Tom's Hardware:

iSCSI The Open-E Way

Storage companies have maintained the same thing now for almost three years: Direct attached storage (DAS) is inflexible and creates bottlenecks, while its administration is time consuming. The future of professional storage, they say, belongs to networked appliances that allow placing storage units anywhere within your network. Switching to iSCSI will thus take storage to the next level, they say.

So if iSCSI really is that good, why has it not taken off yet? In addition to the high price of Fibre Channel infrastructures that are required for high-performance applications, technical barriers remain. A look at iSCSI in detail and its evolution accounts for the protocol's limited adoption.

iSCSI is a shortcut for Internet SCSI, in which SCSI stands for the small computer system interface that is the predominant DAS storage interface in the professional space.

The iSCSI protocol is used to encapsulate native storage data into IP packages, which in turn allows for transferring this data over existing networks as if these were a local storage interface such as UltraSCSI. Due to the large variety of IP networks (PAN, WLAN, LAN by means of Ethernet or Fibre Channel, WAN, MAN and the Internet), a storage area network (SAN) based on the iSCSI protocol can technically bridge any distance, and it is only limited by the performance of the particular network(s) it uses.

Knowing this, it becomes obvious that a SAN is typically kept within the limits of a fast network neighborhood, which brings us back to the barriers we mentioned: Storage devices that respond to multiple client requests require adequate bandwidth to do so. A 100 Mbit Ethernet network, for example, is technically capable of hosting iSCSI, but its performance would not really be up to task. Fibre channel, on the opposite end of the spectrum, often is too expensive for small- and medium-sized businesses.

The rapid deployment of gigabit Ethernet (GbE), however, gradually creates a substantial backbone for iSCSI applications. Although GbE over fiber cable is the state-of-the-art solution, we estimate that GbE over copper wire networks will further accelerate the transition from direct attached storage to storage area network applications; this type of infrastructure is also backwards compatibility and is available at an attractive price point.

Transfer Performance

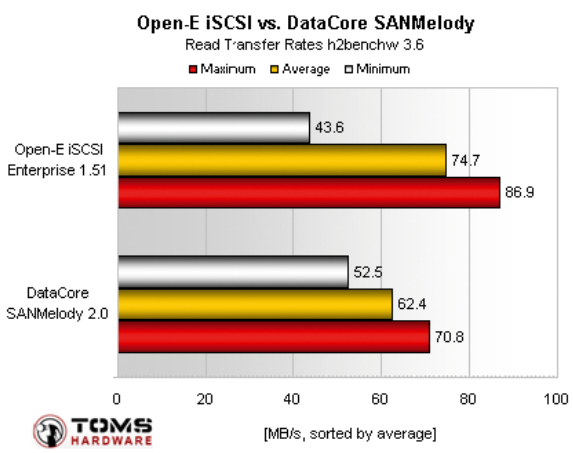

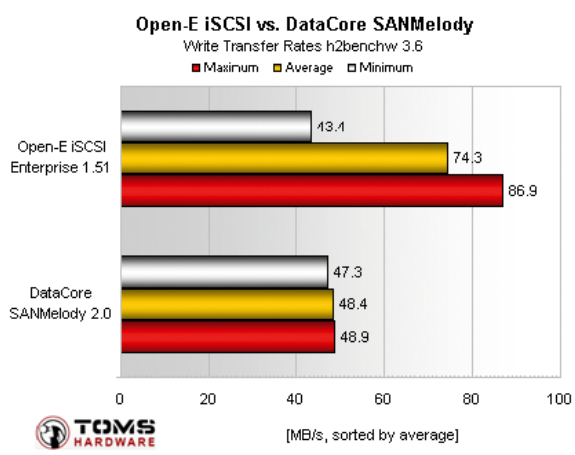

Open-E's data transfer performance is clearly superior, although the DataCore solution is capable of sustaining a much better minimum transfer rate.

Between Fibre Channel And Ethernet

The concept behind iSCSI arose from traditional storage interfaces such as ATA or SCSI being limited to a single machine and a maximum cable length. Flexibility of business and enterprise class storage applications hence is rather limited. So, ideally, storage should be detached from particular servers in order to increase flexibility and to avoid bottlenecks. In the end, administrators should be able to add, move, backup, restore or reconfigure storage without having to change running systems or look after multiple servers. Swapping out storage by using existing network infrastructure thus seems to be a good approach.

It is easy to spend several hundred dollars for SCSI cables or, even worse, a lot of money for Fibre Channel cabling and accessories. Fibre channel became the premier interface for professional applications and SANs because of its high throughput and long operating range. Fibre channel connections can be up to 30 meters in length when twisted pair copper cables are used, or up to 10 kilometers when based on a fiber optics link, with data transfer rates of 2 Gb/s or 4 Gb/s. Fibre channel is thus capable of connecting systems throughout a huge corporate campus for which very high data transfer rates and maximum performance are required.

In addition, Fibre Channel can point-to-point-connect two devices, be set up using switches or be operated in an arbitrated loop similar to Token Ring. Apart from that, it can carry ATM cells, IP packages or SCSI commands by fitting these into its own frames (which are not compatible with Ethernet). For these reasons, Fibre Channel is dominant in the enterprise space.

Fibre Channel storage components are considered the only choice for high-end environments such as data processing or huge database applications. The infrastructure can be used as an interface between hard drives and host bus adapters, or for linking systems and/or storage units within a SAN - or both. However, as soon as an enterprise does not require data to be locally available at a maximum speed, the extra costs do not make sense: Why not opt instead for storage subsystems based on Serial ATA components, and use the existing Gigabit Ethernet infrastructure?

Storage Is No Different Than Printing

The direction that professional storage has taken is similar to that of printers. Previously, at home, we connected printers directly to our computers via a parallel or USB port. In corporations, they were positioned in places where it physically made the most sense. Printers, for example, could be found next to employees that had to print frequently, or attached to computers dedicated only for printing. Then, the introduction of printers with network interfaces such as HP's LaserJets with the JetDirect interface option allowed for a more flexible positioning and administration. Printers became independent from their host computers.

Storage is headed in the same direction. Now, most users install more hard drives when they run out of capacity. However, concerns about where data is stored, available capacity or whether a file server is down or not can be less of an issue.

Indeed, even small businesses should consider moving away from the concept of operating file servers in favor of network attached storage (NAS) devices or flexible storage systems within a storage area network (SAN). These networked storage applications allow for increased flexibility, scalability, availability; backup options are increased and where the storage server is physically located is much less of an issue.

One downside is the more complex interface: While we believe that even a beginner can set up Open-E iSCSI, you should have a certain degree of experience before playing around with SANmelody.

iSCSI Targets And Initiators

iSCSI accommodates a server and multiple clients. The server is called a target and typically is either a dedicated storage computer, or a server that has access to direct attached storage. The storage management software can be part of the storage system (this is the case with Open-E's iSCSI) or is simply installed onto a host system. We used DataCore's SANmelody iSCSI target software to review the Open-E iSCSI system.

The clients are called iSCSI initiators. Provided that you have a user account that has access to the desired iSCSI target, you can simply use the Microsoft iSCSI Initiator: it can be downloaded on Microsoft's Website and is free of charge. It uses CHAP (Challenge Handshake Authentication Protocol) to log onto an iSCSI target. Simply provide the iSCSI target IP and your login data, and you are ready to go. After login, the iSCSI target will appear as a system drive to Windows, because the iSCSI initiator behaves as if it were a storage adapter (which it effectively is). Although the iSCSI target may be far away, you have the same access as you would with a local hard drive.

Test Setup

System Hardware I (for DataCore SANmelody 2.0)	
Processor(s)	2x Intel Xeon Processor (Nocona core) 3.6 GHz, FSB800, 1 MB L2 Cache
Platform	Asus NCL-DS (Socket 604) Intel E7520 Chipset, BIOS 1005
RAM	Corsair CM72DD512AR-400 (DDR2-400 ECC, reg.) 2x 512 MB, CL3-3-3-10 Timings
System Hard Drive	Western Digital Caviar WD1200JB 120 GB, 7,200 rpm, 8 MB Cache, UltraATA/100
Mass Storage Controller(s)	Intel 82801EB UltraATA/100 Controller (ICH5) Adaptec AIC-7902 Ultra320
Networking	Broadcom BCM5721 On-Board Gigabit Ethernet NIC
Graphics Card	On-Board Graphics ATI RageXL, 8 MB

System Hardware II (for Open-E iSCSI Enterprise 1.51)	
Processor(s)	Intel Pentium4 560 Processor (Prescott core) 3.6 GHz, FSB800, 1 MB L2 Cache
Platform	Asus P5MT (Socket 775) Rev. 1.04 Intel E7230 Chipset, BIOS 1001
RAM	Corsair CM2X512A-5400UL (XMS5400 V1.2) 2x 512 MB - DDR2-667 (266 MHz, CL4)
System Hard Drive	Samsung SpinPoint P120 P620JJ 120 GB, 7,200 rpm, 8 MB Cache, SATA/300
Mass Storage Controller(s)	Intel 82801UltraATA/100 Controller (ICH7R) AMCC 3Ware 9550SX
Networking	Broadcom BCM5721 PCI-E Gb LAN
Graphics Card	On-Board Graphics ATI RageXL, 8 MB

System Hardware	
Performance Measurements	c't h2benchw 3.6
I/O Performance	IOMeter 2003.05.10 Fileserver-Benchmark Webserver-Benchmark Database-Benchmark Workstation-Benchmark

System Software & Drivers	
OS	Microsoft Windows Server 2003 Standard Edition, Service Pack 1
Platform Driver	Intel Matrix Storage Manager 5.1.0.1022
Graphics Driver	Default Windows Graphics Driver

Benchmark Results

Data Transfer Diagrams

Open-E iSCSI Enterprise Edition 1.51

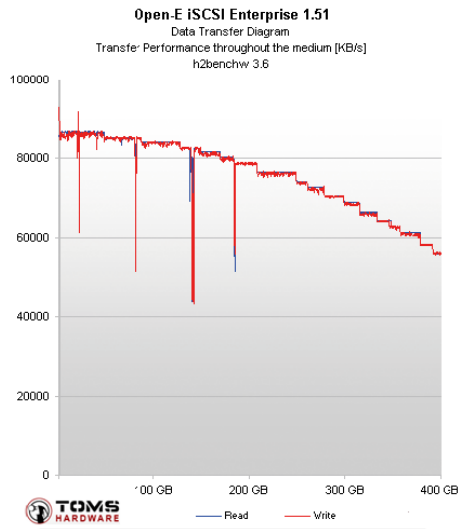

Open-E's iSCSI Enterprise does not seem to be limited when it comes to data transfer performance. Almost 90 MB/s is close to the net maximum that a Gigabit Ethernet is capable of transferring using iSCSI.

DataCore SANmelody 2.0

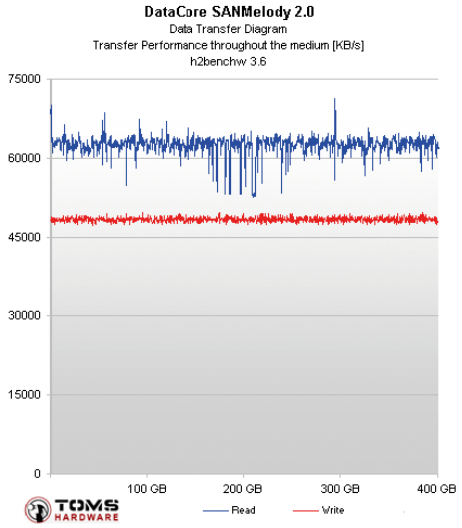

The data transfer performance of the DataCore SANmelody solution seems to be somewhat limited, as it does not exceed 70 MB/s.

Access Time

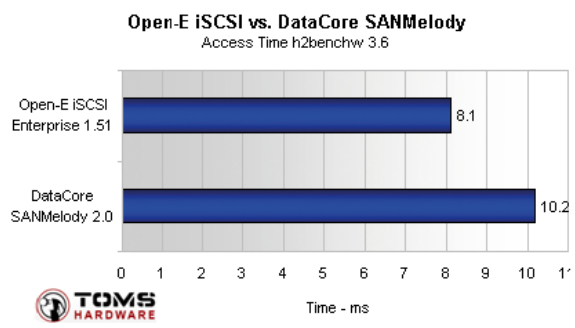

Here, Open-E's iSCSI Enterprise clearly provides the better results. Average access times of 8.1 ms versus 10.2 ms make quite a difference. At this point, you may wonder how such a quick access time is possible, as the 7,200 RPM drives that we are using perform in the area of 12-14 ms. It is the iSCSI solutions' cache algorithms that allow for shortening access times here. While this can sometimes be an advantage, it certainly does not reflect the real average access time of the iSCSI target.

Bandwidth

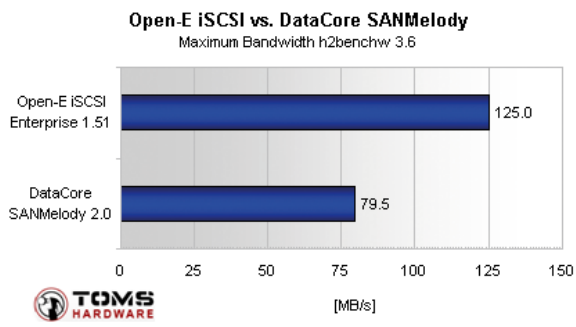

These results basically reflect the maximum bandwidth between the iSCSI target and our initiator system. As you can see, the bandwidth hits the limits of Gigabit Ethernet with Open-E iSCSI Enterprise (125 MB/s)

PCMark05 Windows Startup Performance

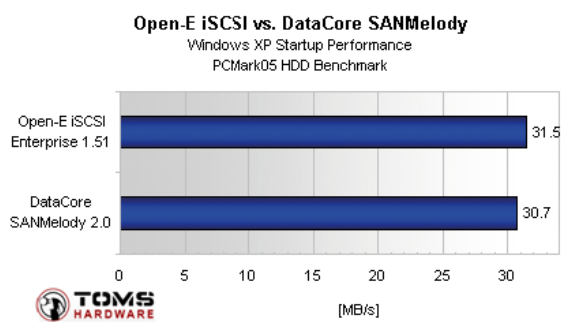

Although we have not seen a solution that would allow booting a system via an iSCSI drive, we still decided to include this benchmark. With this application-based benchmark, there is not much of a difference between DataCore and Open-E. Open-E has a very small lead here.

PCMark05 File Write Performance

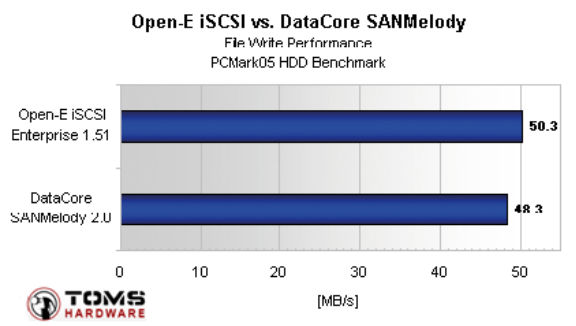

Again, Open-E outperforms its competitor by a very small margin.

Conclusion

While DataCore has been in the iSCSI arena since 1998 (longer than Open-E), both solutions make a very good impression. SANmelody by DataCore is a full-fledged software solution for Windows Server. It is scalable, feature rich and impressively powerful, but costs a fortune if you need more than the basic package. At the same time it should be underlined that SANmelody was intended for enterprise customers.

Open-E's iSCSI modules come with their own system software, which means they have to compete with SANmelody's Windows-based Storage Services features and the massive worldwide driver support that Microsoft enjoys. However, an Open-E iSCSI can turn virtually any computer into a comprehensive iSCSI storage solution without requiring a Windows license, expensive hardware or complicated administration.

While it is the top version of the product family, Open-E's iSCSI Enterprise is only half the price of the basic SANmelody kit. The latter is clearly more sophisticated, but Open-E provides a better price-per-performance ratio for small- to medium-sized enterprises. Proof can also be found in the benchmark results: Open-E iSCSI dominates all disciplines except in I/O performance.

Feature Table

	DataCore SANmelody 2.0	Open-E iSCSI Enterprise 1.51
Solution	Software solution. Requires Windows XP or Windows Server and .net Framework	Linux based software on a Flash storage module, using UltraATA/100
Application	Dedicated or shared storage server	Dedicated storage server
Infrastructure	Ethernet, Fibre Channel	Ethernet, Fibre Channel
iSCSI Accelerator support, 10 GbE support	All hardware with Windows drivers	Chelsio 10 GbE
RAID support	All RAID hardware with Windows drivers, Windows software RAID	Most cards from AMCC/3Ware, Adaptec, ICP Vortex, Intel, LSI Logic. Software RAID support.
Network card support	All hardware with Windows drivers	Multiple cards
Fibre Channel HBA support	All hardware with Windows drivers	Emulex, Qlogic
Multilane support / load balancing	Microsoft Multipath I/O	yes
Data security features	IPsec, Secure LUN allocation, iSCSI CHAP, asynchronous IP replication	IPsec, iSCSI target passwords, iSCSI CHAP
Fault tolerance	Auto failover option based on SANmelody disk servers	Secondary HBA takeover
Snapshot support	Complete image, image update, source update. Can be triggered by Windows Volume Shadow Copy	Complete image with scheduling
Storage volume features	Based on Windows Server 2003 Storage Services. Virtual Disks, Volume Shadow Copies, Multipath I/O Auto provisioning (dynamic allocation of storage)	Multiple volume grouping, online volume expansion
Performance features	I/O read and write caching using system RAM	no information found
Administration	Microsoft Management Console, rights controlled by OS	SSL web interface, multiple access levels
UPS support	yes, based on OS and UPS	Yes, with SNMP network shutdown
Email notification	yes	yes
Logging	yes	yes
Multi processor support	1, 2, or 4 (depending on software package)	up to eight logical processors
Hardware Requirements	x86 PC, 300+ MHz, 512 MB RAM, 65 MB disk space, network card	x86 PC, 1.4+ GHz recommended, 512 MB RAM, RAID controller recommended, network card
Software Requirements	Windows XP Home/Pro or Windows Server 2003 with .net Framework	None

ARTICLE by: Patrick Schmid, Achim Roos

Tom's Hardware: iSCSI The Open-E Way

Open-E GmbH (Headquarters), Lindberghstrasse 5, D-82178 Puchheim, Germany,
T: +49 (89) 800777-0, F: +49 (89) 800777-17, info@open-e.com, http://www.open-e.com