SAS, NAS, SAN

Past, present, and future.

IMEX
RESEARCH • COM

IMEX
RESEARCH◆COM

SAS, NAS, SAN — Past, present and future.

INSIDE:

IMEX Research
1474 Camino Robles Ct.
San Jose, CA 95120
Phone: (408) 268-0800
Fax: (408) 268-2300
anil@imexresearch.com
http://www.imexresearch.com

By **Anil Vasudeva**
President, IMEX Research

Reprinted from SMS Magazine

Since the advent of mainframes, computer scientists have constantly wrestled with various architectures to speed the I/O performance with increasing processor performance. Earlier efforts to improve data access involved tight coupling file systems and I/O with its operating systems. The rise of networked distributed computing brought the challenge of sharing files amongst heterogeneous computers running different operating systems. This gave rise to network-attached-storage servers to be independent of applications servers and dedicated to only serving files to users while offloading data management tasks from the over burdened application servers.

Faced with the lack of a practical technology that would interconnect theses server, the industry gave birth to a high speed fibre-channel technology which in turn provided the impetus for a third generation storage architecture called SAN (or Storage Area Networks) to emerge.

SANs create a dedicated network, focused on creating a universal any to any
connectivity between storage and server nodes - a network that combines the best of mainframe bus and channel's high speed and data integrity benefits with networks' distance benefits, a network that frees the main LAN network from backup duties that consume valuable bandwidth, a network that is scalable allowing increments in capacity without disruptions while leveraging the existing investments in legacy platforms and existing data, a network that provides centralized control while providing remote data vaulting for disaster recovery, a network that offloads storage management tasks from application servers and speeds up the entire network, thus allowing users the benefit of fast data access. SANs will eventually be at the core of every enterprise's data center, allowing companies to design centrally-managed data centers that embrace and interconnect farflung global SANs and provide service to all of their servers, no matter how far or no matter what operating systems they are running on. This new focus on data storage, as a key asset to manage, is obvious given the rise in dollars being spent on storage to the tune of 40-50% of total IT dollars in 1998. The rise in storage requirements, is being fueled by the birth of incessantly newer internet, data warehousing and ERP applications and further stoked by the lure of cheap disk drives at 5 cents per MB at the end-user level today.

"SANs will eventually be at the core of every enterprise's data center"

SAS – Server Attached Storage

Early mainframe storage designs took the premise that disk storage, which was cheaper than main memory, could be treated as a extended virtual memory to swap memory-pages. To achieve the fast data access, the data paths (or channels) between storage and processor were widened, the storage bus kept adjacent to the processor bus for data/sig nal integrity while boosting the channel speeds. Server attached storage architectures dominated the scene for several years from mainframe processor channels to PC Server bus slota and adapters.

One of the handicaps of the traditional server attached storage comes from the tight coupling between storage and the operating system. A general purpose SAS server performed a variety of tasks concurrently from running applications, manipulating databases, file/print serving, providing communications, checking data-integrity to many housekeeping functions. This meant that all data access requests from a client must continuously compete with these tasks continuously. As the number of users accessing the common centralized data storage increases, the file access takes a back seat to other tasks leading to slow response time for queries. For years one of the major jobs of MIS administrators was to keep the storage performance fine tuned to achieve a certain minimum level of user query response time.

Another limitation imposed by the SAS architecture was that of limited distance imposed by the interface - the OEMI wide parallel connections in mainframes and wide differential parallel SCSI connections in servers were limiting the distance between computers and servers to a few meters. This led to the creation of raised-floor data centers but posed a severe constraint and limitation on interconnectivity in multi-site operations. One of the major benefits of fibre channel connectivity that is not fully emphasized, is the removal of spaghetti of OEMI/SCSI wires interconnecting storage to servers and the associated improvement In reliability. This is over and above the advantage of allowing high-speed connectivity and increased distance between centrally managed data repositories and dispersed LAN servers.

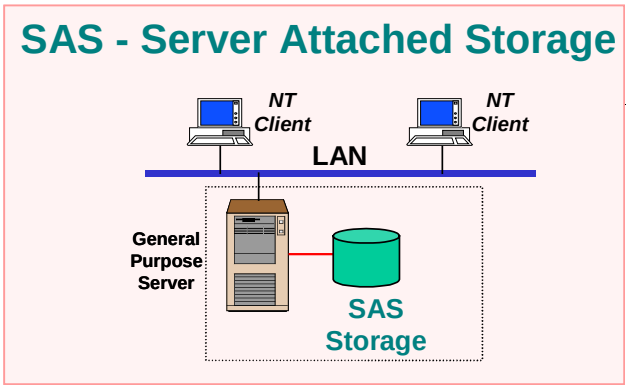

NAS – Network Attached Storage

Network Attached Storage, compared to server attached storage on the other hand is a dedicated file server optimized to do just one function only and do it well - file serving. NAS is a system independent, shareable storage that is connected directly to the network and is accessible directly by any number of heterogeneous clients or other servers. NAS file servers are essentially stripped down servers specifically designed for file serving and offloading file management services from the more expensive application servers.

In NAS, you can add storage at random without disrupting the network. When the storage was on the server as in SAS, the administrator had to take down the system, install or upgrade the drives and bring the system back up again. That created a lot of unacceptable downtime. NAS is being installed increasingly now to mitigate the downtime associated with SAS. NAS is making inroads into the marketplace at different price, performance and size levels. As business operations become more global and around the clock, more and more

applications become mission critical

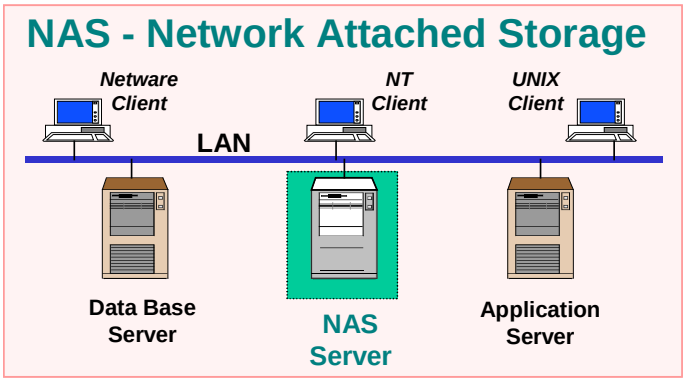

demanding 24x7 uptime. Feeding this frenzy of 24x7 uptime are the obiquitos internet using email messaging and around the clock customer information browsing demanding richer and richer content from text to images to audio/video clips, virtual private nets for e-commerce and data warehousing and ERP applications on the intranet.

NAS architectures generally sport a light proprietary OS kernel and file system able to operate autonomous of other applications and are thus devoid of all overhead from extraneous drivers prevalent in SAS architecture. The NAS operating system is fully compatible with server operating systems such as NT, Unix, Netware etc. Generally called a Network Appliance, NAS devices are relatively easy to set up turning painful storage upgrades into simple plug-and-play devices requiring no server downtime to set up. After plugging a NAS server onto a network and assigning an IP address, setting up control lists and user permissions and voila, all is done. This is because the NAS server boards integrate the Ethernet connection, the SCSI (or Fibre Channel) controller-to-disk connections, the operating system and boot up software all on one simple card. Much as NAS devices have built-in security features, administrators generally choose to rely on existing robust security features of their networks. One of the main benefits of NAS is that it allows clients to directly access data without burdening the application servers.

Factors motivating rise of NAS servers include:

Performance

Stored data supplied directly to clients without server intervention

Performance enhancements for a site can be achieved by dedicating each NAS server for its specific needs (e.g. Publishing/Prepress department can have its own file server dedicated to video/imaging. graphics data using RAID-3 while the e-commerce sales/order processing/ shipping/customer service groups could be running OLTP applications on its own dedicated server, running RAID-5 or RAID-1 attached to the same net.

Availability

- Fault Resiliency -Majority of data has becoming mission critical to run a business and so must be made secure and reliable
- Need for 99.9% availability (8 Hours per year of downtime). Some applications require even higher data availability such as 99.99% (1 hour of downtime per year) and recovery from failure from hardware, software and application switchover within 30 seconds.
- Ease of remote vaulting for data recovery

New architectures such as Wolfpack, the Windows NT 2-node Clustering provides high availability via server failover using MSCS software acting in an active/active mode. Thus when a server fails, the application is switched to the surviving server and so is the storage

Cost

HSM: migration to low cost tape for infrequently used data.

Scalability

Other benefits accruing to NAS architecture include modular scalability by direct attachment of add-on file servers directly to

the net without bringing down any applications running already.

Interoperability

NAS is very capable of supporting heterogeneous clients (such as NT and UNIX workstations) to share same data from network-attached server. Majority of mainframe storage today acts as a physical repository to store different types of file data such as UNIX/NFS or MVS or NT/SMB in different physical locations of the storage system. Some UNIX companies emulate the NT client data into UNIX NFS format and store as NFS data on NAS file server. Others notably, Net Appliance has a data format conversion facility to store NT or UNIX data in a common format, allowing faster data retrievals.

Data sharing and interoperability to serve heterogeneous clients

Manageability

NAS lends itself to dedicated storage management resident on the NAS servers itself to ensure efficient backups,

Challenges/Caveats

One of the major shortcomings of NAS storage architecture is that the network on which NAS runs is also used for data access by clients to retrieve data from the file server or communicate with application servers. The data movement between the disk and tape servers also goes over the same LAN. This creates a major network bottleneck when the number of users increases. Further the overhead of network stack contributes to higher data latency during server or client to storage communications.

While NAS works well for documents, file manipulations and transaction-based applications, it is not necessarily most advantageous for database applications because it is file-oriented. Also for high bandwidth video applications, NAS slows down since the shared network on NAS gets clogged fast with multiple large files and starts to become a bottleneck.

SAN – Storage Area Network

A SAN (Storage Area Network) is a dedicated high performance network to move data between heterogeneous servers and storage resources. Being a separate dedicated network it avoids any traffic conflict between clients and servers. A fibre channel based SAN combines the high performance of an I/O channel (IOPS and bandwidth) and the connectivity (distance) of a network.

To interconnect distributed systems over distance IT system administrators have been forced to use Fast Ethernet links, which are terribly inefficient because of large packet overhead (associated with small 1500 byte transmission packets) and high latency. In smaller computer room environments, short, thick and unwieldy spaghetti of SCSI wires or OEMI copper cables in mainframe environments to connect storage to servers are commonplace.

Adopting SAN technology through the use of Fibre Channel and hubs and switches allows high-speed server to storage, storage-to-storage or server-to-server connectivity using a separate network infrastructure mitigates problems associated with existing network connectivity. SANs have also the potential to allow cable lengths up to 500 meters today and unto 10 km in future so servers in different buildings can share external storage devices. And because the new emerging SAN/VIA (virtual interface architecture) interconnects have low latency and lesser overhead as compared to traditional LAN/WAN networks, they are ideally suited for clustering and mirroring/replication applications. The capability of connecting existing SCSI devices to SAN using SCSI to Fibre Channel bridges also preserves investments made in existing storage devices. This will help fuel growth of SAN infrastructures.

Performance

SAN enables concurrent access of disk or tape arrays by two or more servers at high speeds across fibre channel, providing much enhanced system performance.

Availability

SAN has disaster tolerance built in since data can be mirrored using FC SAN up to 10 km away.

Cost

Since SAN is an independent network, initial costs to set up the infrastructure would be higher but the potential exists for rapid cost erosion as SAN installed base increases.

Manageability

- Data centric
- Part of Server cluster
- Thin protocol for low latency
- DMA to server RAM - direct communication to Data

Future of SAN

- Embedded and Distributed File System
- Intelligent SAN-smart File System where portion of File System is in SAN
- Data routing,
- Storage network management
- Concurrent processing and manipulation of intelligent data streams

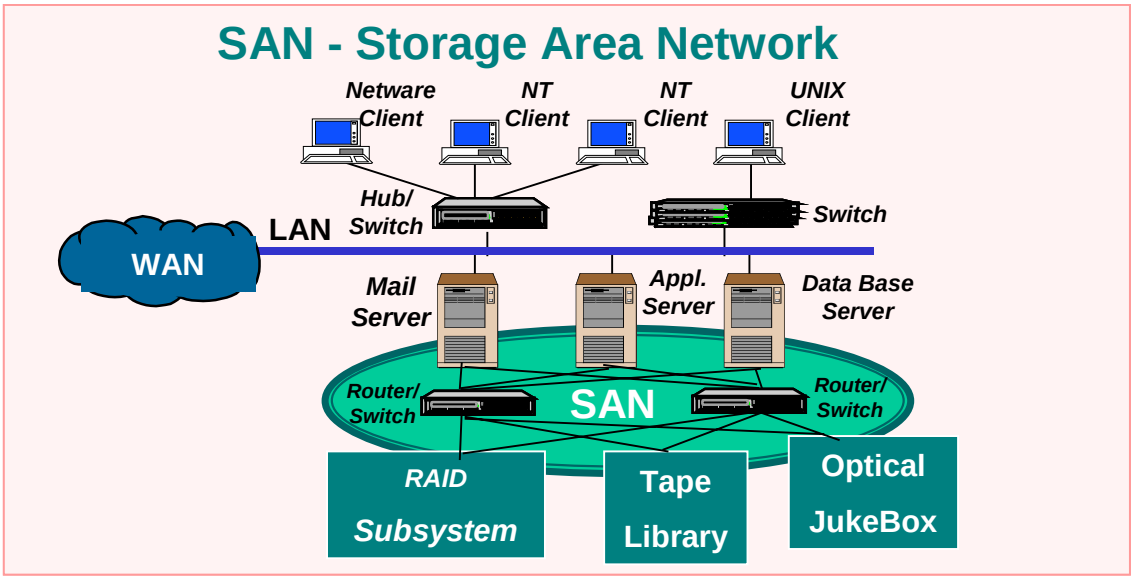

Scalability

Scalability is natural to SAN architecture, depending on the SAN network management tools used.

Interoperability

Like a LAN/WAN it can use a variety of technologies such as serial SCSI, ESCON, FICON, SSA, ATM, SONET etc. This allows easy relocation of backup data, restore operations, file migration and data replication between heterogeneous environments.

- Server Independent Storage Tasks
- Peer to Peer copying
- Peer to Peer backup
- Automatic back up using Fibre Channel
- Data Sharing, Data Formatting
- Security - Authorization, Authentication, Access Control
- SAN technology, in future, may also interconnect worldwide with other SAN intranet sites to provide instantaneous replication of corporate data to these remote sites to create a global information system. This would allow local access to fast while being up-to-date.

Challenges

As with all new technologies, SAN developments must rapidly happen in areas of data management, security features, interoperability test suites, availability of VI adapters to improve latency between interconnected servers and the availability of SCSI/Fibre Channel bridges.

The ability to manage SANs is as vital as having the speed and distance benefits of SANs unless the storage management features are built into the operating systems, customers end up buying them from server vendors or third parties who in turn license them to the server vendors. To simplify management, SAN vendors need to adopt SNMP and WBEM type standards to monitor, alert and manage data on SAN networks. Also the need for dynamic logical partitioning of different network operating systems being managed by the centralized console. Since there are a number of different devices from different vendors, the big challenge facing system administrators end up to make sure that they are interoperable and have one centralized management tools (such as HP Open View) and with which other management software packages are compatible with.

The lack of SAN optimized applications; management utilities, fault-tolerance features and full plug-and-play interoperability at this time are the caveats and cautions for administrators to use before plunging into adopting SANs.

I/O Performance SAS vs. NAS vs. SAN

In a distributed and networked environment, NAS allows better performance as measured by the response time of user queries, as the number of users increases. This is based on the independence of NAS from burdening other application servers and sporting a dedicated light OS to move files.

The SAN related curve shown in the graph relates to the potential benefit the SAN/Fibre Channel architectures will bring in the future.

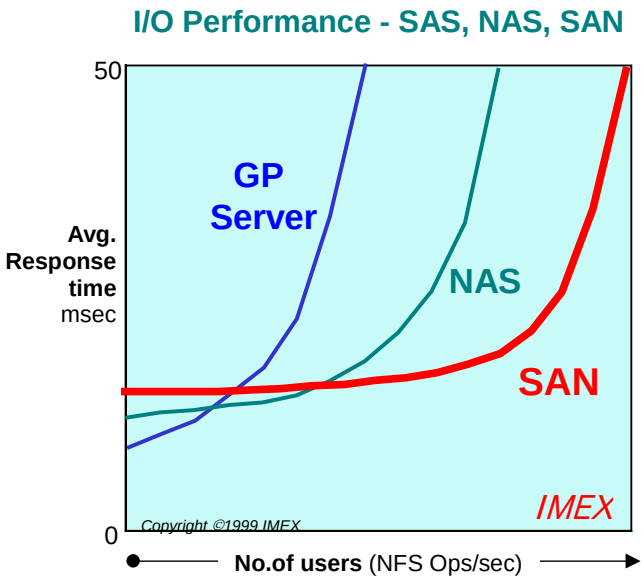

Figure – 4: Performance of SAS vs NAS vs SAN with increasing number of users in an OLTP environment.

NAS – its origin, present adoption and future incarnations

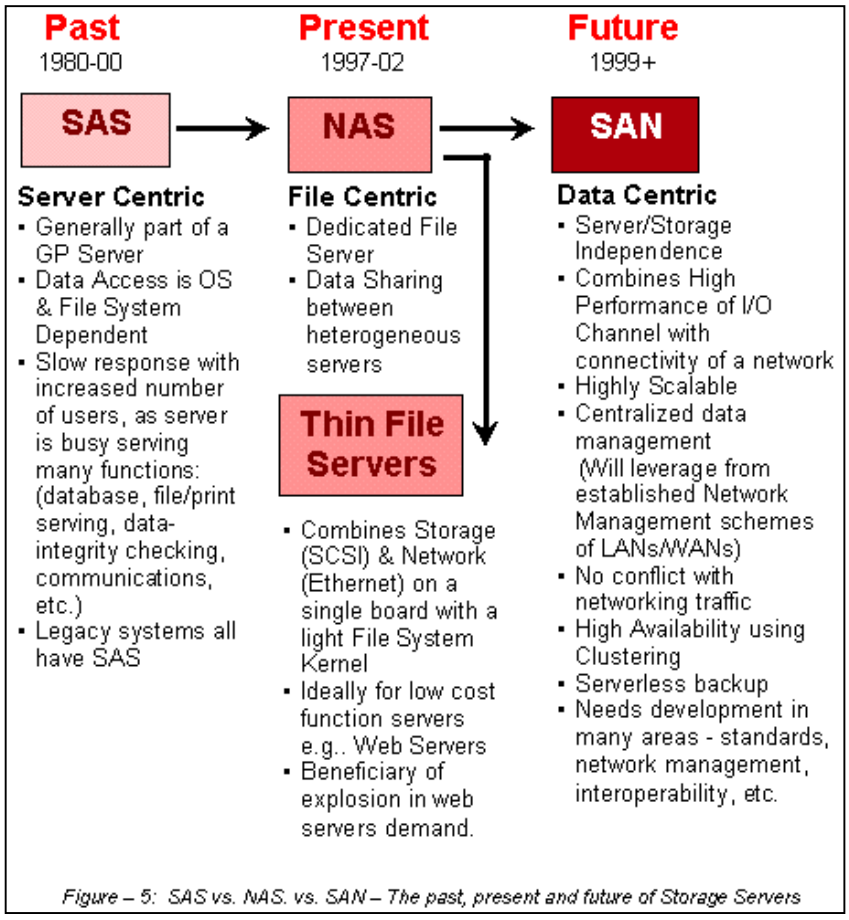

Figure – 5: SAS vs. NAS. vs. SAN – The past, present and future of Storage Servers

Figure 5 summarizes the comparisons between SAS, NAS and SANs. There is a wide proliferation of low cost NAS products to serve the exploding Web Server market.

Additional benefits of NAS for Web Servers accrue when using load balancing and web caching technologies.

Market Outlook for Storage Subsystems, SAS, NAS and SAN

The market for Storage Subsystems is concentrated in top 10 players. The top 3 players including Compaq (including DEC acquisition), IBM, and EMC alone control 50% of the market. Internal RAID is being led by Compaq, HP and Dell servers.

Network attached storage has started to materialize led by Network Appliance at the High end and Meridian Data/Snap products at the lower end.

SAN fibre channel products are increasingly being introduced in the market including fibre-channel adapters, hubs, switches and routers, SCSI/FC bridges, disk drives and testers from multiple sources.

Eyeing the advantages and popularity of FC/SANs the leading data networking vendors like 3COM have jumped in to embrace the technology. Not far behind would be major players like CISCO, Lucent and Nortel/Bay Networks eyeing the convergence of voice and data networks and associated centralized database servers for interactive voice response and e-commerce applications for the global marketplace.

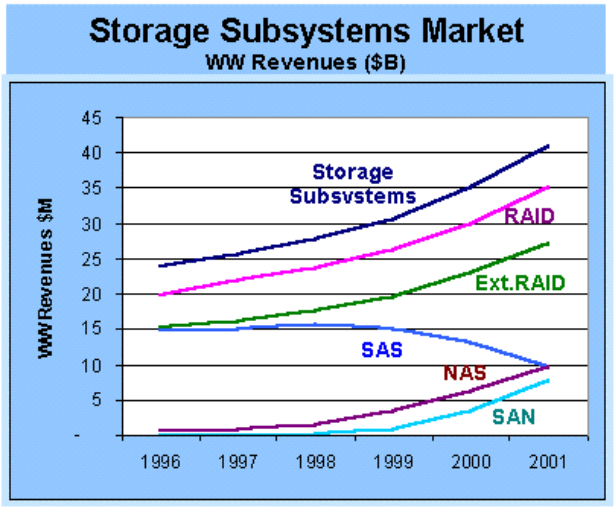

High Availability the next frontier for Computing Systems

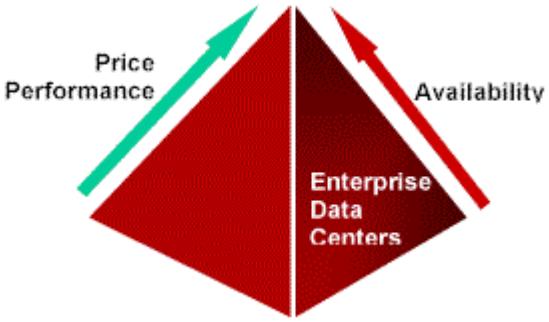

One of the key factors necessary for open systems to succeed in enterprise environments, despite inroads made in price/ performance areas, is high availability.

Given the ever increasing mission critical nature of almost every computing task, high availability ranks at the top of the list today for enterprise data center administrators.

ROI on High Availability

Achieving continuous availability has been the main mission of mainframe s/390 and proprietary computers. While the price/performance/availability afforded by UNIX computers have achieved remarkable success, NT clustering is being adopted by at least 20 major server manufacturers. Given their expertise in mainframe and UNIX type computing, it would not be long before the clustered NT solutions not only provide the performance but the cascaded high availability from cluster of multinode servers.

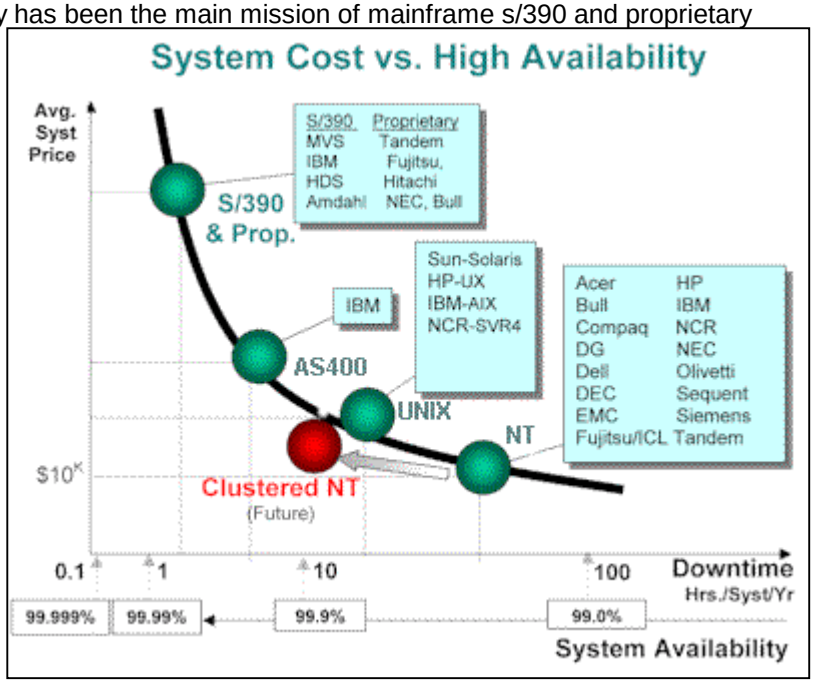

NAS and SAN architectures have a lot more steam than evident at this time.

While NAS will succeed more in the low end dedicated to providing plug and play solutions, SANs would be embraced strongly by the enterprise starting with multinode clusters running Oracle OPS and SQL type of parallel databases. It would proliferate to initially coexist and then displace SAS and NAS architectures. Its real strengths will come from centralized SAN connected to remote SANs in Global 2000 and Fortune 1000 enterprises doing global electronic commerce.

IMEX
Corporate Headquarters: 1474 Camino Robles San Jose, California 95120
Phone: (408) 268-0800 - Fax: (408) 268-2300 Home: http://www.imexresearch.com - E-mail: info@imexresearch.com