Storage Area Networking

Planning on a storage topology to meet current and future needs requires a review of your critical storage requirements — capacity, performance, security, segmentation and latency. CORKY SEEBER from Small Tree Communications looks at the options and explains the differences.

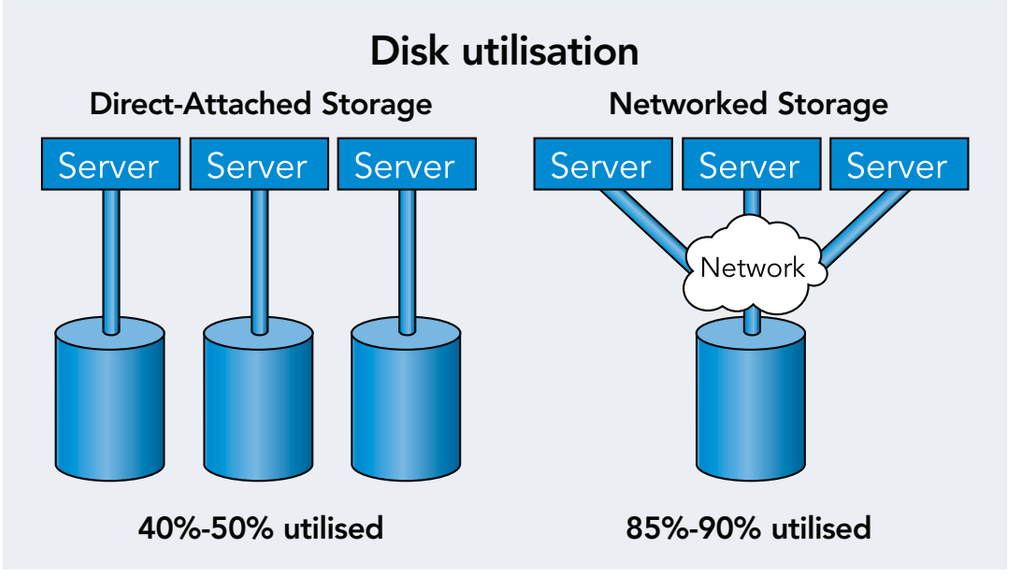

THE GENERALLY ACCEPTED definition of Storage Area Network (SAN) is a network designed to attach computer storage devices such as disk array controllers and tape libraries to servers. SANs can be viewed with an inward perspective of a computer system with a communications path to transfer data between the computer system to the attached storage elements, or with an outward perspective of multiple computer systems transferring data to a storage system across a network.

The distinguishing feature of a SAN and other forms of network storage is the low-level access method that is used. SANs employ formats that are very similar to those found in internal disk drive designs such as ATA or SCSI. This is done to enable the SAN to have higher performance by allowing the server to issue data requests for specific blocks or segments of data from the required disk drive with no intermediate translations resulting in much faster completion of requests. In the non-SAN environment of standard file storage requests, applications such as Common Internet File System (CIFS) or Network File System (NFS), the requesting system will issue a request for a file, which is a component of a larger file system and is often managed by an intermediary system. The intermediary system, acting like a traffic cop, determines the physical location of the requested file from its internal disk drives and sends the complete file across the network. The SAN configured data is able to be sent or retrieved much more quickly than a standard file storage solution because it can accurately request the exact information needed in the disk drive's native format without an intermediary system being required.

Most SANs use the SCSI protocol for communication between the requesting computer system and the storage elements, but they do not use the SCSI low-level physical interface. Typically, SAN's physical interfaces use some performance level of the Fibre Channel Host Bus Adapter (HBA), 1Gbit, 2Gbit or the new 4Gbit port interconnect connected to Fibre Channel Disk Drives for the superior performance and generally improved reliability over SCSI disk drives. The SCSI protocol information is sent via a mapping layer to the Fibre Channel drives, and most SANs in production today use some form of SCSI protocol systems mapped to Fibre Channel drives (using the FCP mapping standards). A well-designed SAN will use infrastructure specifically optimised for typical storage data patterns — sequential I-O to the disk in 4 kilo byte (kb) blocks, which improves latency by minimising the number of missed revolutions of the spinning Disk Drive in the storage system. The combinations of these features allow SAN users to realise additional benefits beyond simple storage performance.

Being able to communicate precise requirements for data across the network enables simplified storage administration by removing the need to cable the needed data from storage device to the requesting system, reducing system to storage distance requirements, eliminating the need to connect and disconnect cables (and the associated cables failures from the increased handling) and the improved workflow speed. SAN storage is still a one-to-one relationship between the storage device (or Logical Unit) and a single computer. The requesting computer is referred to as the initiator and the storage device is assigned a unique number, called the LUN for Logical Unit Number. The exception to this one-to-one relationship is when a SAN is configured with a clustered file system, which is an additional software application allowing multiple systems access to the same LUN concurrently. SANs provide higher overall storage use than standard file systems, which slows the requirement to provide additional storage capacity for the systems being supporting. The SAN can be configured to boot connected systems, allowing faulty systems to be configured off the network and allowing any attached LUN to be reached by the remaining systems on the SAN while repairs are made to the faulty system.

Perhaps the biggest advantage of SANs after performance considerations is data protection. SANs, by their nature of being a network, allow the ability to

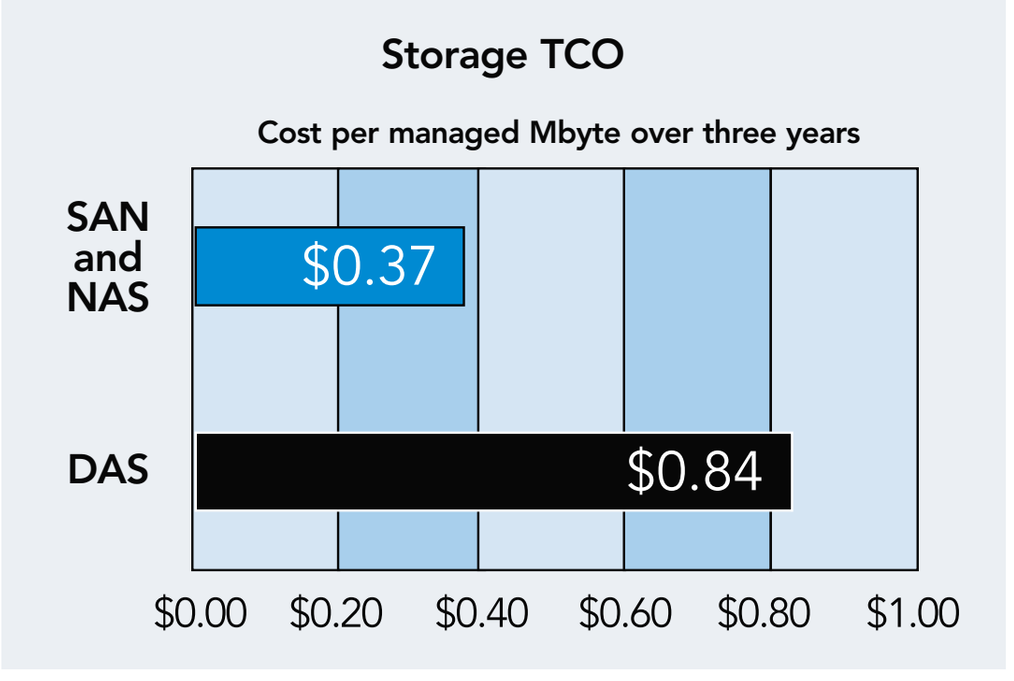

have copies of critical data in geographically distanced locations, providing protection from natural disasters and power outages and increasing security. SAN adoption providing this additional data protection has seen a dramatic increase since the September 11 attacks. There are additional methods to protect your data with SANs or to allow preventative or corrective maintenance, referred to as cloning, mirroring or snapshotting a LUN, each method provides a different level of protection and has a different impact on the system performance to provide this support feature.

SANs provide many advantages over standard file system storage, including performance, flexibility, security, and overall lower total cost of ownership (TCO) by reducing administration support and making higher use of the available storage capacity. Early Fibre Channel based SAN fabrics that incorporated multiple vendors experienced incompatibilities between the various Fibre Channel implementations. The lower level storage protocols typically worked well, but the higher-level applications tended to have the compatibility problems as vendors looked for ways to provide product differentiation in their solutions. This incompatibility created additional costs to customers by forcing them to have detailed knowledge on their fabric design instead of the simple plug-and-go scaling process. The members of the Storage Networking Industry Association (SNIA) worked together at attacking this specific problem and greatly improved the compatibility problems within two years and virtually eliminated the problems noticed in the earlier designs.

There are several different types of SAN fabric available today, with by far the most common SAN solution being based on Fibre Channel. For most of the last 10 years, Fibre Channel has enjoyed a large performance advantage over most other fabric types, specifically 10BaseT and 100BaseT Ethernet. More recently, there have been a couple of other fabric-capable protocol standards developed, iSCSI (Internet SCSI), based on Ethernet TCP/IP and ATA-over-Ethernet (AoE). With Ethernet being limited to 1Gbit while FC was operating with 2Gbit links, these alternative protocols have not received much interest from the market to date. There are other esoteric fabrics used in the Supercomputer High Performance Computing areas such as InfiniBand, which provides an interface to the storage attached to the system. With the release of the 10Gbit Ethernet protocol standard, coupled with the lower switch infrastructure costs, the performance advantage that SAN with Fibre Channel fabrics enjoyed were challenged. Today, many are rethinking what fabric protocol they should be using for their next SAN deployment.

The alternative that is best positioned to challenge the dominant market position of Fibre Channel fabric based SANs is iSCSI. The iSCSI protocol was released in 2003, prior to the introduction of 10Gbit Ethernet. iSCSI fabrics, comprised of 1Gbit Ethernet, did not provide a performance advantage over Fibre Channel fabrics. When 10Gbit Ethernet was introduced, the initial cost of 10Gbit solutions were more that $5000 per port compared to $1000 per port solutions for Fibre Channel. The high cost combined with the fact that most computers were using the PCI-X I-O bus interface inside their systems, which limited the top end of a 10Gbit solution to 8.5Gbits, produced little market interest in an iSCSI storage fabric. With the recent introduction of copper-based CX4 10Gbit Network Interface Cards, 10Gbit Ethernet fabrics have dropped significantly, and today costs compare favourably with Fibre Channel, making a 10Gbit Ethernet or Fibre Channel fabric decision a toss up from the cost perspective. The decision points have moved on to things like overall performance, compatibility between vendors and TCO. 10Gbit Ethernet enjoys a distinct performance advantage over Fibre Channel that is expected to widen as time passes. New systems are being released with the newer higher performance I-O Bus, PCI Express, which allows 10Gbit Ethernet to deliver its full performance potential (20 Gbits). Fibre Channel HBA performance improvements have slowed dramatically in recent years –— it has taken nearly twice as long to develop and release the 4Gbit capable HBA than originally forecasted. The 10Gbit Ethernet NIC is likely to move to 40Gbit performance levels sooner than Fibre Channel will move to 8Gbit.

Another potential competitor to the Fibre Channel fabric is the ATA over Ethernet (AoE) protocol. This protocol was designed for accessing ATA storage devices over an Ethernet fabric. AoE is different from most all Ethernet-based protocols in that is does not rely on the standard network layers such as Internet Protocol (IP) or User Datagram Protocol (UPD). This enables the network to bypass levels of the system's operation software, greatly improving the latency of the data transfers. The AoE standard is a much less complicated design than the iSCSI standard; the AoE specification is a mere 8 pages, whereas the iSCSI specification is 257 pages long.

The other principle difference to consider between a Fibre Channel SAN and iSCSI-based SAN is that iSCSI runs over the standard network protocol Transmission Control Protocol/Internet Protocol (TCP/IP). This enables the storage to be configured as a Network Attach Storage (NAS) solution. This allows the storage system to reside directly on the network in the

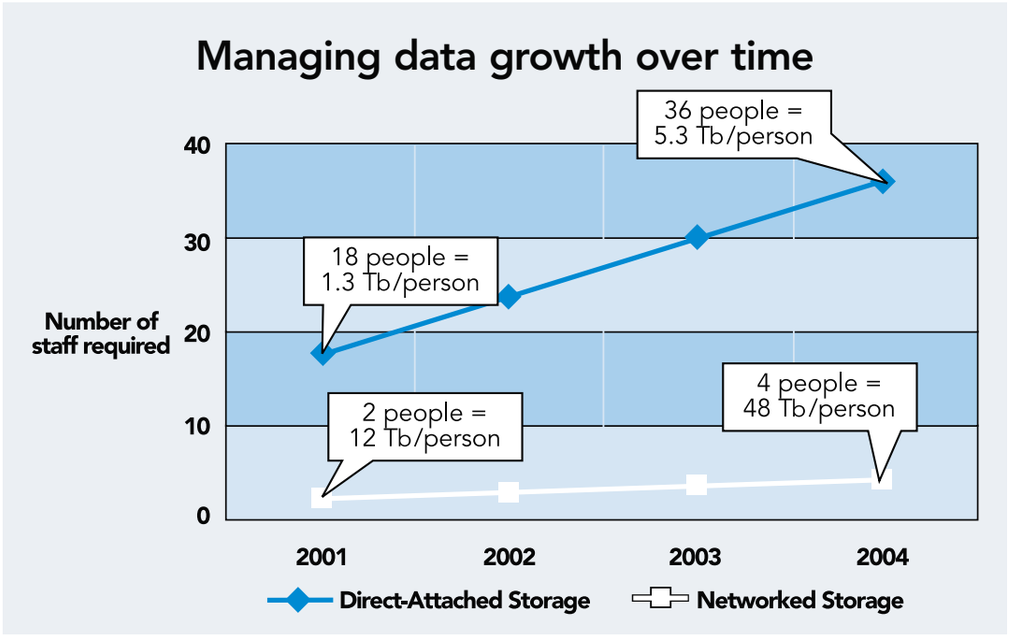

Worldwide factory revenue forecast for Disk Storage Systems by Achitecture

Networked Storage	2001	2002	2003	2004	2005	2006	CAGR 2001-2006
SAN-Attached Storage ($M)	$6,024	$6,594	$7,534	$8,902	$10,704	$13,424	17.40%
NAS ($M)	$1,618	$1,472	$1,681	$2,095	$2,649	$3,302	15.30%
Networked Storage Total ($M)	$7,642	$8,066	$9,215	$10,997	$13,353	$16,726	17.00%
Networked Storage Martket Share	31.9%	37.0%	43.4%	51.7%	60.5%	70.0%	- -
DAS							
DAS-Attached External RAID Storage ($M)	$7,228	$5,599	$4,970	$4,420	$4,040	$3,608	-13.00%
Host-Attached External RAID Storage ($M)	$2,388	$1,814	$1,345	$973	$694	$482	-27.40%
Host-Attached Internal RAID Storage ($M)	$6,687	$6,334	$5,691	$4,872	$3,971	$3,078	-14.40%
DAS Total ($M)	$16,304	$13,746	$12,006	$10,264	$8,705	$7,169	-15.20%
Total Market Share	68.1%	63.0%	56.6%	48.3%	39.5%	30.0%	- -
Total RAID Storage Market ($M)	$23,946	$21,811	$21,221	$21,261	$22,059	$23,895	0.0
Total Y/Y Disk Storage Market Revenue Growth	-18.6%	-8.9%	-2.7%	0.2%	3.7%	8.3%	- -

NAS configuration compared to residing behind one specific system in the SAN topology.

The benefits of NAS over SAN are that your storage access performance is not limited or bottlenecked by the single system that is configured in front of the specific LUN that you are trying to access, nor does it require any administration to get to the LUNs behind a failed system because all the storage is accessible to all the functional systems on the network. Additionally, in a NAS topology, it is easier to scale the size of your storage solution as all storage is attached directly to the network making more uniform access to all of the attached storage. The benefits of the Fibre Channel SAN topology over a NAS solution are the maturity of the protocol that translates into fewer setup problems and the superior performance for sub 10Gbit Ethernet fabrics.

In today's industries, the need for more storage and faster access continues to grow. The Small and Home Offices (SOHO) of less than 10 clients are increasingly looking at attaching their storage via their network to improve their workflow time. The days of sneaker-netting storage back and forth between systems are numbered even for small offices. Larger storage solutions are required to keep pace with the data intensive video and audio industries and to be competitive. To be effective, everyone involved with turning the raw work into the finished product has to have the quickest possible storage access to all the necessary data. Making your storage available to all your clients via a storage network allows the best use of the current storage capacity to hand, and delays costly upgrades. Planning ahead to ensure that you have selected the storage topology that will meet your current needs and that will have the ability to grow with you requires a thorough review of your critical storage requirements — capacity, performance, security, segmentation and latency. Typically, SANs are better suited for Enterprise-size solutions as they offer better segmentation and latency than their NAS counterparts, which are generally better suited for smaller sized configurations, but have naturally superior capacity and security characteristics.

No matter what storage topology you end up selecting, one thing is certain, to get the best use of it you need to make it accessible to all of your systems on the network and today you have options on how to configure that storage network to meet those requirements.