-optimized controllers, efficient internal data architecture and a simplified control interface that makes storage management incredibly simple. Command-line diehards who like to control their environment drive by drive may not embrace EqualLogic's simplified option set, but for



everyone else, it's hard to dislike a system as fast and easy to use as the PS200E.



The PS200E is a completely independent, 3-U, 14-drive iSCSI SAN with the security of fully redundant, hot-swappable cooling, power supplies and controller/communications modules. Each controller contains a 64-bit RISC processor with 2 GB of battery-protected cache memory, dual serial console ports and three Gigabit Ethernet ports whose choice of copper or optical interfaces supports failover, bonding or an internal method of load-balancing across all three. The management software resides on a Compact Flash card mounted on each controller module, and the storage environment can be administered inline, locally or remotely over HTTP, SSH, SSL or telnet.

Our test setup required an initial serial connection, through which we configured the Gigabit Ethernet ports and network addresses, as well as assigned names to the storage device and iSCSI "management group." Each of the two PS200Es we tested came loaded with 14 400-GB, 7,200-rpm SATA drives, for 5.6 TB of raw storage. Upon creating the storage group, we could choose either RAID 10 for optimum performance or RAID 50 for maximum storage.

EqualLogic's initial storage was easier to configure than any other test unit. Disk allocation is handled at group level

































REAL-WORLD LABS® REPORT CARD

iSCSI MODULAR SANs

	EqualLogic PS200E Storage Array	MPC Computers DataFrame 420	LeftHand Networks NSM 150	EMC Clariion CX300i
FEATURES (20%)	5	4	4	4
MANAGEMENT (20%)	4	3.5	4	3.5
PERFORMANCE (20%)	5	3.5	3	2
PRICE (20%)	4	5	3.5	3.5
SCALABILITY (20%)	4	4	3.5	4
TOTAL SCORE (100%)	4.40	4.00	3.60	3.40
<small>A≥4.3, B≥3.5, C≥2.5, D≥1.5, F<1.5 A-C GRADES INCLUDE + OR - IN THEIR RANGES. TOTAL SCORES AND WEIGHTED SCORES ARE BASED ON A SCALE OF 0-5.</small>	A-	B+	B-	C+

Customize the results of this report card using the Interactive Report Card®, a Java applet, at www.storagepipelinemag.com.

iSCSI MODULAR SAN PERFORMANCE

		IOPS
512 b IOPS Read		
EMC*		1,721.7
EqualLogic		44,468.4
LeftHand Networks		8,862.3
MPC		9,545.4
NWC Custom		
EMC*		22.1
EqualLogic		41.1
LeftHand Networks		18.1
MPC		27.5
8-Kbps Read		
EMC*		45.0
EqualLogic		51.7
LeftHand Networks		20.1
MPC		27.3
8-Kbps Write		
EMC*		37.7
EqualLogic		38.1
LeftHand Networks		16.1
MPC		14.1
64-Kbps Read		
EMC*		4.5
EqualLogic		75.7
LeftHand Networks		42.9
MPC		46.0
64-Kbps Write		
EMC*		73.7
EqualLogic		46.8
LeftHand Networks		26.8
MPC		22.6
1-Mbps Read		
EMC*		10.6
EqualLogic		123.0
LeftHand Networks		44.0
MPC		53.6
1-Mbps Write		
EMC*		112.0
EqualLogic		48.7
LeftHand Networks		28.8
MPC		37.4

*Because of unresolved problems during performance tests, the results listed for EMC's CX300i may be in question. These results were repeatable and are published here, as recorded, using the same test configuration for all systems.

automatically, and drives are assigned to arrays. Hot spares are specified based on the storage group's initial RAID configuration. Using this method, any new PS200Es added to the group automatically inherit the existing members' configuration. This sophisticated approach lets group members dynamically re-provision storage across disk resources, seamlessly load balance for efficiency, eliminate hot spots and even remove entire group members by automatically vacating the data to other members if space allows.

Base configuration for both PS200E units took only a few minutes. After assigning both modules to the management group, we logged in to the Web management interface to create volumes for our test servers. A simple three-step wizard helped allocate the volume, set iSCSI access and verify our choices to create the volume.

The PS200E simply blew the competition away. The SAN's read/write statistics exceeded the other competitors' scores by as much

“The PS200E simply blew the competition away”

as 2-to-1, and the 44,468 IOPS score we recorded on our 512-byte read test was not only through the roof, but consistently repeatable.

The management interface provides excellent graphical information on group status and storage allocation, and a huge amount of statistical information on physical member status. In addition, the PS200E offers the snapshot and remote replication features you'd expect from the most sophisticated SAN environments, without any additional software or licensing costs. Remote management groups can be configured as replication partners, with one-way, two-way (reciprocal) or centralized replication, in which multiple groups can replicate to a specified group.

The PS200E offers data-center-class redundancy, impressive performance and high-end software capabilities that are hard to beat. Although each unit priced at the high end of the systems tested, our calculated price per terabyte came in at \$9,250, with all software included. Only MPC's DataFrame came in with a lower cost per terabyte.

» PS200E (5.6TB), \$51,800, EqualLogic, (888) 579-9762, (603) 579-9762 ext. 7792. www.equallogic.com

