White Paper

XIOtech
A Seagate Company

Network-Attached Storage and XIOtech's Virtualized Storage Architecture

Proponents of SAN, NAS, and IP-storage transports often claim that just one networking model can serve the entire organization's storage needs. However, utilizing an integrated model can provide greater benefit than one transport infrastructure alone. Whatever the storage networking model, customers can achieve the greatest return on their enterprise storage investments if the storage foundation platform is independent of the network transport layers above. XIOtech's virtualized storage architecture provides a networked storage foundation that is fully independent of the transport layer, allowing organizations to reap the benefits of storage virtualization regardless of whether the network "plumbing" is SAN, NAS, or IP.

This paper is intended for CIOs, system administrators, and other IT professionals who are familiar with enterprise storage technology. It should be of particular interest to individuals who are investigating different storage network infrastructures from a SAN vs. NAS perspective.

April 2002

Eric Wendel
XIOtech manager of strategic technologies

A n I n f i n i t e l y B e t t e r S t o r a g e S t r a t e g y

IMPORTANT NOTICE

While we believe the information in this white paper is reasonably complete, we know that the needs of customers vary greatly and change rapidly. It is simply not possible to anticipate or address every situation or need that may arise. We provide this white paper “as is,” and reserve the right to change this white paper without notice. We also reserve the right to deviate from this white paper if we alone believe it is necessary to do so.

This white paper provides no warranty of any kind, whether considered alone or in addition to any product warranty. (Limited warranties for XIOtech products are stated in separate documentation accompanying or relating to each product.) No damages or remedy of any kind shall be recoverable by any party for any claim or loss in any way arising or alleged to arise from or as a result of this white paper, whether considered alone or in addition to any other claim.

The configuration or configurations tested or described in this white paper may not be the only available solution(s). This white paper is not intended (nor may it be construed) as an endorsement of any product(s) tested, as a determination of product quality or correctness, or as assurance of compliance with any legal, regulatory, or other requirements.

Introduction

Beginning in 2001, the concept of "NAS-SAN convergence" entered into the storage industry's marketing lexicon. What exactly is going on? Is it really possible for NAS and SAN to "converge" into a single network? In purely technical terms, the answer is no, but it is possible to build both NAS and SAN atop a single storage foundation. In this light, a more accurate description of what is going on is NAS-SAN integration, which is proving to be superior to other NAS implementations.

Whether an organization's storage networking strategy revolves around NAS, SAN, IP, or an integration of several technologies, it is important that the underlying storage platform be chosen carefully. XIOtech Corporation's MAGNITUDE™ provides a flexible and scalable storage foundation for NAS, SAN, IP storage, or whatever storage networking transports are used.

What about IP storage? Will this be the wave of the future, or will it find a more limited niche in the networked storage arena? Both NAS and iSCSI leverage IP networks as a storage transport, but what is the interplay between them? This white paper explains the differences among NAS devices, describes how XIOtech's MAGNITUDE works with all storage network "plumbing," and examines the realities of different storage models, including SAN, NAS, and IP storage.

Contents

What Is NAS?

Basically, network-attached storage (NAS) evolved to provide an easy and inexpensive way to connect distributed users to shared storage. In general, NAS devices utilize an embedded file-serving operating system (OS) and function as platform-independent file servers, but do not run applications or provide services such as file management. They are connected to the local area network (LAN) and support communication protocols such as Transmission Control Protocol/Internet Protocol (TCP/IP) and network file protocols such as Common Internet File System (CIFS) and Network File System (NFS).

However, NAS is not a technology in its own right; rather, it is an integration, or packaging, of technologies for a specific application. This is why NAS products are called *appliances*, which implies that they are designed for a particular use, as in a household or office device.

The chief goal of any appliance is to provide ease of use in a given *use context*. The downside, however, is that appliances do not move easily between different use contexts, nor do they scale well as the use context grows. For example, a 2-slice toaster integrates two technologies (heating element and thermostatic control) that make it effective and easy to use in a home kitchen. However, it is nearly useless—and in fact problematic—in a large restaurant. As toasters are added to meet the demand for toast, electrical outlets become overloaded and the mass of electrical cords makes using these appliances complex (and dangerous). The appliance that was designed to provide ease of use in one context has led to complexity and inefficiency as the context has grown. The user now has to throw away the toasters and find a better way to make large volumes of toast.

Generally, appliances that are easiest to use and most specialized are also the least scalable. This is true in the NAS market as well. The oft-heard complaint about NAS filers that are currently available is, “love your first, hate your third.” In order to address the scalability issues of either toasters or NAS, a new approach to technology integration must be found. With enterprise storage requirements increasing at nearly exponential rates, the use context for NAS appliances must include scalability. The toaster must be reinvented.

NAS Product Categories

It is important to understand the pieces that make up NAS before a new, scalable approach to NAS integration can be defined. The technologies that are integrated to form NAS appliances include—depending upon the degree of specialization—a network interface, file serving OS, compute platform (processor), RAID controller, storage enclosure, and disks. To understand how these pieces fit into products available today, it will be helpful to break NAS into product categories.

Three types of NAS products currently are available and differ mainly in the degree to which physical storage is integrated into the “box.” All three types include file systems, network protocol processing functionality, and integrated storage management services commonly associated with the concept of an “appliance.”

The following table describes the core functions associated with NAS appliances and the degree to which each of these functions is integrated in each NAS product category. File-server NAS functionality in both storage area network (SAN) and direct-attached storage (DAS) contexts also is included for comparison.

NAS Function or Component	NAS Filer or Head	DAS File Server	NAS Gateway	SAN-Attached File Server	Comments
Data Network Interface	Yes	Yes	Yes	Yes	TCP offload engine (TOE) may be used to increase throughput
File Serving, Network Protocol Processing Software	Yes	Yes	Yes	Yes	NAS appliances offer integrated support for multiple network file protocols; file servers offer this via additional software
RAID Controller	Yes (pro-prietary)	Yes (open)	No	No	NAS filers/heads typically support proprietary storage or JBOD, resulting in captive customer
Storage Backplane and Enclosure	Yes	Yes	No	No	NAS filers generally use proprietary storage to ensure customers remain captive
Physical Disks	Yes	Yes	No	No	NAS filers generally use proprietary storage to ensure customers remain captive

The key point here is the similarity between NAS filers and DAS file servers and, likewise, the similarity between NAS gateways and SAN-attached file servers.

NAS Filers

The NAS filer (Fig. 1) integrates all NAS functions into a single box. The box can be physical, as in low-end NAS devices, or the “box” can be a construct imposed on the customer through the use of proprietary interconnects between functional elements. While all three NAS product categories may be considered appliances, customers most often associate NAS filers with the appliance concept due to the high degree of integration and plug-and-play nature of these devices.

NAS filers in the current market use system-proprietary physical storage—whether included in the box, or sold in “shelves.” For the leading filer vendors, a substantial portion of revenues and profits come from sales of proprietary disk storage to a captive customer base. While it’s true that many enterprise-class storage vendors use vendor-specific disk modules, the difference between this model and the NAS filer proprietary storage model is that NAS storage also is *application proprietary*—the only application that can use filer storage directly is the filer application itself.

Today, the NAS filer is the only NAS configuration that spans the entire market space from small office/home office (SOHO) to high-end enterprises. NAS filers that currently are on the market include Network Appliance's NetApp® filers and Maxtor Corporation's MaxAttach™ NAS series.

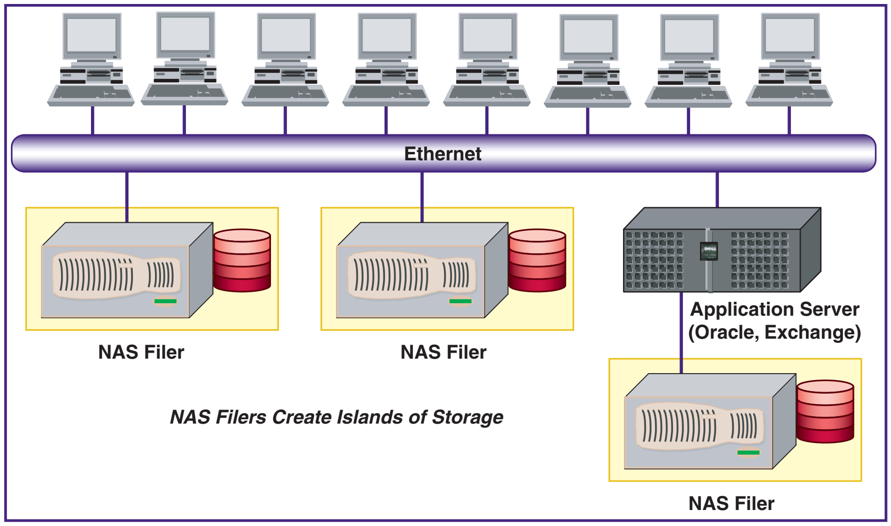

Fig. 1. NAS filers and heads follow the direct-attached storage model; this leads to "islands" of storage that must be administered individually

NAS Heads

The NAS head includes all the functions of the NAS filer, but excludes the storage backplane, disk enclosure, and physical disks.

An integrated RAID controller enables the vendor to employ mechanisms that enforce the use of proprietary storage, as in the case of the NAS filer. Some of these mechanisms include nonstandard sector size and format (Network Appliance) and special signatures placed in drive firmware or on track zero. Other NAS heads can use off-the-shelf disk enclosures and disk drives, but no leading NAS vendor currently offers this open design.

NAS Gateways

The newest NAS appliance type—and the most compelling in the enterprise space—is the NAS gateway (Fig. 2). In a NAS gateway, the proprietary RAID controller is replaced with an open, ANSI-standard SAN interface, usually a Fibre Channel host bus adapter (HBA).

From a systems architecture perspective, the NAS gateway looks exactly like an application server attached to SAN storage, so the NAS gateway also can be called a *server appliance*.

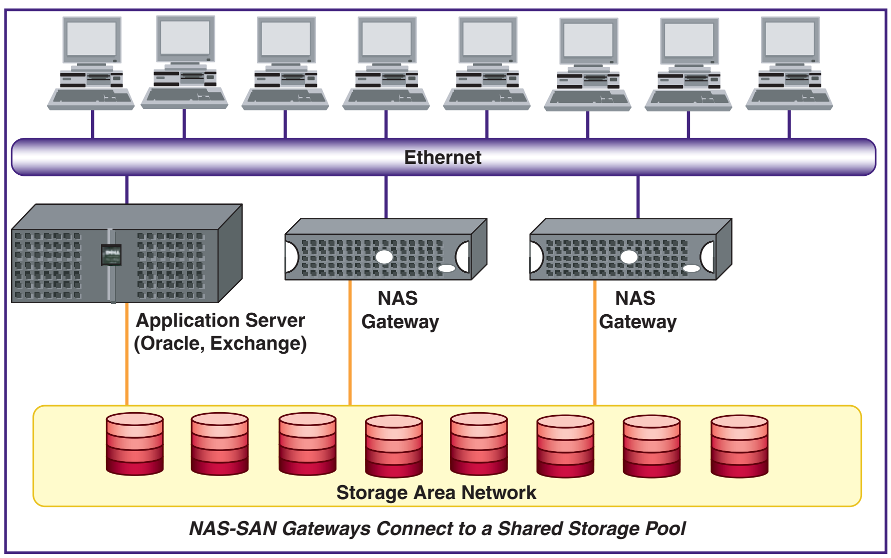

Fig. 2. NAS gateways allow attached users to access a shared storage pool, which is centrally managed and administered

The NAS gateway provides functionality to meet different users' needs. For example:

- From an end user's perspective, the NAS gateway looks just like a NAS filer
- From a network administrator's perspective, the NAS gateway looks exactly like a NAS filer
- From a storage manager's perspective, the NAS gateway offers ease of use and manageability that exceeds that of a NAS filer

In actuality, NAS gateway *functionality* has been available since the introduction of SANs, in the form of SAN-attached NFS and CIFS file servers—the NAS gateway provides added value, however, through *ease of use*.

The NAS gateway is differentiated from other NAS product categories by its ability to use shared, nonproprietary, or third-party storage. The NAS gateway provides users with the greatest flexibility for designing an overall storage architecture that meets enterprise requirements for both SAN and NAS. Because of this key advantage, proponents often position NAS gateways as a "best of both worlds" approach to NAS.

Because the NAS gateway relies on Fibre Channel-attached storage, installation is generally considered to be more complex and expensive than NAS filer implementations. But the costs associated with this additional complexity can be more than offset by the ability of a NAS gateway to take advantage of storage pooling and other return on investment (ROI) drivers associated with storage virtualization and SANs.

The NAS gateway provides a very compelling case in the mid-range and high-end enterprise space. All of the major server vendors introduced NAS gateway products in 2001, several of which are based on the Microsoft® Windows® 2000 Server Appliance Kit (SAK). In addition, VERITAS Software and Novell have launched "soft appliances." The soft appliance concept frees the customer even further, permitting a choice of server hardware platform to fit various performance and availability requirements, and allowing NAS performance to scale even within the server itself. The NAS gateway soft appliance enables customers to view NAS as an *application* that can run on their existing base of server platforms, rather than as a separate category of hardware device. This approach drives down total cost of ownership (TCO) by leveraging existing server investments, vendor discounts, and service arrangements.

When viewed in this simplified context, it's easy to see that, in functional terms, NAS and SAN never really diverged from one another in the first place, and that NAS is really an application—a file-serving application. The NAS application can be packaged as a soft appliance, a server appliance, or a filer, but in any case, NAS is an application that sits on top of block storage. Further, it is now apparent that only the filer implementations of NAS have become fenced off from the benefits of SAN. Thus, the concept of *NAS versus SAN* is about the packaging of technologies, not the benefits of the technologies themselves.

A Closer Look at NAS Filer Limitations

Islands of Storage, Islands of NAS

The primary weakness in the leading filer offerings is in scalability. In a NAS filer, the compute resources available in any given filer platform are essentially fixed, and cannot scale effectively as the amount *and usage* of the direct-attached storage "behind" the network processor scales upward. In business, it is rarely the case that the most critical data has a steady-state or static usage pattern, and so NAS filers are particularly problematic when the usage of a given segment of storage changes—as when reporting cycles or payroll processing is in progress.

Individual filers cannot share storage with other filers, so users of NAS filers inevitably discover that they either run out storage capacity before the compute resources are fully utilized or, more likely, run out of network processor bandwidth before disk space is used up. At the core of the problem is the fact that NAS filer appliances use a DAS model for storage.

These scalability problems are exactly what led to the need for SANs in the first place. It is easy to understand the scalability problems that storage planners face in the NAS filer context once the essential equivalence of the NAS filer and DAS file server is understood. There is quite literally no difference between attempts to scale a NAS filer "farm" and a DAS server farm. Either way, the storage planner winds up with islands of storage and must manage each filer's individual storage separately. At the same time, applications engineers must manage individual compute performance bottlenecks at the NAS-server level as application usage changes over time.

Finally, the NAS filer model requires that customers introduce yet another "box" into their mix of devices—another box that must be managed, upgraded, and maintained. While a single NAS filer may be simple to manage, the costs of managing filers increases dramatically as the number of filers increases.

And yet, storage must be made available to users via the data network, and that storage simply must scale. Fortunately, the answer is already here. The solution, as alluded to above, is the NAS gateway. By separating the NAS filer into its two major components—server and storage—the confinements of the filer, namely the server-DAS association, are eliminated.

Processing Power Cannot "Stuff" the Network

NAS filer scalability issues stem from the fundamental performance bottleneck of NAS—TCP stack processing. This is true even if the NAS device uses a TCP offload engine (TOE) on its network interface card (NIC) to redistribute the processing load between the main and offload processors. Simply adding more NICs doesn't work; each has its own IP address, which forces the user to further subdivide the storage inside the filer.

The industry has been hard pressed to solve the problem posed by TCP stack processing, but the ubiquitous presence of Ethernet in the market overall and the recent focus on Internet SCSI (iSCSI) has prompted massive investments in TCP offload accelerators by Intel, Alacritech, and others. Despite these efforts, full offload TOEs that enable Gigabit Ethernet (GbE) to run at line-speed throughput are only now becoming available. These new products are working at the very limits of current technology.

The trouble on the horizon is at the intersection of Moore's Law and 10GbE. In 1965, Gordon Moore, cofounder of Intel, stated that the number of transistors per square inch on integrated circuits would double every year through 1974. This formula, dubbed "Moore's Law," has been used to extrapolate processing performance gains over time as well.

But this equation began running out of steam (in terms of processing performance, if not transistor density) in 1994. According to NASA Advanced Supercomputing (NAS) research, the growth of reduced instruction set computer (RISC) and RISC symmetric multiprocessing (SMP) compute power between 1994 and 1998 fell to about one third of the growth that occurred between 1990 and 1994. It has continued to flatten further since then (Fig. 3).[1] In order to sustain Moore's Law of doubling speed (not density) every 18 months, new processing technologies and software optimization techniques have been required.

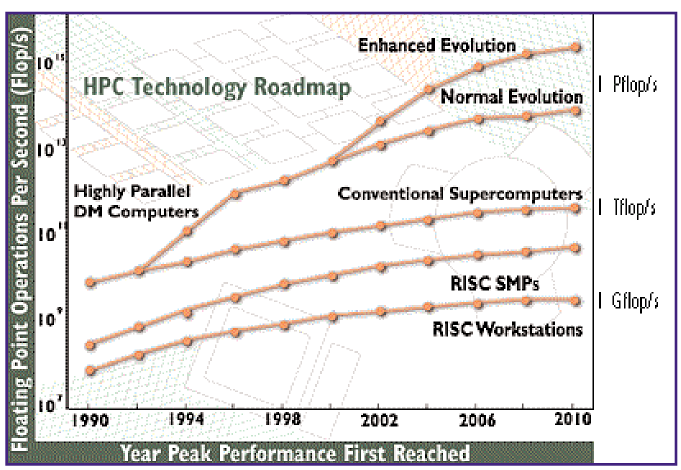

Fig. 3. NASA's high-performance computing technology roadmap.[1]

The TOE is an example of one of these new technologies, and it has finally enabled the processing of storage traffic over IP at GbE rates. However, the TCP stack processing power required to process data at 10GbE speeds is tenfold that required for GbE. Absent some quantum leapfrogging of Moore's Law, it will be more than four years *after* NAS filers using GbE TOEs reach full line speed that they will be able to "stuff" the 10GbE network.

1. "Petaflops: Supercomputing for the Next Century," *Currents*, vol. 2 (April 1997).

Meanwhile, even today's low- and mid-range, enterprise class SANs deliver more than enough bandwidth to fully saturate GbE. In 2002, Seagate Technology, IBM, and others will deliver disk drives that will be individually capable of saturating a GbE interface. In less than two years, mid-range storage subsystems will be capable of gigabyte-per-second throughput, enough to saturate 10GbE networks several years before this level of TCP processing power will be available.

Today, even high-end NAS front-end throughput performance significantly lags behind the throughput available from low-end SAN-attached storage devices. As TOE performance lags, NAS throughput performance will continue to fall further and further behind that of the storage pool.

In the history of the industry to date, storage interface speeds have always outpaced the ability of storage devices to "stuff" them. TCP-based storage transports, including both NAS and iSCSI, are the first ever to throttle the storage throughput available from disk.

One proposed solution to the latency, throughput, and processor utilization problems inherent in NAS involves the Direct-Access File System (DAFS) protocol in conjunction with new network transports based on Virtual Interface Architecture (VI). This new combination of lightweight protocol and low-latency transport permits data from storage devices to be copied directly into application memory and vice-versa.

From a technological standpoint, this approach clearly has the potential to solve NAS performance issues and perhaps scalability issues as well—but the costs of adopting these new technologies will outweigh their advantages until both DAFS and VI are widely accepted. Apart from the lack of broad acceptance of either DAFS or VI, customers should note that, while these technologies will run on virtually any physical network (i.e., Fibre Channel, GbE, Infiniband), DAFS/VI traffic will not coexist with IP traffic. Furthermore, the security risks of running direct memory access (DMA) transfers directly into application memory over GbE networks and the lack of security protocols to handle these risks mean this approach to NAS will require yet another discrete network in the data center for customers who wish to deploy it on an Ethernet infrastructure.[2]

While DAFS and VI may solve some NAS issues in the future, clearly some means of scaling the file-serving performance of NAS processors—independently of the capacity and bandwidth of the storage behind them—is necessary, and is necessary now. Thankfully, NAS gateways do just that, and they are available today.

How XIOtech's Virtualized Storage Architecture Fits into the Equation

Unlike other storage devices available today, XIOtech's virtualized storage architecture—composed of the MAGNITUDE hardware platform and REDI™ software family—provides a transport-independent networked storage foundation for NAS, SAN, IP, or just about any other type of storage network connectivity.

Since the first MAGNITUDE hardware platform was deployed, customers have attached "file-and-print" servers to it—first via point-to-point connections, next via Fibre Channel-arbitrated loop (FC-AL), and today via Fibre-Channel SAN fabrics and IP. For file sharing, users simply deployed either off-the-shelf OS modules or standard file systems software (e.g., plug-in modules for NFS or CIFS) on their existing file-and-print servers. When connected to a MAGNITUDE, these servers fill all the essential functions that NAS filers provide, including ease-of-use.

2. History suggests that many early adopters of DAFS and VI will be the same early adopters of large, Fibre-Channel SANs. The bounded and deterministic nature of Fibre-Channel SANs, mature and well understood pathing-zoning capabilities, and the low-latency of Fibre-Channel switches make Fibre-Channel SANs an ideal place for early-adopters to deploy VI and DAFS, which can readily coexist with other Fibre Channel protocols in the same SAN.

But XIOtech's virtualized storage architecture takes customers even further by allowing them to:

- Scale storage capacity independently of file-serving bandwidth requirements—with or without a SAN and regardless of network OS or file system
- Consolidate multiple file-and-print servers onto a single server
- Deliver the business benefits of storage virtualization to servers directly, through a SAN, via NAS, or through all communication models simultaneously
- Easily and quickly match file-serving performance with user needs for storage without having to purchase and engineer the storage on a server-by-server basis

MAGNITUDE users regularly report that they have little or no interest in deploying NAS in their environments. This is because many of the problems "solved" by NAS filers are problems that MAGNITUDE users simply do not have. The MAGNITUDE solves these problems at the storage foundation layer, and MAGNITUDE users have little interest in going backward to a DAS model.

NAS, SAN, IP Storage...It's All the Same to the MAGNITUDE

XIOtech's virtualized storage architecture was designed to deliver value by eliminating complexity and simplifying storage and server management.

XIOtech's MAGNITUDE hardware platform delivers the business value of virtualization fully independent of whatever network plumbing is most appropriate for each application or user. SAN, NAS, IP storage, or any combination of these different "pipes" (protocols, transports, and file systems) can serve alone or coexist with one another, on top of a MAGNITUDE-based networked storage foundation.

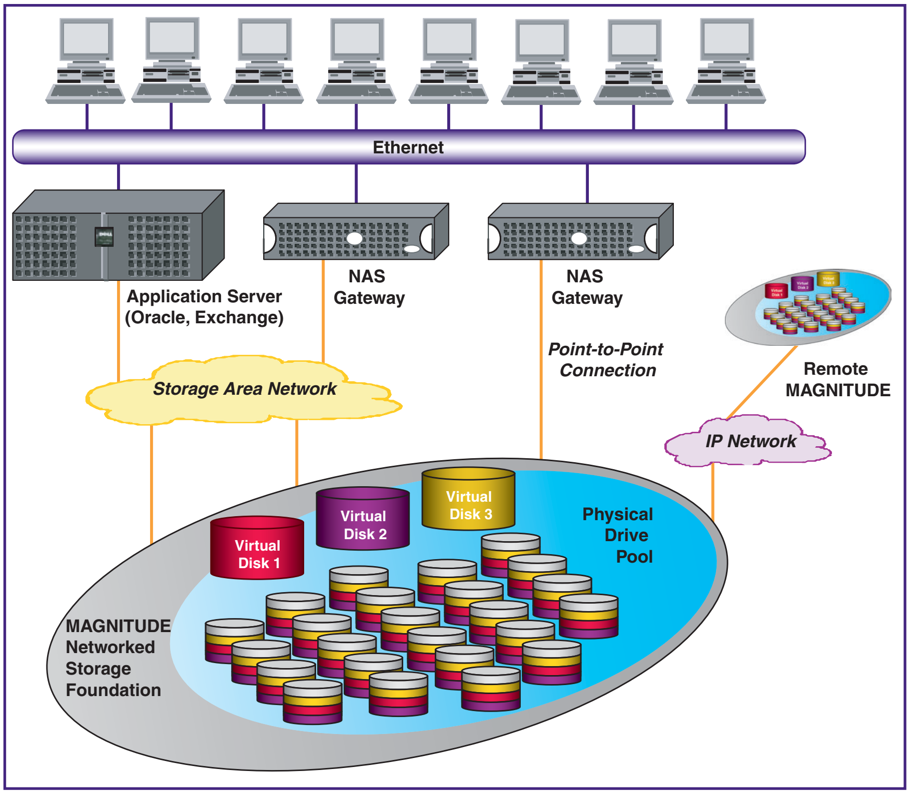

Fig. 4. XIOtech's MAGNITUDE provides a networked storage foundation regardless of network "plumbing."

Conclusion

NAS filers provide low-cost, file-level access to reasonably small amounts of storage, but they cannot scale to meet the performance and capacity demands of enterprise IT environments and applications. NAS gateways, when connected to a shared, centralized storage platform, are well-suited to the file serving demands of an enterprise. They allow storage and performance to be scaled independently, providing the foundation for a storage strategy that can adapt as an organization's needs evolve.

XIOtech's MAGNITUDE solves storage problems at the foundation level, and these problems stay solved as the storage networking infrastructure atop this foundation evolves through SAN and NAS, to IP storage, Infiniband, and whatever advancements lie beyond. The type of network(s) deployed atop the MAGNITUDE and the benefits associated with those infrastructures can evolve and scale independently, or asynchronously, from the storage foundation itself.

An investment in the MAGNITUDE hardware platform ensures that the customer will have a networked storage foundation that can integrate with new technologies and adapt to meet the organization's changing storage needs.

Appendix: High-Performance Storage for Enterprise Applications: The Domain of Fibre Channel Storage Area Networks

The performance of enterprise applications such as databases and email, serving hundreds or thousands of users, is measured in terms of the number of transactions per unit of time as well as transaction response time—the time it takes to complete each transaction from the user's perspective. The Transaction Processing Council's (www.tpc.org) well-known "tpmC" benchmarks are good reflections of this measurement. In the enterprise applications context, the performance required from storage is likewise defined in terms of input/output operations per second (IOPS) and response time. In general, the amount of data transferred in any given transaction or I/O is very small when compared to the size of the database from which that transaction is served. Therefore, the critical performance metrics for enterprise application storage are small-block IOPS and response time.

Investigating NAS/iSCSI Performance Claims

Unfortunately for customers trying to sort through storage networking performance claims, IP storage proponents—including the NAS and iSCSI camps—generally ignore IOPS and response time and instead focus on bandwidth measured in terms of megabytes per second.

In January 2002, a well-publicized startup company developing a TOE for both NAS and iSCSI announced that it had attained "wire-speed iSCSI throughput on a single Gigabit Ethernet connection." The following is from a Technical Note that accompanied the announcement.

> *Armed with these impressive test results, IT professionals now have the assurance that they can freely choose any combination of iSCSI and Fibre Channel interfaces for their servers and storage devices, while enjoying wire-speed performance in their storage networks.*[3]

In the announcement, "wire speed" throughput was defined only in terms of megabytes per second, so from an applications performance perspective, this statement is somewhat misleading. In the testing documentation, response time and small-block I/O performance were not tested, nor even mentioned. In addition, the test load appeared to be carefully selected to mask the impacts of TCP-induced latency on throughput.[4]

The performance claims of most NAS vendors have followed this pattern for years, and most iSCSI vendors seem to be following suit. The following quote from a September 2001 Gartner report echoes this concern:

> *Users should beware of TOE vendor over-promising. Most claims of "full line rate" or "full server offloads" are not true or assume ideal conditions.*[5]

Gartner released the report in response to the TOEs that were being announced beginning in mid-2001. While the report refers to iSCSI specifically, the TCP issue applies equally to NAS. The Gartner report continues:

> *...latency is usually a more important factor than bandwidth in application performance, and most TCP/IP vendors either ignore or deal poorly with latency issues.*[5]

Proponents of TCP/IP as a transport for storage (including both the NAS and the iSCSI camps) continue to ignore latency and response time because, like an 800-pound gorilla sitting quietly in your living room, there is not much that can be done about it. The issues are fundamental.

3. *Achieving Wire Speed iSCSI Performance*. Alacritech and Nishan Systems, 2002.
4. 64KB block size, constant queue depth of 8, Ethernet frame/packet size not disclosed
5. *TCP/IP Offload for IP Storage Networking: A New Necessity*. Gartner, 2001.

Feeling for the Floor

Faster processors and TOEs can solve TCP bandwidth and processor utilization issues by offloading the process, but this does not solve the fundamental latency issues associated with TCP—issues that are essentially "carved in stone." The Internet Engineering Task Force (IETF) and the entire Internet relies on TCP/IP sessions that are required to coexist with other traffic.

Essentially, TCP/IP sessions do this by "feeling" for bandwidth before using that bandwidth. Like a blindfolded person with a cane, who must feel for the floor, TCP sessions feel their way around other traffic. When a TCP session "bumps" into other traffic, it must stop or slow down and feel again. Even if there is no other traffic on the network, or if other traffic is at a lower priority, TCP sessions are blind and have to feel for network bandwidth anyway. This presents an unavoidable delay between the time bandwidth is required and the time it can be used. This problem cannot be solved short of a complete reengineering of the Internet itself, and quite simply, the SAN market will never be big enough to justify this.

Some NAS vendors have left their customers to discover these issues on their own. The recent "collision" between Microsoft's enterprise applications (Exchange 2000 Server and SQL Server 2000) and NAS devices is one example. While some NAS vendors have indicted that Microsoft's withdrawal of support for NAS is an attempt to establish its hegemony in file-serving operating systems, it should be noted that Microsoft also withdrew applications support for its own SAK-based NAS devices as well.

The TCP protocol is a fundamental enabling technology of the Internet. TCP enables hundreds to share access to a LAN, thousands to share a MAN or WAN, and millions to share the Internet. The price is paid in the latency impact on small data transfers. Does it make sense to place TCP/IP in the path between enterprise applications and storage? Can this question be answered before meaningful performance comparisons have been made? The latency overhead imposed by TCP processing—even at data-center distances—must be understood fully in terms of its impact on applications performance before NAS or iSCSI can be considered as a storage network for enterprise applications. Real-world collisions between NAS and enterprise applications underscore the market realities that already have resulted from placing TCP in the storage I/O path. Unwary NAS customers have been left to pay the price; unwary iSCSI customers may discover the same fate.

The Good News for IP Storage

Customers already are deploying IP-based storage transports where TCP is needed, but will not suffer application performance penalties where TCP in not needed.[6] Storage transports based on TCP/IP have found and will continue to find further applications in the storage networking market:

- NAS, in the form of the NAS gateway, will continue to replace general purpose DAS file-and-print servers due to lower TCO associated with ease of use and the ability to share storage via a SAN. The price/performance of NAS gateways will improve dramatically as TOEs mature.
- Microsoft and others will follow Novell's and VERITAS' lead in incorporating NAS functionality into server OS-based soft appliances. The server OS will, once again, become the NAS gateway.
- As TOEs mature, IP-storage transports including iSCSI will gain wide acceptance in data replication, backup/archiving, disaster recovery, and other applications where longer distance is required and network latency is relatively unimportant.

6. Products based on the iSNS, iFCP, and mFCP standards saw their first production implementations in the first half of 2001, nearly a year before the iSCSI specification is due to be finalized.

- In the post-9/11 world, disaster tolerance will replace fault tolerance as a primary requirement for high-availability storage infrastructures. To accomplish this, high-availability storage architectures must be distributed across wider and wider distances. The Fibre Channel-to-IP gateway device or switch will grow to be an important technology for *IP-enabled Fibre-Channel SANs*.

However, high-performance storage for enterprise applications is, and for the foreseeable future will remain, the domain of Fibre-Channel SAN-attached storage.

A Simple Strategy for Customers

Regardless of whether the pipes are NAS or iSCSI, any storage plumbing that puts a pair of TCP stacks between the application and the data cannot provide latency performance equivalent to Fibre-Channel SANs. This is not to say that the performance of every SAN implementation must be equal to that of Fibre Channel, but there is a performance penalty associated with TCP that must be recognized and dealt with.

For enterprise applications, performance and productivity go hand in hand. Productivity gained, or lost, by deploying any new technology is an important component of ROI. Therefore, the performance penalties of TCP must be measured in terms of application performance impact before customers can decide if the cost savings are worth the penalty.

Out of the clouds of hype, complexity, and uncertainty, there are a few simple facts that customers can use to make fail-safe decisions on strategic deployments of enterprise SAN, NAS, and IP storage technology.

- DAS is the common enemy that inhibits scalability, availability, and storage utilization.
- NAS is an application—a file serving application. Like other applications, NAS can be deployed using either a DAS model (NAS filer) or a SAN model (NAS gateway). All applications, including NAS, can benefit from improved storage utilization, uptime, and scalability via SANs.
- Many storage problems currently reside at the storage foundation layer...and must be solved at the storage foundation layer. Problems solved at the foundation layer tend to stay solved as use-context-appropriate storage network infrastructures evolve atop the foundation.
- The ROI associated with Fibre-Channel SANs overwhelms deployment costs, especially where the storage foundation has been well designed.
- Fibre-Channel SANs are safe investments for the near and long term, and prices are likely to come down.
- A SAN strategy does not preclude a NAS strategy. Enterprise SAN investments deliver even better ROI when they can be leveraged to serve NAS requirements as well.
- Fibre-Channel SANs will be the underlying transport for most NAS gateway deployments.
- TCP should be deployed in the storage transport layer only where it is needed.

XIOtech®
A Seagate Company

6455 Flying Cloud Drive
Eden Prairie, MN 55344-3305
phone: 952.983.3000 fax: 952.983.2320
www.xiotech.com