Validation Report

ESG Lab Report™

NetworkAppliance

iSCSI SAN

A validation study
by
ESG Lab

April 2004

Authors:
Tony Asaro Brian Garrett

Table of Contents

Executive Summary

Network Appliance (NetApp) is the leader in providing network attached storage (NAS) solutions and has been building fast, simple and reliable storage subsystems since 1992. The NetApp Filers are easy to manage, utilize standard and mature network protocols, support multiple operating systems, and are built on a highly reliable and proven platform. NetApp has developed its own file system and sophisticated data management software, the most notable being its snapshot capability.

NetApp has recently added support of the iSCSI protocol across all of their product lines so that nearly every Filer deployed in the field, and all systems being shipped today, can use the block-based SCSI protocol over existing commodity Ethernet networks. As of the writing of this report, NetApp provides their Filer-side iSCSI target license free of charge to customers. NetApp customers can try out and deploy iSCSI on a cost free basis.

The use of iSCSI extends the simplicity and reliability of NetApp Filers to support applications that traditionally use block-based storage (e.g. databases and email). iSCSI support should be particularly compelling to IT organizations that:

- Have concerns around the costs and complexities of deploying and managing Fibre Channel (FC) SANs.
- Haven't deployed FC SANs.
- Haven't extended their existing FC SANs to encompass all of their application servers.
- Are eval\uating a tiered storage strategy and migrating less important, least access data onto less expensive and easier to manage storage systems.

Why not just store everything using NAS? Why did NetApp bother to support iSCSI and block-based storage? NAS is effective for file sharing but isn't always the best solution for storing database or email data. NAS is an application that uses high-level protocols to transfer data while database and email applications use the lower-level block-based SCSI protocol. Storing data onto NAS adds an additional layer that results in simplifying the storage management process. However simplification does not come without a cost – and that cost is often a performance hit due to the high-level NAS protocol. By combining both NAS and block-based storage onto one platform enables the customer to choose which protocol is best-suited for their needs.

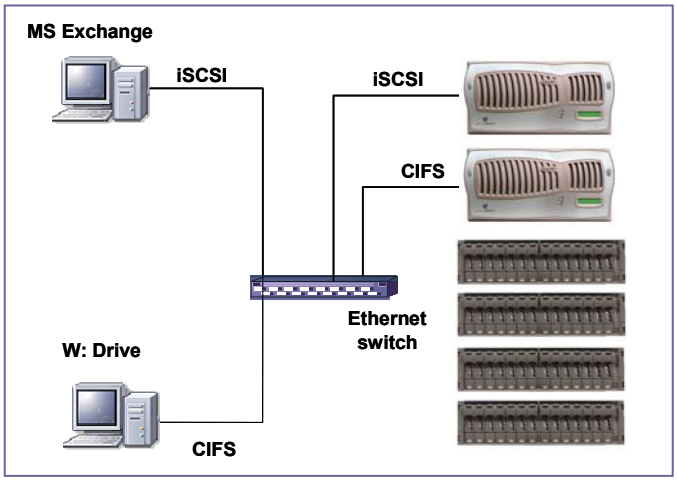

NetApp has combined CIFS, NFS, FCP SAN and iSCSI SAN within their Filer products in order to support all applications and data types within one platform. Using the same Filer, Ethernet network, and storage capacity, a customer can have world-class NAS for file sharing and iSCSI SAN support for applications such as databases and email, and manage it all as one system. NetApp has focused on ease of use and integrated management and is in a unique position, since their customers have been using NetApp products for IP storage for over a decade.

ESG Lab Test Results

ESG Lab spent two and a half days testing the iSCSI solution from NetApp at their RTP facility in North Carolina with the following results:

- ☑ Installing and configuring a NetApp Filer is quick and easy, whether it is used for file sharing or for iSCSI SAN.
- ☑ NetApp has simplified iSCSI configuration and management on servers running Microsoft operating systems by creating the SnapDrive plug-in GUI, which will have a familiar and intuitive look and feel to Windows administrators.
- ☑ Microsoft Exchange configurations can be easily created and managed using the NetApp wizard-driven SnapManager for Exchange plug-in.
- ☑ NetApp iSCSI support is being extended beyond NT and Linux to include Novell Netware and various UNIX operating systems. For example, ESG Lab installed and tested iSCSI devices on a server running HP-UX.
- ☑ ESG Lab conducted true "apples-to-apples" performance testing with the same server accessing pools of the same type of drives comparing iSCSI performance to the use of Direct Attached Storage (DAS). ESG Lab tests showed that iSCSI outperformed DAS with up to 41% faster throughput and 63% better response time using consolidated real-world and simulated OLTP workloads (details are available later in this report).
- ☑ Based on results obtained during simulated and real-world application workload testing, ESG Lab has determined that the performance of iSCSI is more than fast enough for departmental servers running e-mail and database applications.
- ☑ ESG Lab spoke with users who are using iSCSI for e-mail, database, and other applications. These customers were using NetApp Filers for NAS and are now supporting database and email applications with the same level of customer satisfaction.

Issues to Be Addressed

More SnapDrive Functionality

A number of popular operations can be performed with SnapDrive using the same GUI for both iSCSI and NAS. Some operations however require the use of FilerView, a second GUI. For example, ESG Lab needed to use FilerView to define a new RAID group. NetApp should consider consolidating all management functionality into a single management GUI.

HBA enabled iSCSI Boot

NetApp should continue to work with the standards committees and hardware vendors to enable users to boot from iSCSI devices. NetApp is testing and qualifying hardware adapters that support remote booting (a.k.a. iSCSI HBAs). ESG Lab planned to test this capability with a qualified iSCSI HBA but ran short on resources and time.

Wider UNIX support

NetApp is qualifying iSCSI initiator drivers that bring iSCSI capability to UNIX variants other than HP-UX and Linux, which are supported today. It is understood that this effort is not necessarily in NetApp's control, but with a heritage of providing solutions to the UNIX community and their dedication to iSCSI qualification, ESG Lab believes NetApp should continue to drive the proliferation of iSCSI into more UNIX environments.

ESG Lab's View[1]

iSCSI SANs are emerging as a competitive price/performance alternative to FC-based SANs. There are a large quantity of application and database servers on which companies are still using DAS, for two primary reasons – the cost and complexity of FC SANs. Customers now have an alternative to FC SANs by using iSCSI, their existing Ethernet infrastructure, and any one of the growing number of iSCSI storage systems.

NetApp is the only storage system that natively supports NAS and iSCSI today and is also one of the few established and leading storage companies providing an iSCSI storage system. NetApp has an excellent support staff, as well as world-class quality assurance, testing, and interoperability labs. A major and crucial benefit of NetApp supporting iSCSI SANs is that it leverages almost all of their 40,000 NAS Filers already being used in production environments. The customers have accepted NetApp as one of their strategic vendors, and iSCSI uses the same Ethernet infrastructure as NAS.

ESG Lab validation testing allowed us first hand experience with NetApp's implementation of iSCSI and we found the installation to be simple; management was easy, including taking snapshots, adding storage capacity online, and migrating data; and their iSCSI performance was impressive. iSCSI SANs are a compelling price/performance alternative to FC SANs and even more so when NetApp Filers are already installed and being used by customers for NAS. NetApp has a brilliant strategy that is advantageous to customers, NetApp themselves, and iSCSI adoption in general. Every current NetApp NAS customer and those evaluating NetApp for the first time should consider implementing Filers for both file sharing and deploying an iSCSI SAN. NetApp has leapfrogged its competition by supporting NAS and SAN (iSCSI and FC) with the same ease of use, reliability, and functionality that has made them one of the leading storage system companies in the market today.

[1] ESG Lab Disclaimer: It should be noted that ESG Lab did not use and is not using NetApp products in a production environment day-in and day-out. ESG Lab reports are not designed as a substitute for formal engineering testing by potential customers. Each customer should do his or her own evaluations. However, ESG Lab believes its tests provide a solid basis for the opinions expressed in this report.

ESG Lab Validation

The goal of the ESG Lab project with NetApp was to validate their iSCSI solution in terms of ease of use, software features, performance, and reliability. NetApp Filers can be configured to support NFS and CIFS for file sharing as well as Fibre Channel (FC) and iSCSI for block-based applications. NetApp refers to this as their "unified storage vision." This study focuses on the unified storage implementations of iSCSI and NAS.

The number of NetApp units shipped configured to support mixed protocol combinations is steadily increasing. High-end Filers are being shipped with FC support for databases and e-mail servers with NFS support for file services on UNIX platforms. **Mid-range and entry level Filers are often shipped with iSCSI support for use with Exchange and SQL server with CIFS for shared drive access under Windows 2000, NT and Windows 2003.**

NAS solutions communicate with application servers using a relatively high-level file-oriented protocol (e.g. Windows uses CIFS and UNIX uses NFS). SCSI, FC channel drives, and storage subsystems use the lower level block-based SCSI protocol. Block level protocols are typically used by databases like Oracle and SQL server and email applications, including MS Exchange.

A server that reads or writes to an iSCSI storage subsystem is called an initiator. Logical drives within an iSCSI storage subsystem are called iSCSI targets. ESG Lab performed initial testing using the following configuration with a Dell 2650 server as the iSCSI initiator, connected through a Dell GbE switch to a NetApp Filer configured with iSCSI targets.

Figure 1: NetApp Filer Supporting iSCSI Targets

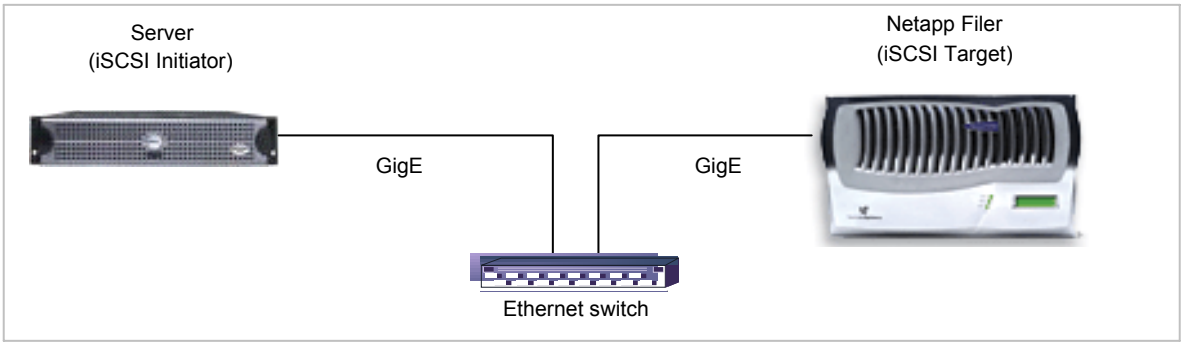

iSCSI SANs Need to be Simple

FC SANs are typically deployed and managed by experienced IT staff within corporate data centers. iSCSI SANs enable the same block level storage traffic as FC SANs but instead run over a "cheaper wire," a.k.a. Ethernet. NetApp realized that the "cheaper wire" of iSCSI will at first be mostly used for departmental servers and small/medium sized businesses that have limited resources. The conclusion was "iSCSI must be simple" for the first adopters of this technology that don't necessarily have SAN expertise.

ESG Lab began with a focus on testing the simplicity of the NetApp iSCSI solution for the following:

- Installation and initial configuration
- The SnapDrive Plug-in
- SnapManager for Exchange

Installation and Initial Configuration

One of the reasons NetApp has been so successful is that their NAS products are easy to install and manage. It takes less than an hour to have a NetApp Filer up and running for NAS services. ESG Lab performed an installation, wiring, and configuration of a Filer for iSCSI, and the experience was exactly the same as the process of installing, wiring and configuring a Filer supporting basic NAS services. The only additional step required was obtaining an iSCSI license key from the NetApp website (free for all licensed NetApp customers). The process is even easier if the customer already has an existing NetApp Filer being used in production for NAS. In the latter case, a customer can have iSCSI services up and running on a production Filer within an hour.

SnapDrive

ESG Lab began iSCSI testing on a Dell server running Microsoft Windows 2000. The NetApp SnapDrive plug-in was used to configure and manage NetApp iSCSI devices. SnapDrive is a native Microsoft Management Console (MMC) plug-in that can be launched with a right click from the My Computer "Manage" console, much like the Microsoft Disk Administrator and Disk Defragmenter utilities.

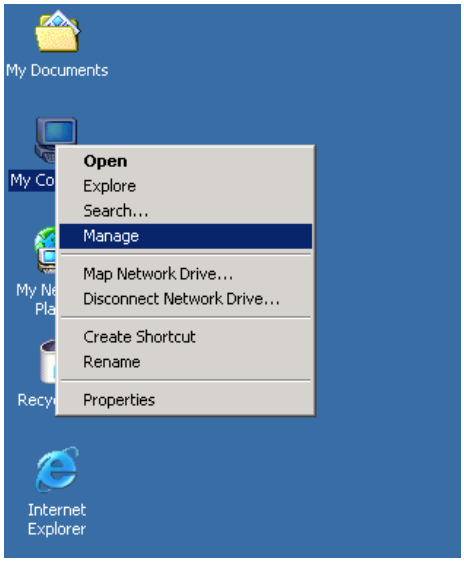

ESG Lab used SnapDrive to define iSCSI virtual disks that were formatted and presented to the Windows operating system as basic logical disks. Once created, the disks were managed with SnapDrive, including the ability to expand virtual disks on demand and instantly present the new disk space to Windows, all while data was being copied to the drive being expanded. NetApp has the unique capability of being able to recognize and use increased drive capacity online without having to run the additional command line utility (diskpart) that is needed for other vendors' storage systems.

Why This Matters

Storage is becoming plug-and-play for both NAS and SAN. NetApp has taken this one step further by supporting NAS and SAN over the same network infrastructure.

The NetApp SnapDrive interface is used to perform iSCSI (and FC) management and maintenance from a simple "single pane of glass" with a familiar NT look and feel. In contrast, when ESG Lab has configured solutions from other iSCSI vendors to date, three interfaces were used instead of one as depicted below. ESG Lab found the SnapDrive interface simple and easy to use.

Figure 2: SnapDrive Native MMC Plug-In

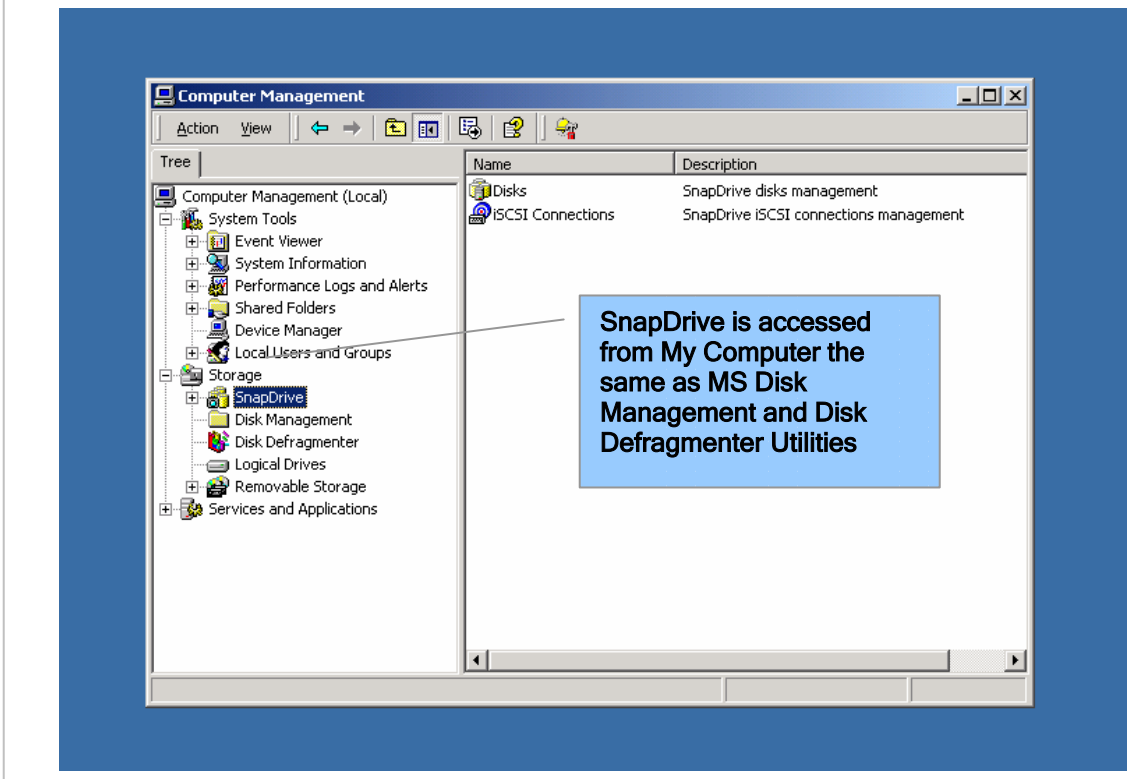

Figure 3: SnapDrive - A Single Pane of Glass

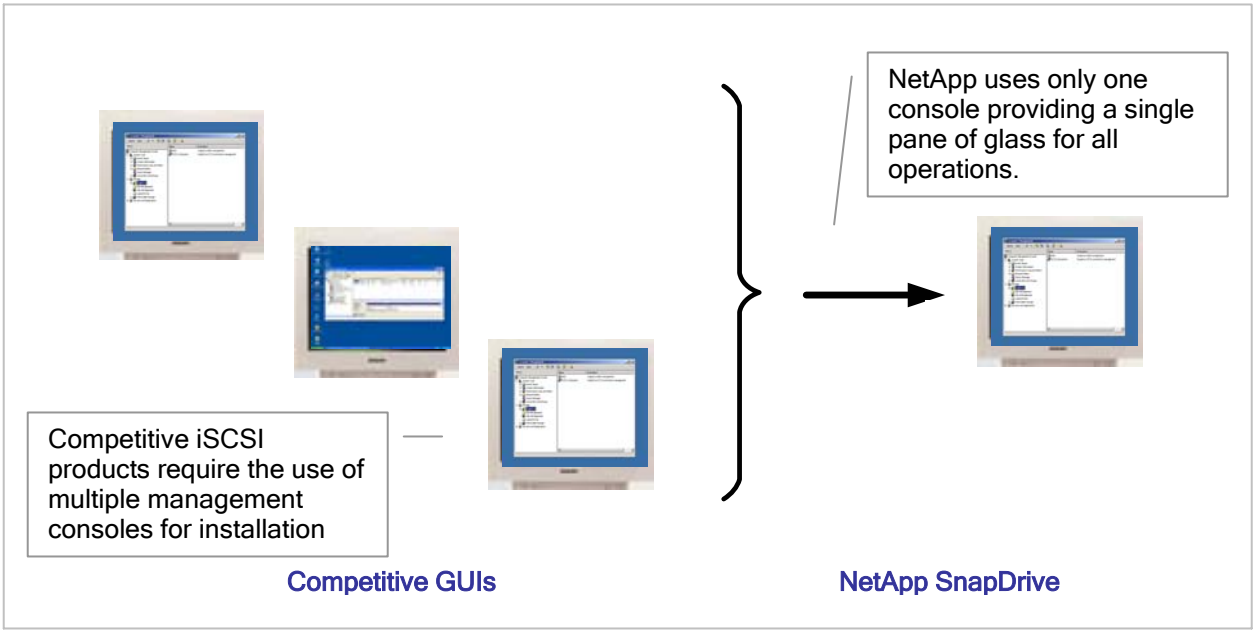

iSCSI Snapshots and SnapManager for Exchange

One of the most popular data management tools used with NetApp NAS Filers is the ability to make snapshot copies. Snapshots are used as point-in-time images for quickly restoring deleted or corrupt files and shrinking backup windows.

The amount of storage consumed for each incremental snapshot is only the amount of data that was written since the last snapshot was taken. This makes snapshots, sometimes referred to as differential snapshots, cost effective, and system administrators can keep multiple versions online for quick restores. Having lots of copies from different points-in-time means that companies have a better chance of recovering the "best" version of the data before a disaster occurred. Because only the incremental data is being copied, snapshots typically take up approximately 20% of the storage capacity compared to full volume copies.

Using snapshots on block-based devices is a bit more complicated than making a snapshot on a NAS system. When a snapshot needs to be created, the application (e.g. Microsoft Exchange) might be in the middle of a transaction, such as updating calendars for a meeting. To ensure that the snapshot is a consistent and useable image, integration with applications is needed to make sure any data is flushed from memory and written to disk in a consistent state. One way to get this done is to use scripts to shut down the application or put it into a "hot backup" mode and then "split a mirror" or "take a snapshot."

Another approach integrates the application and the storage system's snapshot capability using a single management interface for control. For example, NetApp SnapManager for Microsoft Exchange is a well integrated and simple plug-in. SnapManager for Exchange integrates Microsoft Volume Shadow Copy Service (VSS) technology with the NetApp Filer and SnapDrive capabilities to enable easy configuration and fast backup/recovery of Exchange databases.

Why This Matters

Making multiple copies of data through snapshots is probably one of the most important features of a storage subsystem, allowing customers to recover from file corruption and deletion.

Differential snapshots make more efficient use of storage capacity than full volume copies and save customers money while providing more copies of the data at different points-in-time to recover from.

Integrating snapshot technology with applications is the next step, allowing copies to be made without impacting user data availability or performance.

ESG Lab walked through the Configuration, Backup, and Restore wizards, noting the Exchange-centric interface and familiar Microsoft look and feel. Creating snapshots of the Exchange environment was extremely easy using the wizard-driven SnapManager GUI.

The following is a real life example of how system administrators can use SnapManager for Exchange to manage e-mail accounts:

1. The system administrator takes three snapshots per day and retains two days of snapshots as depicted in the diagram below[2].
2. A user lost her entire calendar during a software installation of a new PDA. She needed her Exchange environment restored from the most recent backup ASAP.
3. The system administrator was able to restore her .PST file from a recent snapshot in less than an hour.
4. With their old tape-based restore process, it was a challenge to get restores from the previous night's backup completed within four hours.

[2] Beyond the retention of two days of snapshots for typical restore requests, the customer utilizes an advanced Exchange backup scheme which mirrors snapshot images nightly between three data centers on three continents in a round-robin fashion. In this manner, the customer is protected from a site failure and the use of tape has been rendered obsolete.

Figure 4: SnapManager for Exchange

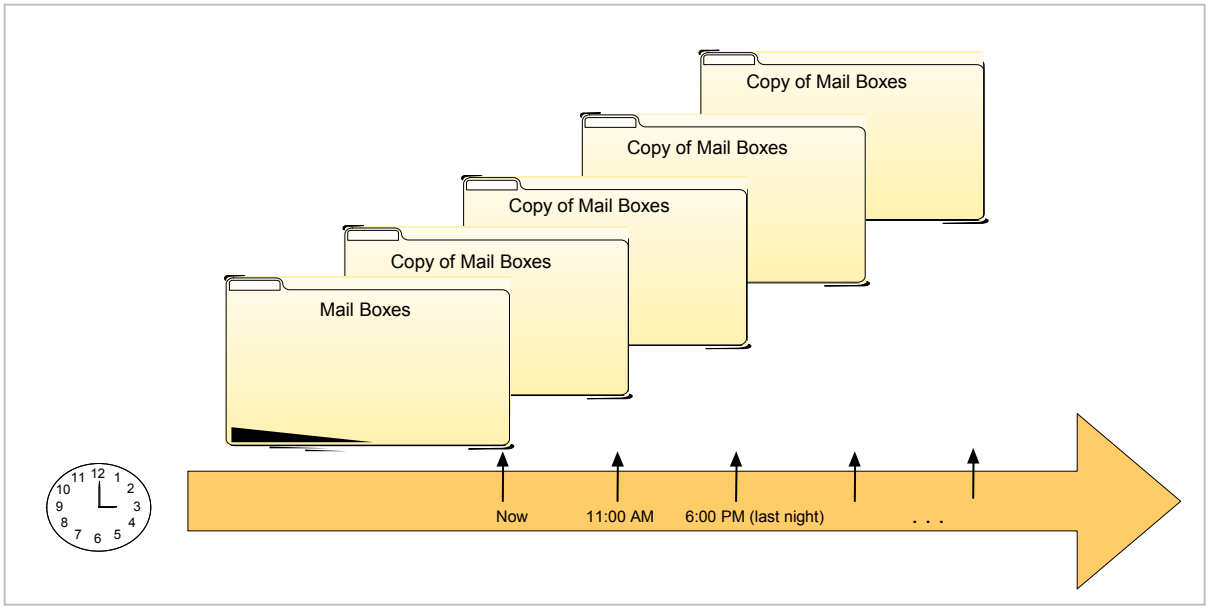

NetApp supports using scripts and integration with MS VSS for MS Exchange for block-based snapshots as described above. NetApp recognizes the advantages of integrating with the applications and will continue to do so as demonstrated with MS Exchange.

ESG Lab was also impressed with the user interface and attention spent in the migration capability of the configuration wizard. The migration capability makes it easy for system administrators to move an existing Exchange server running on DAS to NetApp iSCSI SAN storage.

Use Case: NetApp iSCSI Snapshot

ESG Lab spoke with a NetApp customer who is using an iSCSI SAN for an application called Clearcase.

The Clearcase server is used by software developers to manage revisions and releases of software code. The Clearcase application was running on internal SCSI block storage. Whenever the customer had to perform a restore, the entire application database needed to be recovered to be useful. This could take up to half a day using the tape backup equipment.

The user had prior experience with NetApp snapshots for their NAS environment. Snapshots enabled him to quickly dial the clock backwards in time to restore files without the use of tape. They wanted to use the same recoverability capability for the Clearcase server. During a server upgrade the Clearcase database was moved to iSCSI on an existing Filer being used for NAS.

Clearcase recovery times have gone from as long as four hours to practically instantaneous. Based on the success of this project, the user is planning to move their MS Exchange to iSCSI.

iSCSI Performance

ESG Lab wanted to conduct iSCSI performance testing using the NetApp Filer, since some customers have expressed concerns about performance and whether using an Ethernet network for storage would impact their productivity. ESG Lab is convinced that iSCSI SAN performance is more than sufficient in mid-range applications and will continue to get better over time.

NetApp iSCSI performance testing was performed with the following questions in mind:

1. How does iSCSI performance compare to Direct Attached Storage (DAS)?
2. What happens when multiple servers, each with their own DAS, are consolidated to run instead on a shared iSCSI enabled NetApp Filer?
3. How much faster is FC than iSCSI?
4. For what applications and configurations is iSCSI performance well suited?

ESG Lab iSCSI Performance Conclusions

The results of ESG Lab testing, along with selected ESG Lab audited results of the NetApp consolidation test are summarized as follows:

- iSCSI performance for consolidated real-world and simulated OLTP workloads is faster than similarly configured DAS:
  - iSCSI had up to 41% faster throughput
  - iSCSI had up to 63% better response time
- The DAS configuration tested had problems scaling beyond 2,500 interactive database users as evidenced by a drastic rise in application level response times. Similarly tested NetApp configurations scaled flawlessly through 4,000 users while maintaining acceptable response times, even while other consolidated workloads driven by multiple hosts were directed at the same pool of NetApp storage.
- ESG Lab testing showed that the DAS configuration did not scale as well as the NetApp iSCSI solution due to a lack of CPU horsepower in the PCI-X RAID FC controller tested[3]
- NetApp FC performs and scales somewhat better than iSCSI
  - FC had 11% more OLTP throughput
  - FC had 25% better OLTP response times
- For the configurations and real-world workloads tested, ESG Lab concludes that the performance of iSCSI is well suited for departmental applications including e-mail and interactive database applications.

Why This Matters

iSCSI SANs outperformed a similar configuration of DAS. iSCSI performance is well suited for many storage environments and is being widely used for MS Exchange and SQL applications. An iSCSI SAN will provide huge management benefits to customers who are today using islands of DAS to support their applications.

The performance of FC SANs versus iSCSI was not earth-shattering. This confirms other performance testing performed by ESG Lab. When performance is crucial to the application, then FC is the requisite choice, but for many environments, iSCSI provides an excellent price/performance alternative to FC.

There may be other reasons why customers do not want to use an iSCSI SAN – performance should not be one of them.

[3] The RAID controller performed at about 25% the speed of the NetApp iSCSI and FC tests when performing pure IOPS testing of 4K reads from a small 64 KB cached region. The embedded processor on the RAID adapter is much weaker than the processing power of a Filer. Further, during this testing ESG Lab learned that the number of IOPS out of cache that a NetApp Filer can perform is roughly constant regardless of whether the I/O is performed over an FC interface or a GbE iSCSI pipe.

- ESG Lab believes that the price/performance of a NetApp iSCSI solution for the tested application workloads is better than a comparable FC solution due to the fact that the cost of FC infrastructure is more expensive than iSCSI over Ethernet infrastructure.

ESG Lab Performance Testing

ESG Lab reviewed the results of an exhaustive series of performance consolidation tests performed by NetApp during the summer of 2003. The performance of real-world applications running on servers each with their own DAS was compared to a configuration where those same servers were configured to instead use centrally managed NetApp Filers. A diagram of the consolidated iSCSI configuration tested is depicted in Figure 5 below[4]. Three servers share the same pool of consolidated NetApp Filer storage over Ethernet using the iSCSI protocol, and testing was performed with each of the following real-world applications running alone and then together:

- On Line Transaction Processing (SQL server order entry with 1,000 through 4,000 users)
- E-mail (Lotus Notes with 1,000 through 4,000 users)
- Business Intelligence (long running SAS queries)

NetApp consolidation testing was performed over a period of months involving multiple teams and a large amount of equipment. ESG Lab set out to validate the results of that consolidation testing by creating a similar configuration and performing a subset of comparable tests. In addition, ESG Lab interviewed performance engineers involved in the testing, and performed a detailed audit of the test configuration, methodology, and results.

Figure 5: iSCSI Performance Test Configuration

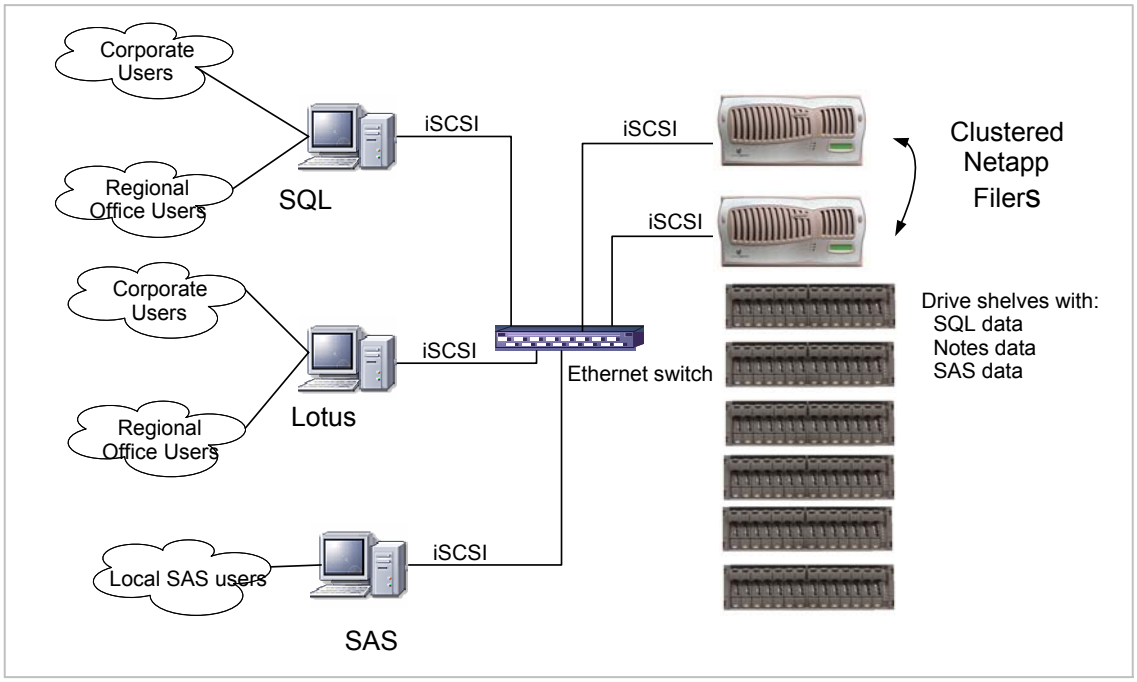

[4] See Appendix for configuration details.

ESG Lab performance testing began with a simulated OLTP workload for comparison to results obtained during consolidation testing. The standard 8K IOMETER OLTP workload originally distributed by Intel with the IOMETER toolkit was run against similarly configured DAS, iSCSI, and FC drive configurations. The results presented below were compared to those presented in the Chart 1, which represents the performance of an OLTP SQL server application gathered during a consolidation run that consisted of a mix of real-world applications.

Chart 1: iSCSI versus DAS

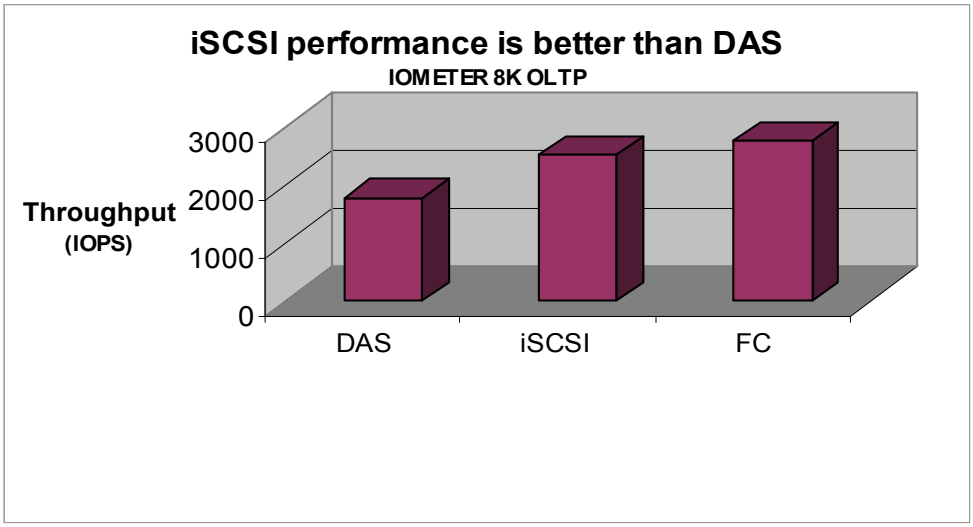

What the numbers mean:

- The DAS configuration tested was a Mylex PCI-X FC RAID controller connected to two trays of FC drives of the same type with a comparable configuration to that used in the NetApp iSCSI and FC tests.[5]
- The relative throughput measured in I/O's per second is indicated on the Y-axis, with the FC configuration performing best and the iSCSI configuration coming in a close second.
- iSCSI performance for simulated OLTP workloads is faster than similarly configured DAS:
  - 41% more throughput
  - 30% faster response time
- FC performance for the configuration tested was slightly better than iSCSI.

[5] See Appendix for more configuration details.

Chart 2: iSCSI versus FC

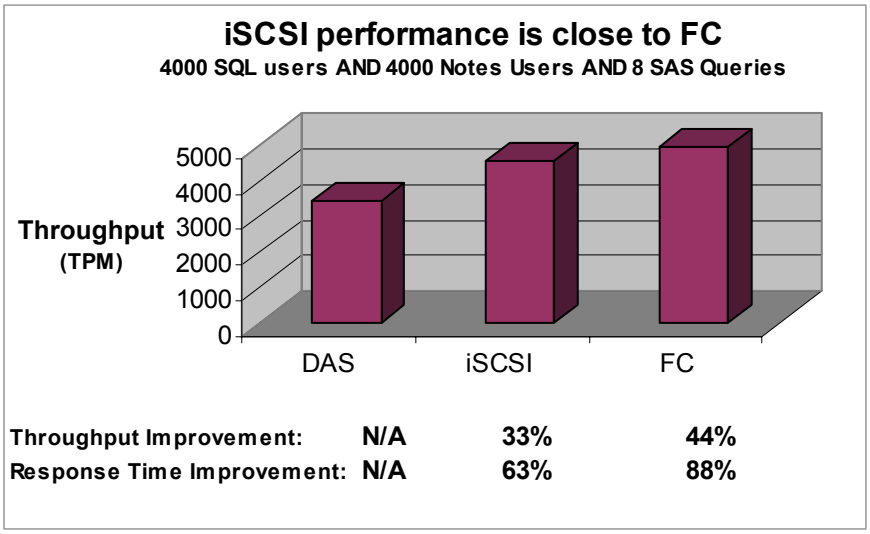

What the numbers mean:

- Performance for 4,000 simulated interactive SQL server users was measured while 1,000 simulated Lotus Notes e-mail users and 2 SAS queries ran in parallel.
- Relative throughput measured in database transactions per minute is indicated on the Y-axis.
- iSCSI outperformed DAS and FC performed better than iSCSI, as in the IOMETER results presented above.
- The performance advantage of iSCSI over DAS dropped a bit compared to the IOMETER results and the advantage of FC over iSCSI narrowed slightly. The likely reason for the variance is a less random mix of I/O in the real-world consolidation test run as compared to the simulated IOMETER test run.

Why This Matters

The ESG Lab performance testing and research showed that the iSCSI SANs consistently outperformed DAS and DAS islands and performed only slightly less than FC SANs.

ESG Lab showed that DAS hit a performance wall that iSCSI easily surpassed. This is important since iSCSI SANs are a great alternative to DAS, especially in environments that cannot cost justify a FC SAN.

Chart 3: iSCSI versus DAS Islands

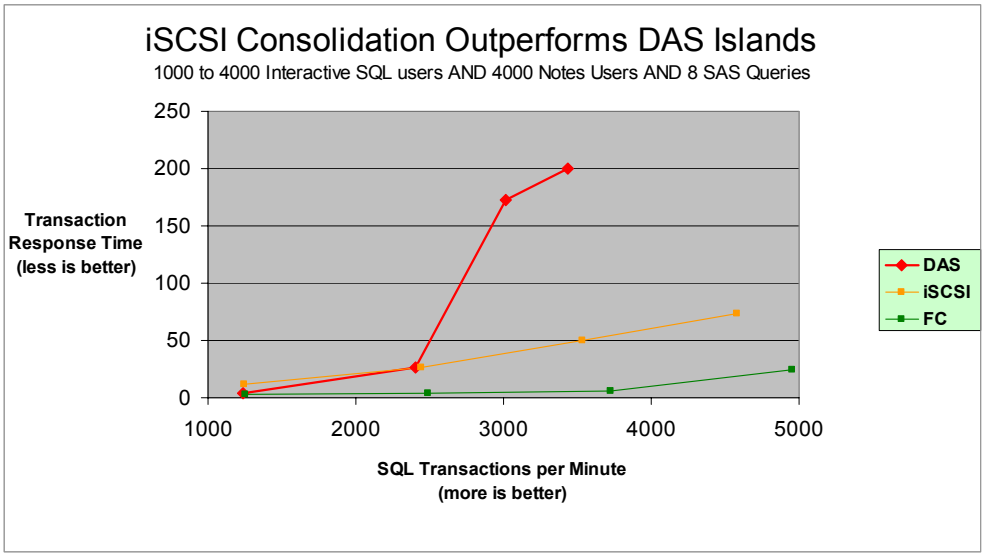

What the numbers mean:

- From 1,000 to 4,000 interactive SQL server users were measured while 1,000 simulated Lotus Notes e-mail users and 2 SAS queries were running at the same time.
- The four data points depicted for each configuration moving from left to right indicate performance measured at 1,000, 2,000, 3,000, and 4,000 interactive SQL server users respectively.
- As the number of users increased, the amount of throughput, measured in SQL transactions per minute, increased.
- As the throughput increased, transaction response time tended to increase.
- When the system becomes overloaded, then the "knee of the curve" shows up, as in the DAS line, which trends quickly upward at about 2,500 transactions per minutes. This sharp increase in response time indicates that the system, in this case the PCI-X FC RAID controller, was swamped with I/O and couldn't keep up. This resulted in slow application response times for users with performance becoming increasingly sluggish. This unpleasant condition is sometimes referred to as an "application brownout."
- The DAS configuration couldn't handle much more than 2,000 interactive SQL users while the NetApp iSCSI and FC solutions handled 4,000 users and could possibly handle more, even as they simultaneously serviced SAS queries and thousands of e-mail users.
- As traffic levels increased, iSCSI performance tended to slow down (rising relative response times) in a manageably linear fashion, while FC performance tended to support more traffic before encountering a similarly manageable linear slow down.

ESG Lab Performance Testing Notes:

1) Streaming - ESG Lab measured large block sequential read performance over a 2 Gb FC pipe at almost twice as fast as iSCSI over 1Gb Ethernet. Those results are not presented in the graphs above. Applications with streaming characteristics that demand high throughput (e.g. a video-on-demand server, some scientific applications) may be better served by FC.
2) Mission critical interactive - FC may be more appropriate than iSCSI for mission-critical interactive applications where very fast response times are imperative.
3) Network utilization - Network utilization was not an issue for iSCSI performance tests presented in this report. Testing occurred in a small isolated GbE network with jumbo frames enabled. Ethernet traffic between clients and servers was simulated and avoided. As iSCSI matures and is deployed in large and active networks, increased network utilization may become an issue. Standard Ethernet VLAN features can be used today to isolate and increase the performance of iSCSI networks. In addition, physically isolating an iSCSI network, as well as trunking and teaming, all of which NetApp's iSCSI implementation supports, can also be used to minimize such congestion and to isolate data traffic from general network traffic for increased security. The evolution and deployment of 10 GbE Ethernet will positively impact this issue as well.
4) CPU utilization - iSCSI performance test results presented in this report were obtained while testing relatively new servers with quad 2.4 GHz Xeon processors and off-the-shelf GbE adapters. The CPU utilization of the servers during IOMETER driven workloads tests was generally low (under 20%), with the worst case tested (100% 4KB reads from storage subsystem cache) at 42% on one processor and in the low single digits for the other three processors. CPU hungry applications and/or applications running on relatively older servers with less CPU horsepower may require hardware assistance in the form of a TCP/IP offload (TOE) engine and/or iSCSI protocol offload in the form of an iSCSI HBA. An evaluation of the benefits of a TOE enabled GbE adapter or an iSCSI HBA is beyond the scope of this evaluation.
5) Mileage may vary - Storage benchmark testing with standardized workloads, by definition, does not necessarily represent results you will see in your installation. Although the caveat of "your mileage may vary" holds, ESG Lab warrants that the test methodology and results presented in this paper are a sound basis for the observations presented.

Reliability

ESG Lab performed error injection testing to test their latest iSCSI offerings. A redundant power cable and a drive carrier were pulled while a copy of a large amount of data to an NT iSCSI drive was in progress. Correct error notification and a drive rebuild were duly noted. The iSCSI Ethernet cable was pulled and NT timeouts were observed. A complete power failure and reboot of the NetApp Filer were then performed. Correct operation was observed after the NT server was shut down and rebooted.[6]

iSCSI Qualification

iSCSI vendors, including NetApp, focused first on the Windows and Linux markets where early adoption is first taking place. NetApp is working diligently with partners and operating system vendors to bring iSCSI to new platforms and users. For example, NetApp is working closely with Novell to bring iSCSI support to Netware users, and major UNIX vendors like HP are working closely with NetApp to qualify iSCSI drivers and configurations. In addition, NetApp is working with partners to qualify iSCSI configurations containing TCP/IP offload engines, iSCSI Host Bus Adapters, clustered configurations, etc.

ESG Lab began HP-UX iSCSI testing with a visit to the NetApp On the Web (NOW) support site where documentation and the free iSCSI software initiator driver from HP-UX were obtained. An HP-RP5470 server running HP-UX version 11i v2 was used for testing. The HP-UX server was connected to a NetApp FAS250 using standard GbE adapters and cabling (i.e. no special TCP/IP offload or iSCSI adapter hardware was needed). Downloading and installing the HP-UX iSCSI driver using the HP-UX

[6] NetApp's error handling is fully compliant with the iSCSI specification as it is written today, but enhancements to the specification are needed to ride predictably and transparently through catastrophic failures, particularly in clustered configurations designed for high availability. Efforts in the standards committee and between vendors are progressing rapidly towards that goal.

standard Software Distributor (sd) tool was simply uneventful. The sd utility completed the installation with an automatic reboot. After the iSCSI driver loaded automatically during the next reboot, iSCSI targets were discovered and configured using the familiar HP-UX ioscan and SAM utilities. In other words, once the HP-UX iSCSI driver had been installed, working with an iSCSI device felt exactly like dealing with a local or FC drive.

Figure 6: HP UX Qualification

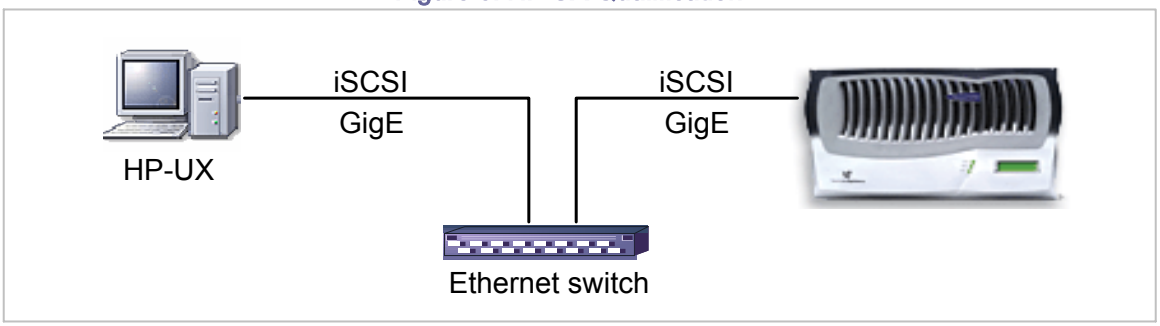

iSCSI Support

ESG Lab was impressed with the NetApp NOW support site mentioned previously in this report. ESG Lab found the site helpful and easy to use with an open approach to sharing information, including access to existing iSCSI related bug details, iSCSI specific knowledgebase articles, an iSCSI forum, chat access to live support representatives, and operating specific iSCSI documentation, including a very concise and helpful quick reference sheet for each operating system. NetApp has done an excellent job of making iSCSI installation and support information readily available.

Another example of the richness and simplicity of the site is that it can be used to download and upgrade new versions of NetApp embedded software. Embedded software upgrades are downloaded directly from Network corporate headquarters to a NetApp system, eliminating slow and cumbersome maintenance interface upgrades.

Customers interviewed by ESG Lab identify their sales and support team as one of the reasons they continue to buy from NetApp. ESG Lab interviewed one of the 120 NetApp support specialists. NetApp has four support centers spread over three continents so they can "follow the sun." Escalation procedures are codified and enforced with a one-hour escalation procedure in place for high priority problems. A comprehensive knowledgebase of over 7,000 articles is available for support staff and customers. NetApp was the first vendor to receive Support Center Practices (SCP) certification on three continents. SCP is an internationally recognized standard created by the Service & Support Professionals Association (SSPA), which has defined best practices for delivering world-class technology support. Much like ISO certification is to the manufacturing industry, SCP certification is the gold standard of the support industry.

Conclusion

ESG Lab validation goals for NetApp and their iSCSI support were met, including a simple installation process, easy to use management tools, excellent performance, support, and reliability. ESG Lab interviewed NetApp customers who are using Filers for both NAS and iSCSI and have a unified product for file sharing, database, and email storage. NetApp has a long history of providing easy to use NAS solutions, which they have successfully carried over to their SAN offerings.

Appendix

Consolidation Performance Testing

Consolidation testing was performed on three servers with DAS, FC, and iSCSI attached storage. The following diagram, which depicts the iSCSI configuration tested, is followed by a summary of the configuration and methodology. Note that the same servers and workloads were used for each of the DAS, iSCSI, and FC tests. Note also that the number and type of drives (32 FC drives spinning at 10K RPM) were used for each test for an apples-to-apples comparison.

Figure 7: Consolidation Test Configuration

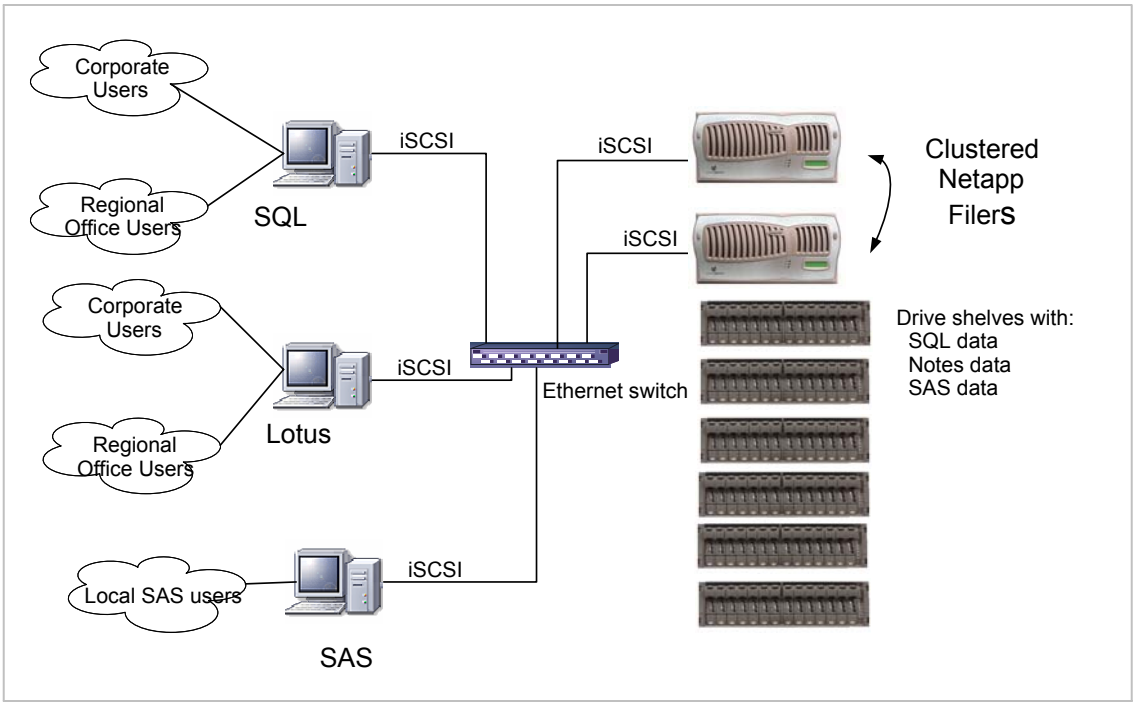

- **Three** Dell 2650 servers running three different applications:
  - On Line Transaction Processing (order entry application with 1000 to 4000 users)
  - Business Intelligence (long running SAS queries)
  - E-mail (Lotus Notes from 1000 to 4000 users)
- Each server was tested with three different storage technologies:
  - Direct Attached Mylex PCI-X RAID controllers
  - iSCSI attached NetApp Filers
  - FC attached NetApp Filers
- **Four** trays of 10K RPM FC drives (32 drives in total) were used for all tests
- Mylex drives were configured as RAID-10 (mirrored groups of striped drives) for optimal performance
- Filers were configured to the NetApp RAID-4 default
- A Mylex PCI-X RAID controller connected to four trays of RAID-10 protected 10K RPM FC drives
- For more details, see http://www.NetApp.com/tech_library/3273.html

IOMETER Performance Testing

IOMETER performance testing was performed using a similar, but smaller, configuration to that used during consolidation testing described above. Note that only one server and 16 drives were used for testing, compared to the three servers and 32 drives that were used for consolidation testing.

Figure 8: IOMETER Test Configuration

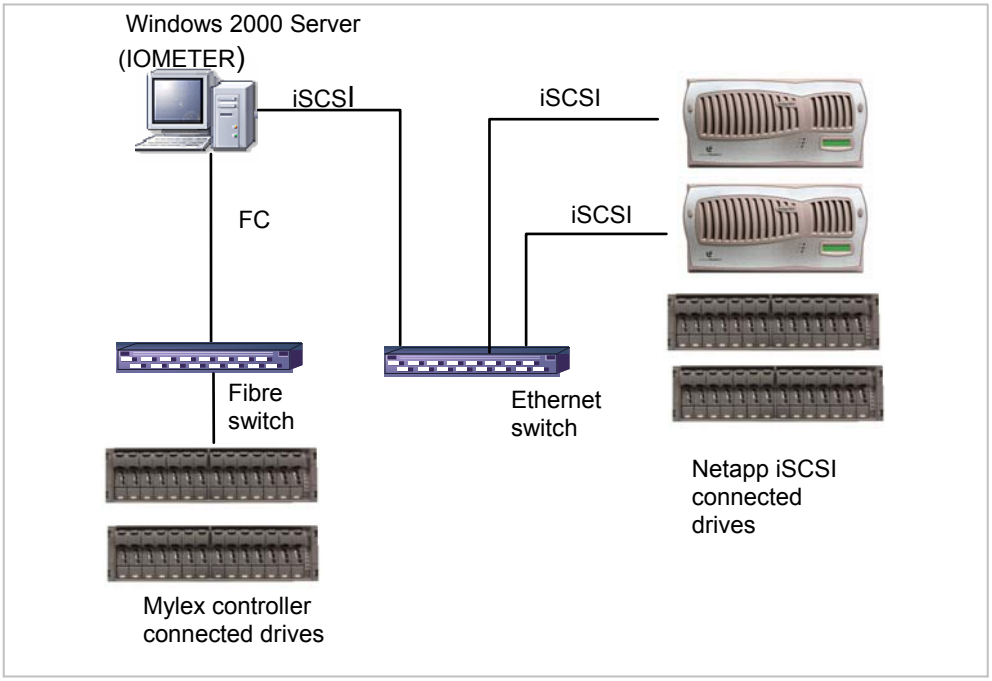

- **One** Dell 2650 server running IOMETER simulated workloads:
  - Intel standard 8K OLTP (simulating On Line Transaction Processing)
  - Large block sequential (simulating Business Intelligence queries)
  - ESG Lab Exchange simulation (simulated e-mail)
- Tested with three different storage technologies:
  - Direct Attached Mylex PCI-X RAID controllers
  - iSCSI attached NetApp Filers
  - FC attached NetApp Filers
- **Two** trays of 10K RPM FC drives (16 drives in total) were used for all tests
- Mylex was configured as RAID-10 (mirrored groups of striped drives) for optimal performance
- Filers were configured to the NetApp RAID-4 default

IOMETER test configuration details:

1. IOMETER version 2003.05.10
2. Dell 2650 server, Dual 2.4 Ghz CPU's, 4 GB of memory
3. Windows 2000, Service Pack 3
4. Direct attached RAID controller: Mylex ExtremeRaid 3000 PCI-X FC, Driver version 7.0.15.0, Firmware version 7.01-00
5. Direct attached drive configuration: Hitachi DK32EJ-72FC FC drives, 32 drives, RAID-10 over 32 GB useable

6. NetApp iSCSI / FC drive configuration: Seagate ST336704FC FC drives, 32 drives, RAID-4 over 32 GB useable
7. GbE adapter: Intel(R) Pro/1000 XT server adapter (copper), Driver version 7.2.17.0
8. FC HBA: Emulex LP9002, Firmware 3.90a7, Device Driver 5.5.1.0
9. GbE switch: Dell PowerConnect 5224 GbE switch with jumbo frames enabled
10. FC switch: Brocade 3200

Table 1: IOMETER Configuration

IOMETER Access Name	Size	% Access	% Read	% Random	Max Size	Outstanding I/O's
8 KB OLTP	8 KB	100	67	0	0 (all avail)	1, 4, 16, 32
IOPS	4 KB	100	100	100	128 sectors	128
Streaming	512 KB	100	100	0	0 (all avail)	1,4,16,32