ESG Lab Validation Report™

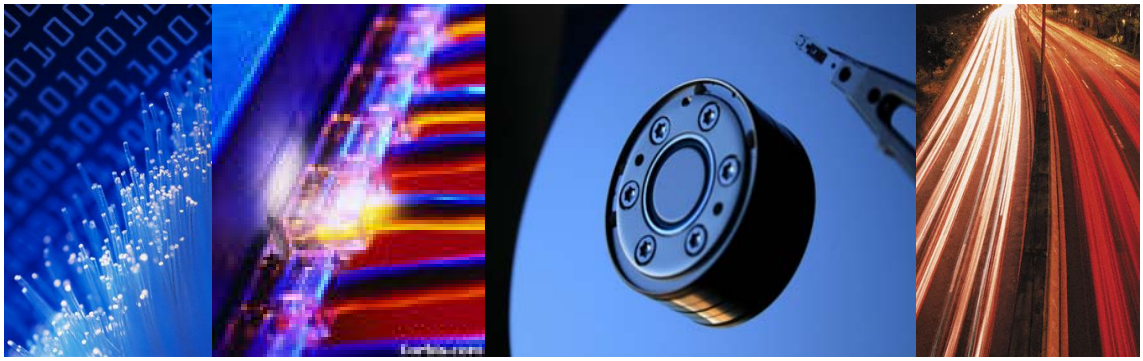

LeftHand Networks® 100 TB Enterprise SAN

A validation study
by
ESG Lab
July 2007

Authors:
Brian Garrett, Tony Palmer and Tony Asaro

Table of Contents

ESG Validation Reports

The goal of ESG Lab Validation reports is to educate customers about various storage and storage-related products, including storage systems, backup-to-disk solutions, storage management applications, backup and recovery software, storage virtualization platforms, etc. ESG Lab reports are not meant to replace the necessary evaluation process that end-users should conduct before making purchasing decisions, but rather to provide insight into these technologies. Our objective is to go over some of the highlighted features/functions of such products, show how they can be used to solve real customer problems, and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was commissioned by LeftHand Networks.

Introduction

When ESG Lab analyzed LeftHand Networks in June of 2006, the company was emerging as one of the leading next generation SAN storage system vendors. LeftHand is one of the best kept secrets in storage with over 2,000 end-user customers utilizing LeftHand products as their core SAN solution. Additionally, they have qualified and are partnered with major server vendors including HP, IBM and Intel, supporting a strategy of using sophisticated and feature-rich storage software on open platforms. Recently, ESG Lab revisited LeftHand to test their SAN/iQ software with a focus on scalable high-end application performance and improved ease of use.

Background

The biggest problem with managing storage environments is the complexity of the storage systems themselves. Most leading storage systems are complicated machines that require deep expertise and ongoing manual administration to maintain operational excellence. This complexity exists because these storage systems are not highly virtualized.

The more virtualized the storage system, the easier it is to manage and maintain operational excellence. Whenever we require a human to perform manual functions, regardless of their level of expertise, operational efficiency decreases and the chance of error increases as the system grows in size and complexity.

Clustering is one of the most powerful and useful forms of virtualization. A cluster enables multiple physical systems to be seen as a single logical system. As a result, clustered storage removes the limitations of individual physical devices, effectively eliminating the boundaries of the "box." Once these limitations are removed and the storage system becomes boundless, IT can transcend the daily grind of managing physical assets and become truly strategic to the business.

N-way clustered architectures support multiple storage controllers in a single cluster. Though these clusters may contain many storage controllers, they still appear to the applications and users as a single logical system. As such, clustered storage systems allow users to add CPU, memory, and bandwidth transparently and enable users to scale based on the needs of the business without having to buy a whole new storage system. In traditional dual node storage systems with fixed architectures, when a user's environment outgrows their storage system, they may be forced to buy another system to achieve greater performance or capacity. With clustered architectures that allow for the aggregation of all hardware resources, performance and capacity can be scaled in a linear fashion–if and as needed.

LeftHand Clustering

LeftHand provides a true clustered architecture with advanced features and functions that make it a compelling solution. The architecture was developed with scalability, availability and ease of use firmly in mind. While LeftHand's offering has been known as a midrange solution, the SAN/iQ cluster scales to meet the needs of Enterprise-class applications with relative ease. LeftHand advises ESG that 40% of sales are additions to existing clusters while the average cluster size sold is 4-6 nodes, the average cluster size deployed in the field is 15-20TB and 20% of the clusters deployed in production contain more than 10 nodes. To offer a platform that can truly scale for the enterprise, LeftHand integrates Enterprise-class servers such as the HP DL320s and the IBM system x3650. Disk drive options support capacity and performance configurations including SATA, SCSI and 15,000 RPM Enterprise-class Serial Attached SCSI (SAS) drives. The system is offered in 10Gb Ethernet ready configurations as well.

The core of LeftHand's value is its SAN/iQ storage software platform. SAN/iQ runs on enterprise-class x86-based server platforms. As LeftHand and its partners provide customers with fully integrated SAN solutions, it's important to be aware that LeftHand's open iSCSI SAN is a software solution that leverages enterprise-class x86 server platforms to provide advanced storage system solutions. LeftHand is highly focused on its Open iSCSI SAN initiative, with the goal of providing open iSCSI SANS through its SAN/iQ Storage Software Platform running on standards-based servers that support different drives and interconnects. LeftHand is putting a stake in the ground by focusing all of its research and development at its advanced software, while leaving HP, IBM and others to invest in storage server hardware.

The LeftHand Open iSCSI SAN architecture leverages the advancements of standards-based technologies including processors, memory, disk drives and Ethernet networking. In doing so, LeftHand rides the wave of massive industry innovations. Additionally, it drives down cost because of the commoditization of these technologies and products. LeftHand is then able to focus on its major strength—the SAN/iQ storage software platform.

Traditional storage systems typically cluster two controllers together, creating an active-active or active-passive pair that appears as a single logical storage system. Dual-node controller configurations have become the predominant storage system architecture, used by all of the leading storage vendors including EMC, HDS, IBM, Network Appliance, etc. These solutions are limited by the number of host ports, cache memory and storage that they support, ranging from two to eight host ports, two to four processors, 1 GB to 16 GB of cache memory and 500 GB to 500 TB of storage. These storage systems are typically classified as midrange or modular and are experiencing the greatest market growth in network storage.

A high-end midrange storage system that supports eight host ports may not provide enough connectivity for a large server environment. Additionally, since midrange storage systems typically support only two processors, they may not have adequate compute power to handle a large number of transactions, which can impact application performance and ultimately user/customer satisfaction. This is also true of cache memory, which has been proven to greatly improve application performance. If customers run into performance and scalability bottlenecks using traditional active-active storage systems, their only recourse may be to buy another system.

LeftHand open iSCSI SANs are comprised of enterprise-class, industry-standard servers configured as fully contained storage modules that provide CPU, memory, bandwidth and capacity. The LeftHand SAN/iQ storage software platform runs on the storage module, providing intelligent storage system functionality. Customers can add storage modules for more scalability as needed. The LeftHand cluster will remain a single logical system regardless of how many storage modules are added to it, making it just as easy to manage a two-node cluster as a 30-node cluster. Additionally, adding modules to the cluster is a transparent and online process.

Customers can scale and grow the LeftHand network clustered storage systems from small through medium to large configurations just by adding more storage modules to the cluster. Every node added to the cluster increases the number of processors, host ports, cache memory and supported capacity. Many of the enterprise-class, midrange and low-end systems from leading storage vendors are, for the most part, incompatible with the same storage vendors' systems at different levels and use different software and hardware. This increases the complexity and cost of migrating up (or down) to meet ever-changing business requirements. LeftHand provides the ability to grow from a low- though mid- to high-end storage system easily, using the same equipment.

The economics of a clustered network storage system are compelling, with the potential to significantly reduce capital and operational costs. With LeftHand's SAN/iQ, the customer only has to add another storage module—which costs far less than acquiring a whole new system—and its associated software, licensing and maintenance charges. Most midrange storage systems support at best a dual-node clustered configuration. This limits scalability and flexibility.

It is important to note the other benefits of a highly virtualized storage cluster, such as thin provisioning. LeftHand has supported this incredibly valuable capability since day one. Thin provisioning minimizes the cumbersome task of storage provisioning while improving capacity utilization.

When provisioning storage using traditional methods, system administrators typically dedicate a fixed amount of storage to each particular application. For example, if a 500 GB volume is allocated to an application that requires only 100 GB of actual data, 400 GB are left with no data stored on it. That unused capacity is still dedicated to that particular application and no other application can use it. This means that the unused capacity of 400 GB is wasted storage, which represents wasted money. Though all of that storage capacity may eventually be used, it might take years to do so. This is a major problem when managing storage capacity, often referred to as stranded storage or allocated-but-unused storage. The problem is exacerbated by local and remote copy facilities that consume additional storage by duplicating

allocated-but-unused capacity (copying "nothing" to "nothing").

Thin provisioning offers a simple solution to the problem of stranded capacity. It is a storage system technology that allows users to safely allocate as much logical capacity as needed to an application over its lifetime. Meanwhile, physical capacity is drawn from a common pool of storage on an as-needed basis. That is, only when an application performs writes is physical capacity drawn from the storage pool. Additionally, physical capacity can be added to the storage pool non-disruptively at any time.

With thin provisioning, the storage provisioning process begins the same as traditional provisioning. Going back to the example in which the system administrator provisions 500 GB to an application with only 100 GB of actual data, the unused 400 GB can still be made available for other applications or need not be purchased at all. This is because thin provisioning is really Just-in-Time storage. The application thinks that 500 GB of storage is available, but the storage system only provides the capacity when needed. The rest of it stays in the available pool. This approach allows the application to grow transparently, while at the same time ensuring that capacity is not wasted. System administrators can set thresholds and are alerted when more disks need to be added into the storage pool.

ESG Lab conducted testing on the 30 node, 100 TB LeftHand storage cluster shown in Figure One. Thirty HP ProLiant DL320s servers running the LeftHand SAN/iQ storage software platform were used as storage modules. Sixteen 1U servers running the Microsoft Windows operating system were configured as iSCSI clients. A Force10 Networks switch supporting one and ten Gigabit Ethernet interfaces was used to create an IP SAN between the iSCSI clients and the storage modules.

Figure One: The LeftHand 100 TB SAN

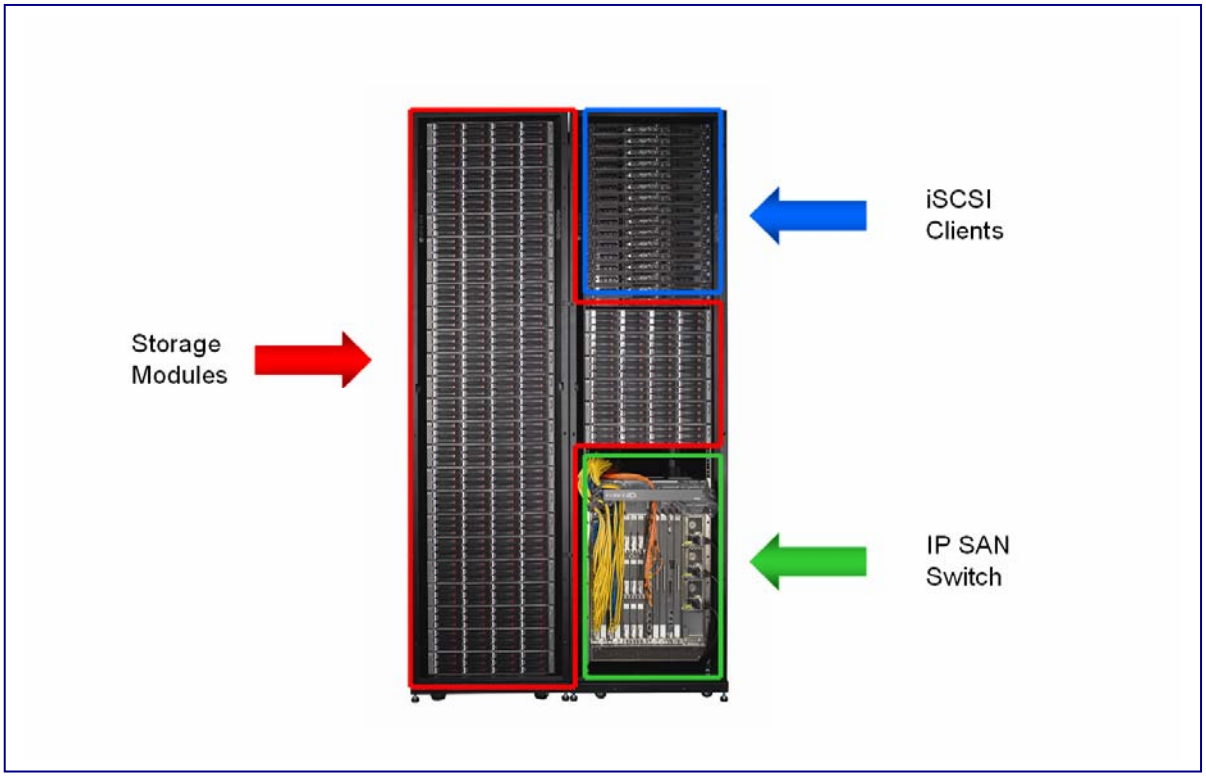

LeftHand describes the Enterprise-class attributes of the SAN/iQ architecture as Scalability, High Availability, Reliability and Performance—or SHARP.

Some of the Enterprise Class features and differentiators of LeftHand's highly virtualized SAN include:

- ☑ Multiple site SANs with scalable capacity, performance and redundancy via LeftHand's Cluster.
- ☑ Automatic cluster failover/failback providing maximum availability to users and applications.
- ☑ Flexible, robust replication allowing for simultaneous local, campus-wide and remote mirroring of critical data sets.
- ☑ Network RAID providing advanced data protection and enhancing performance.
- ☑ Easy administration via a single management console, regardless of the number of sites.
- ☑ Automatic, integrated thin provisioning, reducing utilization and administrative effort.
- ☑ Thin Provisioned Snapshots providing point in time copies of data with the most efficient use of capacity.

LeftHand's execution on this vision is further confirmed by their growing presence in Enterprise-class data centers. This ESG Lab Validation Report is focused on confirming the Enterprise-class attributes of the LeftHand SAN/iQ architecture, including performance, scalability, functionality and ease of management.

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of LeftHand Networks 100TB SAN at Percept Labs, an independent product test and consulting company in Boulder, Colorado, specializing in Data Storage, Regulatory/Compliance and product testing. The tests were executed on a LeftHand Networks 100TB SAN/iQ Cluster comprised of 30 HP DL320s servers running SAN/iQ software version 7.0. Each of the DL320s servers was equipped with twelve 300 GB Enterprise-class 15,000 RPM SAS hard drives[1].

Connectivity was provided using an Ethernet Switch from Force10 Networks. Fifteen Dell 1850 systems were used as load generators. While Figure Two shows all components separated out for clarity, it is interesting to note that the entire 100TB configuration (storage, switch and clients) fit into just two industry standard 19 inch racks. Testing was designed to demonstrate system scalability and performance using industry standard tools and methodologies in addition to an audit of results obtained by LeftHand in Microsoft's ESRP program for Exchange. ESG Lab also tested the administration interface to evaluate the improved manageability and usability of the SAN/iQ management interface.

Figure Two: The ESG Lab Test Bed

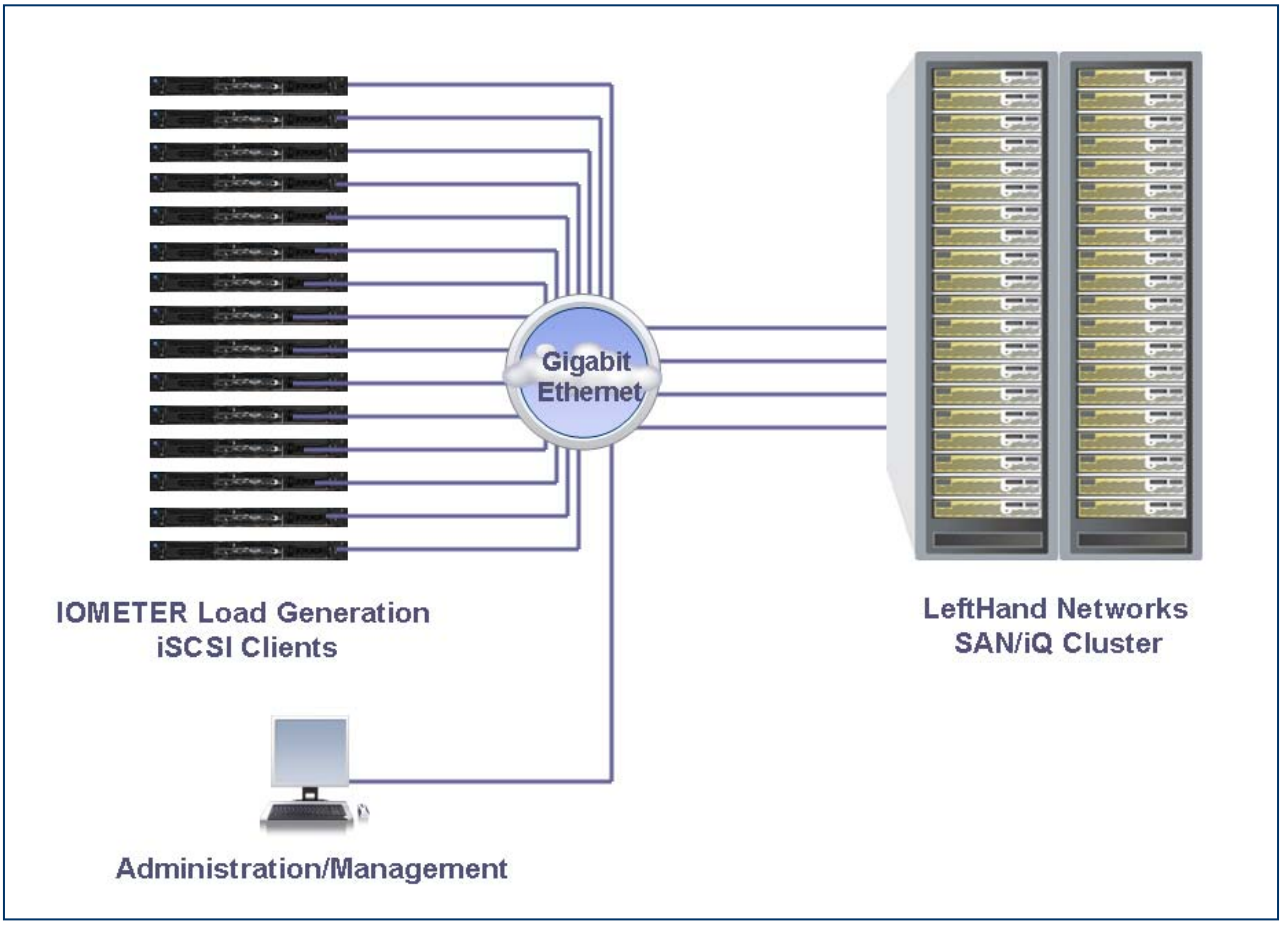

[1] Configuration details can be found in the Appendix.

E-mail Performance Scalability

ESG Lab validated the Enterprise-class performance capability of the LeftHand's SAN/iQ architecture via an audit of Microsoft ESRP 2.0 results and hands-on testing using IOMETER workloads designed to simulate a Microsoft Exchange 2003 environment.

The Exchange Solution Review Program (ESRP) is a Microsoft program designed to facilitate third party storage testing and solution publishing for Exchange Server. ESRP version 1.2 is targeted at Exchange 2003, while ESRP version 2.0 focuses on Exchange 2007. The programs combine a storage testing harness (Jetstress) with publishing guidelines for Microsoft Gold Certified and Storage OEM Partners. Manufacturers use the ESRP framework to test storage solutions and then submit results to Microsoft for review. Approved solution results are posted on the Microsoft Exchange ESRP website[2].

ESRP is different from standard I/O generation benchmarking suites in two important ways. First, ESRP employs the Jetstress utility to create real exchange traffic that runs against real exchange databases–with logging and file attachments–exactly as in the real world. Second, the testing is designed to measure both the performance and reliability of a given solution. The performance test runs for two hours while the reliability test runs for 24 hours. Both tests must run without exceeding a prescribed disk latency threshold (20 milliseconds) and a reliability test is performed to check for database and log corruption at the end of the run. If a storage system supports replication, it must be validated with replication enabled.

The goal of ESRP is to verify that a vendor's storage solution can handle high I/O load for a long period of time. Microsoft makes it quite clear that these tests should not be used for performance comparisons, but in practice, end-users routinely look to these tests to gauge how well a given configuration will perform. The configuration used by LeftHand and approved by Microsoft for Microsoft Exchange 2007 ESRP - Storage v2.0 testing is shown in Figure Three.

Figure Three: The LeftHand ESRP –Storage v2.0 Test Configuration

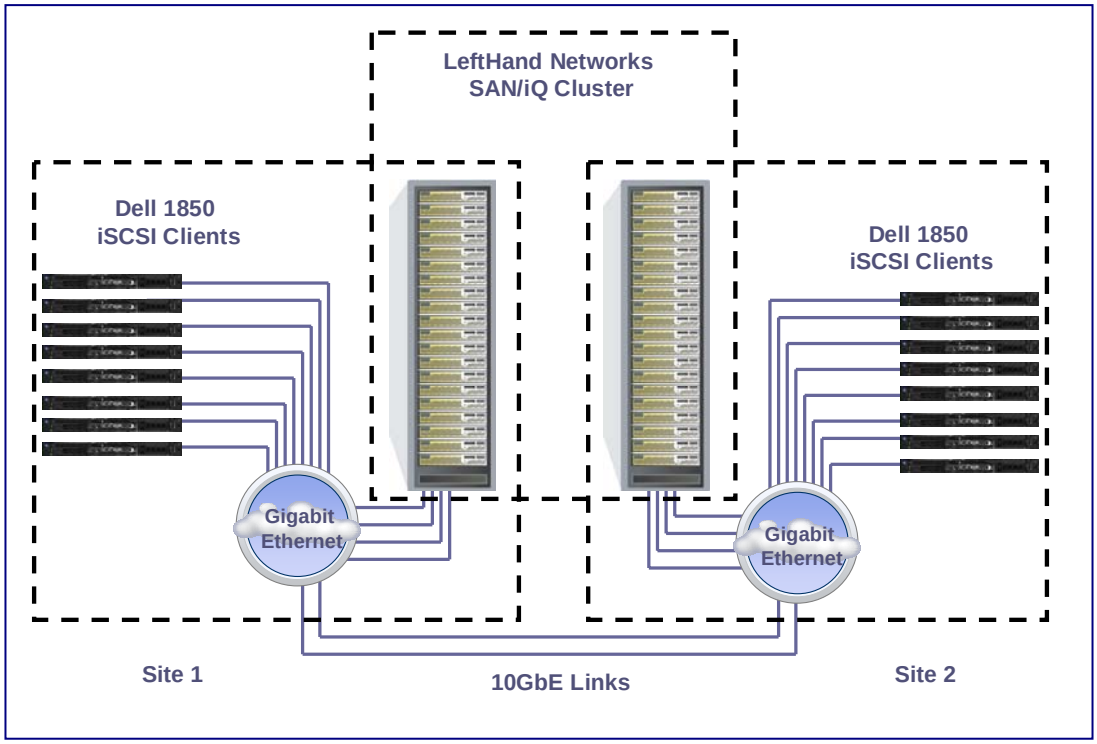

[2] http://technet.microsoft.com/en-us/exchange/bb412164.aspx

As shown in Figure Three, a thirty node LeftHand cluster was split in half and 2ms of latency were injected to simulate a second site located across a campus or metro area. Site One and Site Two were connected by two IP links over a 10 Gbps Ethernet connection. All thirty nodes were configured as part of the same cluster in what is generally known in the industry as a "stretched cluster." LeftHand calls this configuration a "Multi-Site SAN." As such, all thirty nodes are part of the same logical system and volumes can be created anywhere in the cluster. It should be noted that LeftHand also supports Asynchronous replication between two clusters, which can be used to provide data protection and fault tolerance between environments over long distances.

The configuration details for the LeftHand ESRP v2.0 test are summarized in Figure Four.

Figure Four: ESRP Configuration Highlights

Number of Exchange mailboxes simulated	50,000
Number of hosts	16
Number of mailboxes/host	3125
Number of storage groups/host	3
Number of mailbox stores/storage group	5
Number of mailboxes/mailbox store	~208
Number of mailbox store LUNs/storage group	2
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	.4 IOPS / Exchange 2007 'heavy' user
Database LUN size	500GB
Log LUN size	10GB
Backup LUN size/storage group	NA
Total database size for performance testing	19,563 GB
% storage capacity used by Exchange database**	80%

What the Numbers Mean

- ☑ The system was configured with 50,000 simulated Exchange 2007 users–a number firmly in the realm of an Enterprise-class deployment.
- ☑ The I/O profile used for this test was the Microsoft "Heavy" profile for Exchange 2007 with 0.4 IOPS, including a 20% overhead as prescribed by Microsoft.
- ☑ The storage was filled to 80% usable capacity. This is a good thing. Often, benchmark results are posted with a system at 25% capacity or less. This technique artificially enhances results as the disk actuators avoid full seeks across the entire surface of the drive (this practice is sometimes referred to as "short-stroking").

The results of the LeftHand ESRP v2.0 test are summarized in Figure Five.

Figure Five: ESRP Results Summary

Database I/O	
Average Database Disk Transfers/sec	19687.56
Average Database Disk Reads/sec	10832
Average Database Disk Writes/sec	8855
Average Database Disk Read Latency (ms)	17
Average Database Disk Write Latency (ms)	9.75
Transaction Log I/O	
Average Log Disk Writes/sec	4578
Average Log Disk Write Latency (ms)	2

What the Numbers Mean

- ☑ The system completed the certification test with an average database read response time of 17 milliseconds, which is less than the maximum allowed value of 20 milliseconds.
- ☑ A response time of 17 milliseconds means that with 50,000 active "heavy" Exchange users all sending and receiving e-mails during a normal eight hour shift, it would feel that the system was performing as expected.
- ☑ The only other published ESRP 2.0 result as of the writing of this report was for a much smaller number of users. While the configurations were quite different and are not directly comparable, ESG Lab finds it telling that a traditional dual node Fibre Channel attached storage array delivered less than 1/6th the number of Exchange users as the LeftHand clustered solution—*without* replication.

ESG Lab Testing

ESG Lab performed a series of tests to simulate a Microsoft Exchange workload using the industry standard IOMETER benchmarking tool. The testing was performed on the same 30 node cluster used for the ESRP test. Unlike the configuration used during the Exchange ESRP test, all 30 nodes were in the same location for this round of testing. IOMETER was used to provide quick estimate of the number of Exchange 2003 users that the system could support when the number of nodes in the cluster was increased from five to 30.

Two workloads ran simultaneously to mimic a real-world Microsoft Exchange environment: one to simulate Exchange 2003 database volume activity[3] and another to simulate log activity. Based on the Microsoft MMB3 guideline indicating that each "heavy" Exchange 2003 user sends 40 e-mails and receives 100 e-mails per day and performs one I/O per second, the results reported by IOMETER for the database volume were translated into the number of "heavy" Exchange 2003 e-mail users that the system can support as shown in Figure Six.

[3] 90% random 4 KB I/Os, 73% reads, 27% writes; 7% sequential reads, 3% sequential writes, 16 I/Os outstanding

Figure Six: IOMETER Simulated Microsoft Exchange Performance Scalability

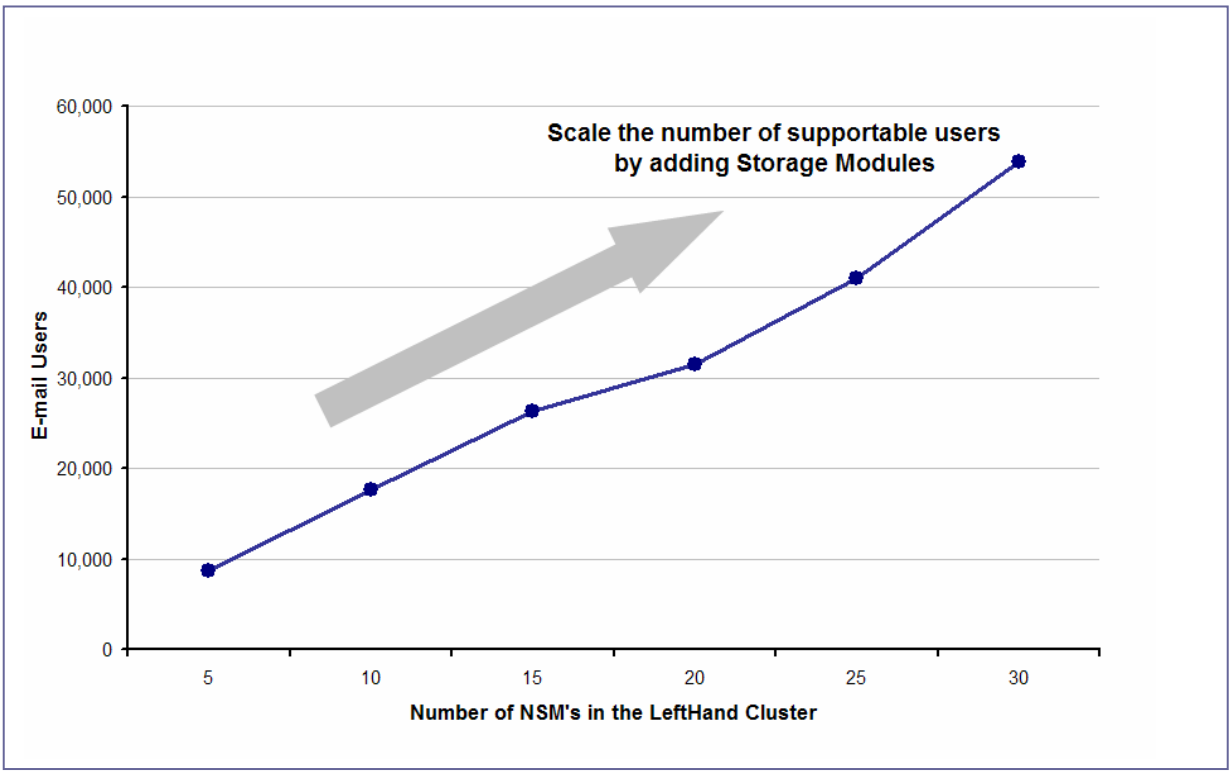

What the Numbers Mean

- ☑ The average response time reported by IOMETER was 6 milliseconds when servicing 53,000 simulated "heavy" Exchange 2003 e-mail users on a thirty node LeftHand cluster. This is well below the Microsoft guideline of 20 milliseconds.
- ☑ ESG Lab testing confirmed that the number of e-mail users that can be serviced from a single centrally managed LeftHand cluster increases nearly linearly as nodes, processors and drives are added to the cluster.

Database Performance Scalability

Performance Scalability in a storage environment is best measured with the metrics used by the applications organizations actually run. For an e-mail application, that measurement is the number of users or mailboxes a given system can support. For an interactive database application, the number of transactions per second that can be sustained during peak periods of activity is the measurement that matters.

ESG Lab Testing

Database performance was evaluated using the IOMETER workload generator and a simulated on-line transaction processing (OLTP) database workload with an I/O size of 8 KB. Tests driven from multiple clients were performed to verify the ability of a SAN/iQ storage cluster to deliver predictably scalable performance in a clustered scale-out server environment. The OLTP workload is random in nature and very disk intensive.

Tests were executed on cluster sizes from five nodes through thirty nodes. As shown in Figure Seven, performance scaled in a predictably linear fashion as nodes were added to the LeftHand cluster.

Figure Seven: IOMETER OLTP (8 KB) Results

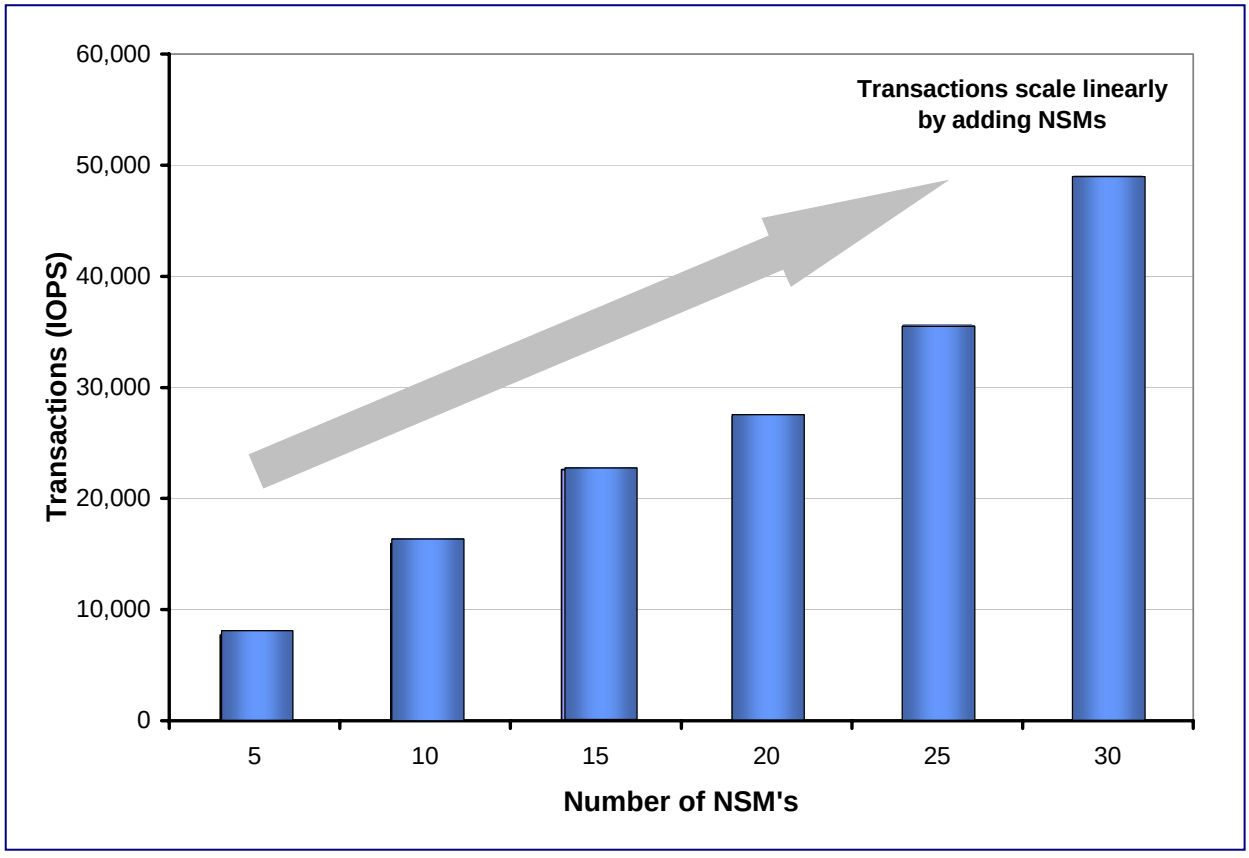

Streaming Media Performance Scalability

ESG Lab tested streaming media performance to evaluate the throughput capabilities of a LeftHand cluster and quantify the advantages of having data striped over all drives within the cluster. The IOMETER workload used during this phase of testing was created after observing the I/O that is generated when videos are viewed with Windows Media Player. Each stream was observed performing 32 KB sequential reads at a rate of 1.5 I/Os per second. This workload was used to simulate the effect of many users within a company accessing videos stored on a centrally managed LeftHand cluster.

Figure Eight: Streaming Media (32 KB sequential reads)

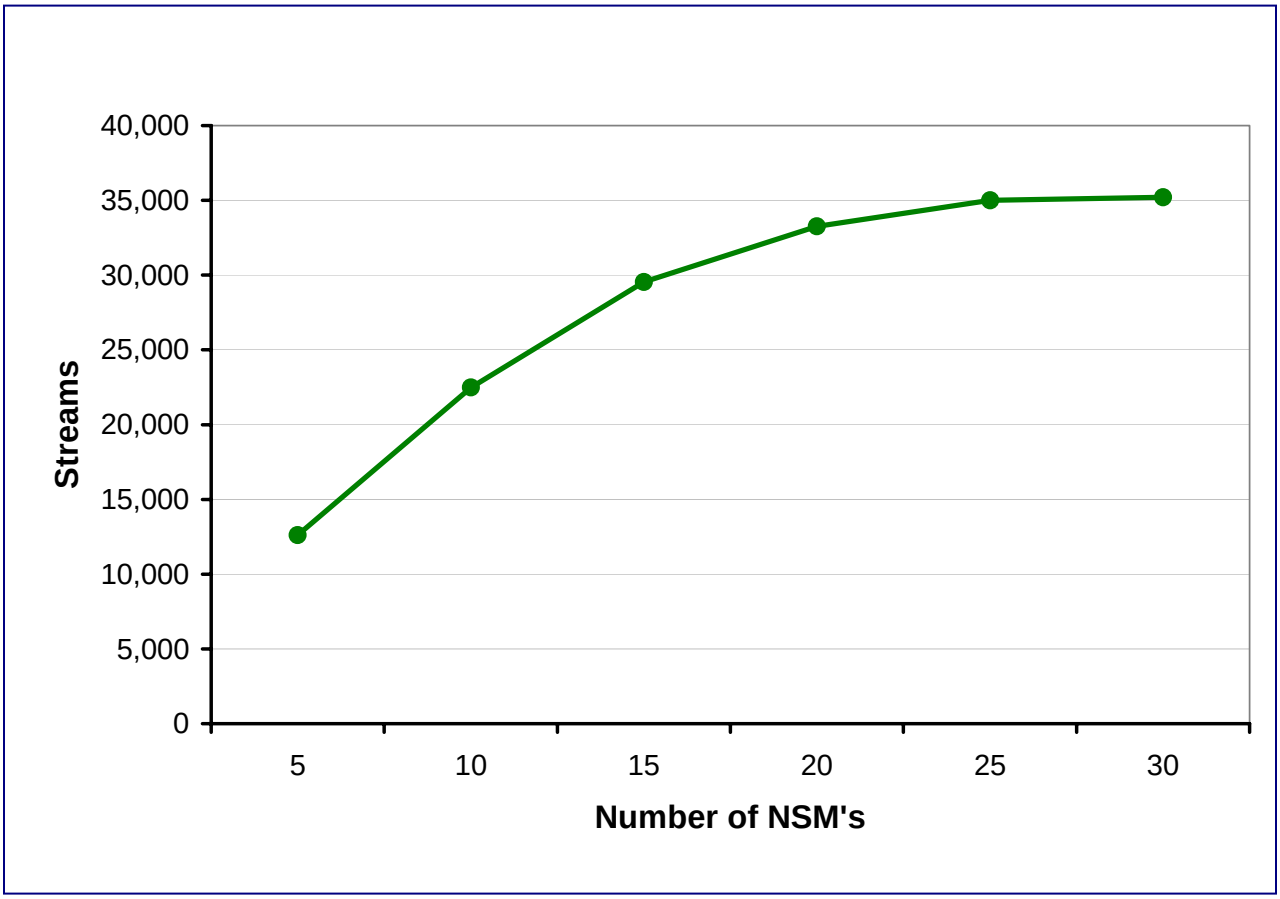

The green line in Figure Eight shows that the number of simulated video streams that can be supported increases as the number of nodes in the LeftHand cluster is increased. At twenty five nodes, network saturation was achieved at the client and adding additional nodes had little effect. However, adding more clients, aggregating Gb Ethernet channels or moving to 10Gb Ethernet would have removed this bottleneck. Despite the client limitations of the test, the LeftHand cluster showed it could support in excess of 35,000 users concurrently viewing videos stored on a 30 node cluster.

Why This Matters

ESG Lab found that LeftHand performance was extremely impressive even when compared to high-end FC-based storage systems. Additionally, LeftHand provides excellent performance with a variety of different workloads including e-mail, database and streaming data.

What is important to note is that because of LeftHand's N-way clustering capability, users can scale performance as needed. If dual-node storage systems hit a ceiling, typically there is no way to add more performance. LeftHand end-users can add more nodes as needed, giving them the flexibility to grow.

Enterprises have traditionally turned to FC SANs for scalability and performance, which comes with a premium in both cost and complexity. The LeftHand solution offers an iSCSI SAN that meets or exceeds the performance of traditional FC solutions by using commodity-based hardware and low-cost SATA and SAS drives. ESG Lab testing has validated that the efficiency and cost effective scalability of the LeftHand architecture can be used to meet the performance needs of real-world applications deployed in any environment, from modest to high-end and everything in between. Additionally, LeftHand provides compelling price/performance, regardless of your requirements.

Device Specific Module for MPIO

MPIO is the name given to multi-path availability software from operating system vendors that support rudimentary load balancing and failover for the connection between servers and network attached storage. Traditional host-based MPIO drivers typically use a round-robin algorithm when deciding which path to use next. MPIO drivers only know the paths that volumes can be found on. There is no visibility into the locality of data stored on those volumes. LeftHand has developed a plug-in for MPIO (the Device Specific Module or DSM) that provides data location intelligence at the client.

To understand the implications of this technology, consider Figure Nine, which shows how a traditional host-based MPIO driver accesses a clustered storage system with three nodes without DSM technology. An application is looking for data that may be located on cluster node A, B or C. A read request is sent to a node in the cluster–in this case node A. Node A consults a routing table to determine which node the IO should be routed to and forwards the IO to the appropriate node. This redirection causes an extra "hop" as the client's I/O is forwarded to the correct node.

Figure Nine: MPIO without DSM, using IO Forwarding

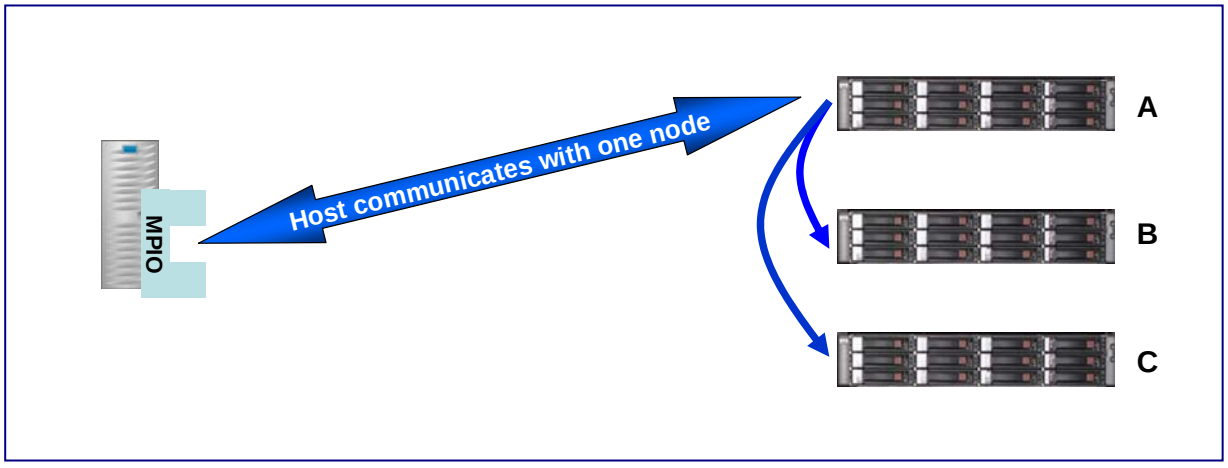

The LeftHand DSM solution is shown in Figure Ten. In this example, the same application is looking for the same data that may be located on cluster node A, B or C. The DSM plug-in to the MPIO driver contains a block map that allows the client to identify the node where the data is stored. The read request is sent directly to the node that has the data that the application needs, with no redirection and no 'hops'.

Figure Ten: MPIO with DSM

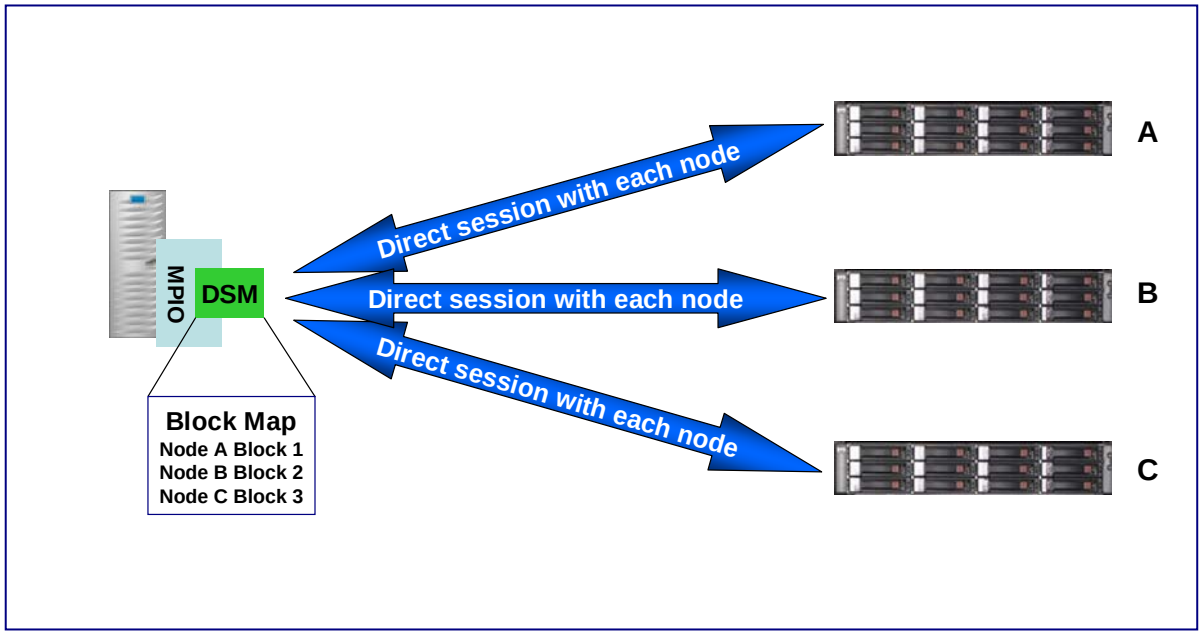

Note that in the LeftHand DSM implementation, a session is established between the client and each node. What this means for a busy system handling many I/Os is greater efficiency of operation. Eliminating unnecessary hops reduces the network traffic within the storage cluster and less traffic within the storage cluster means less work for each node to do. The bottom line is efficiency and performance that gets better as more nodes are added to the cluster.

ESG Lab Testing

ESG Lab compared the performance of the native Microsoft iSCSI MPIO driver to the same driver equipped with a LeftHand DSM plug-in. IOMETER workloads running on 15 Windows servers were serviced by a 30 node LeftHand cluster. A variety of workloads were tested, and in all cases, there was a noticeable performance advantage. As shown in Figure Eleven, a performance advantage of up to 100% was observed for bandwidth intensive 512 KB sequential read workload. This workload is similar to a nightly backup job.

Figure Eleven: The LeftHand DSM Performance Advantage (512k Sequential Reads)

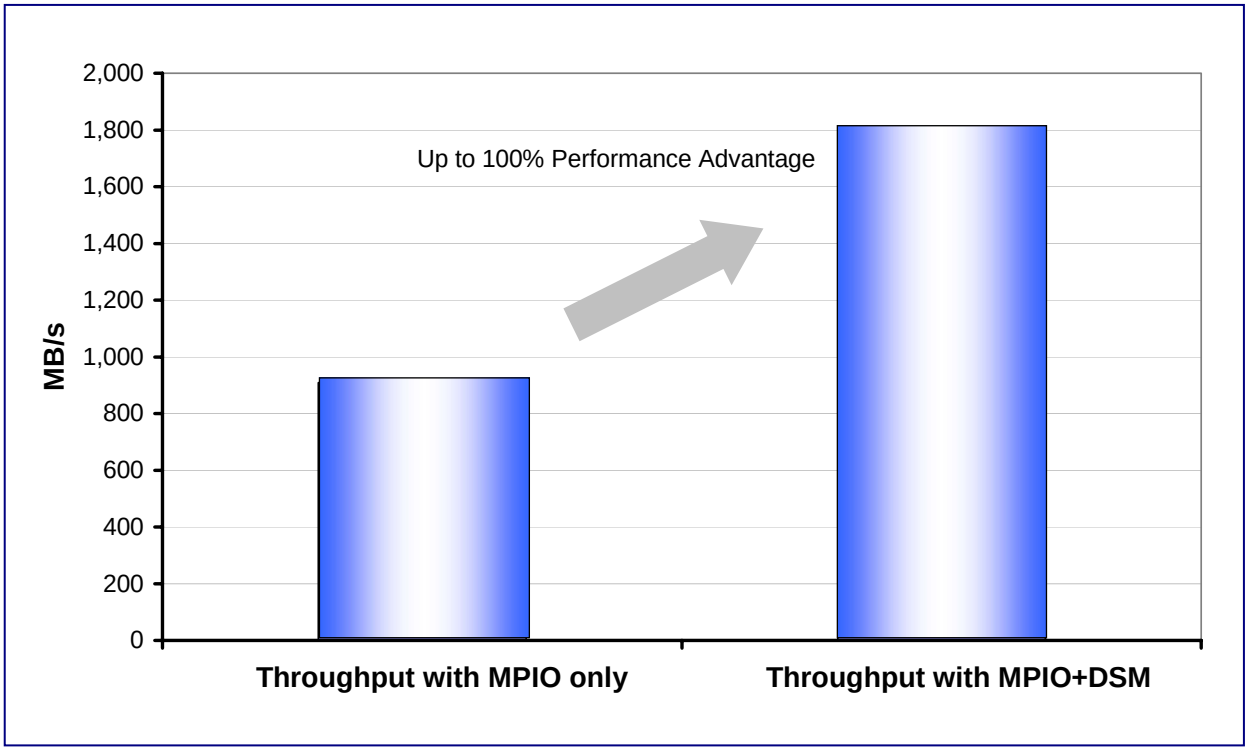

Why This Matters

ESG Lab has found that iSCSI performs as well as FC in most cases. The one area that is arguably in FC's favor is throughput for bandwidth intensive applications. The LeftHand DSM plug-in for Microsoft MPIO provides an impressive performance boost that not only rivals the throughput capabilities of traditional FC arrays, but can exceed it due to LeftHand's clustering capability.

Usability and Reliability

When ESG Lab tested the SAN/iQ storage software platform in June of 2006, one of our key suggestions was for LeftHand to move toward the industry-accepted tree navigation presentation model when navigating clusters, nodes, volumes and snapshots. That is exactly what LeftHand did. The new interface is clean and intuitive, providing easy access to all features and functionality, as shown in Figure Twelve.

Figure Twelve: The LeftHand GUI

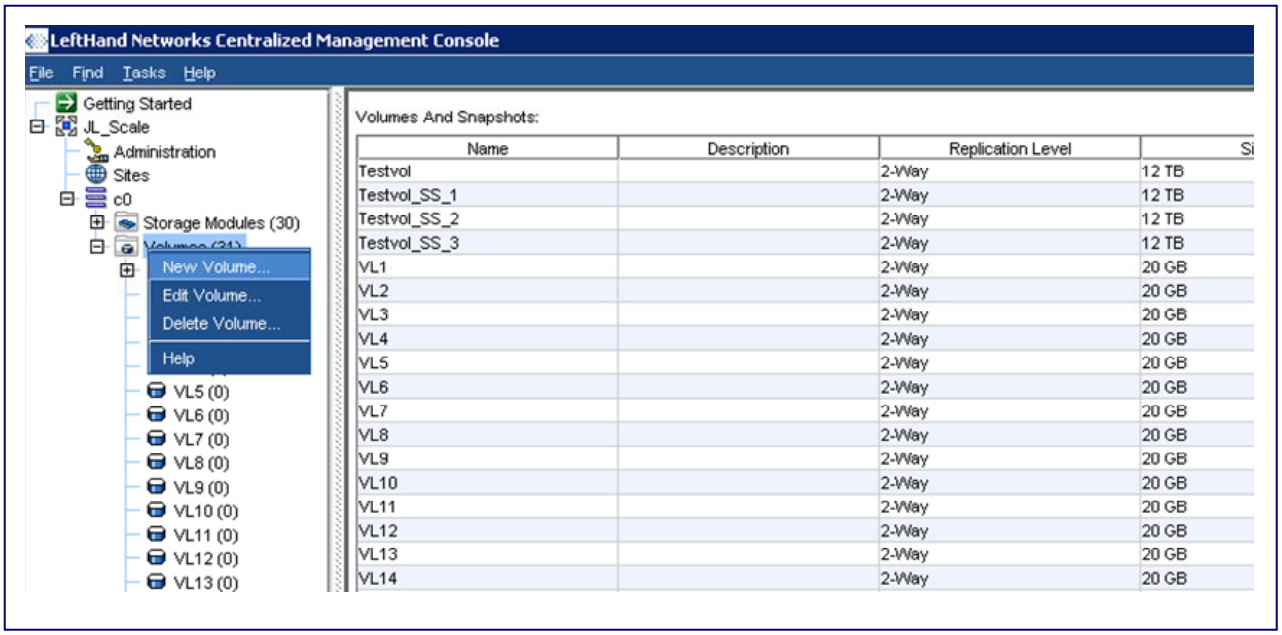

The LeftHand interface displays Sites at the highest level with Clusters, storage modules and volumes underneath. To create a new volume, ESG Lab right-clicked the "Volumes" branch and selected "New Volume" from the context sensitive menu, as shown in Figure Thirteen.

Figure Thirteen: Creating a Volume

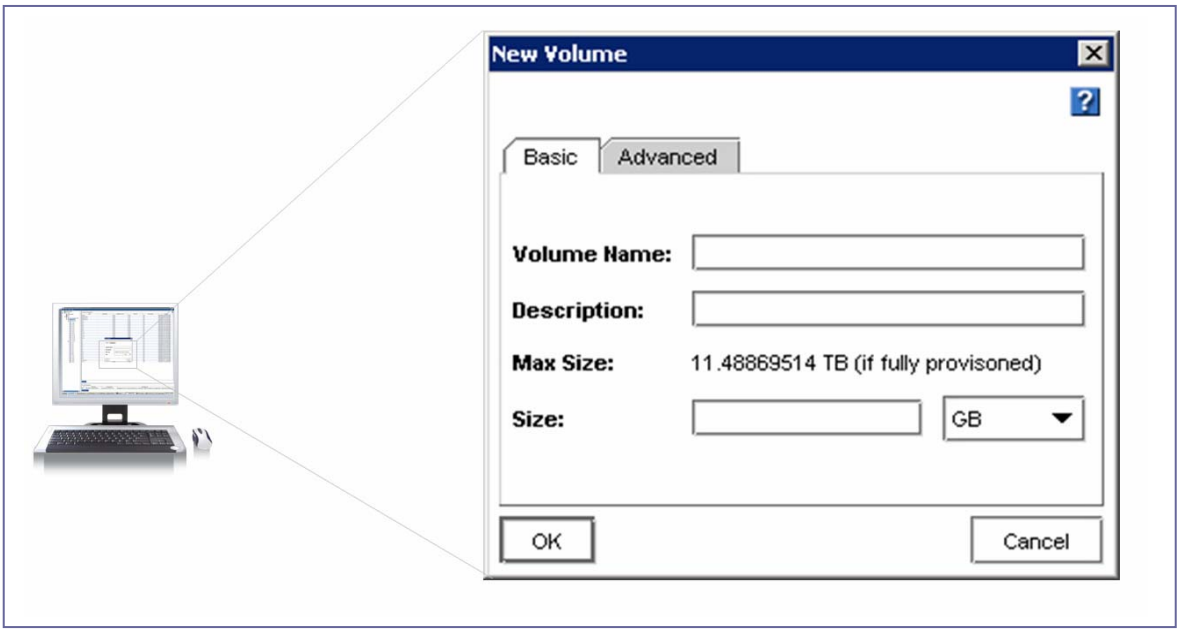

The only information required to create a volume is its name and size. The Advanced tab allows users to access additional optional features like replication levels and the volume type–full or thin provisioned. ESG Lab found thin provisioning in LeftHand's world to be extremely simple, with no calculations or decisions to be made for storage allocation. The size of the volume as entered in Figure Thirteen becomes its maximum size and storage is only allocated as data is actually written by applications or clients.

Another critical feature of enterprise-class arrays is the ability to make local point-in-time snapshots of production volumes. LeftHand's SAN/iQ Snap allows administrators to take multiple point-in-time snapshots of any volume in the system. Some examples of uses for snapshots are backup, fast recovery and file level undelete. Figure Fourteen shows ESG Lab creating a snapshot using LeftHand's SAN/iQ Snap feature.

Figure Fourteen: LeftHand SAN/iQ Snap

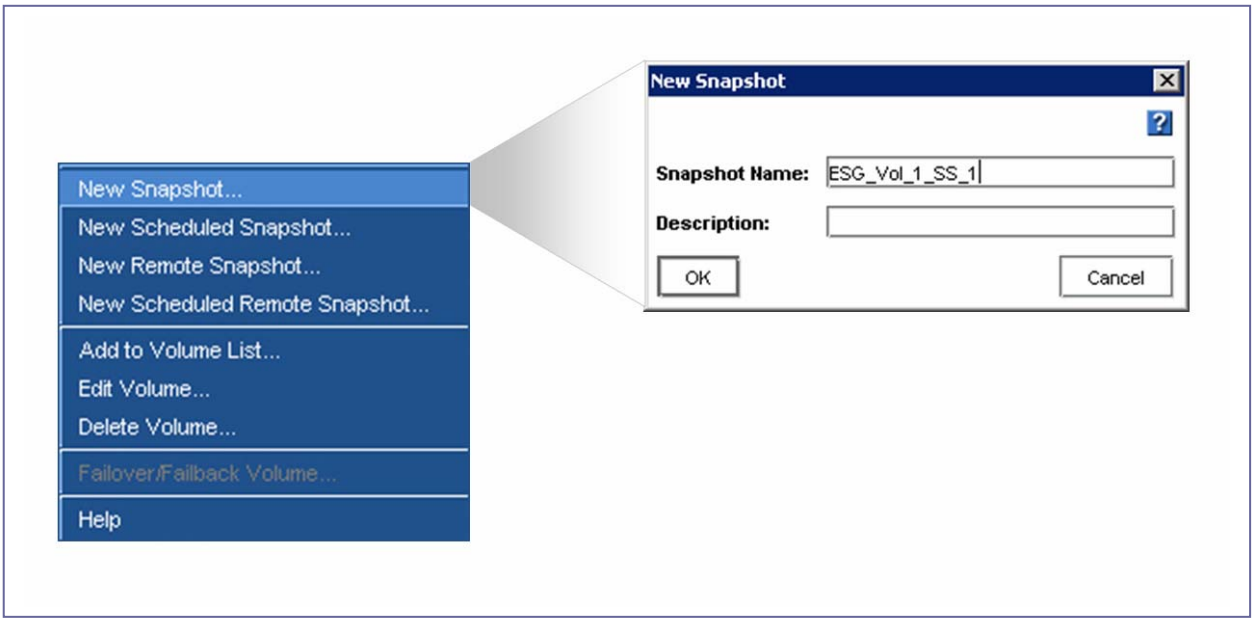

ESG Lab selected a volume to snapshot, right-clicked the volume name in the tree and selected '"New Snapshot" from the right-click pull down menu. When defining a snapshot, the only required information is the snapshot name. LeftHand's SAN/iQ Snap provisioning is completely virtual–snapshots are all thin provisioned and do not require pre-allocated or reserved space. This makes management of snapshot copies for backups, file level undelete and quick restores extremely simple.

Figure Fifteen: The Utilization Summary Screen

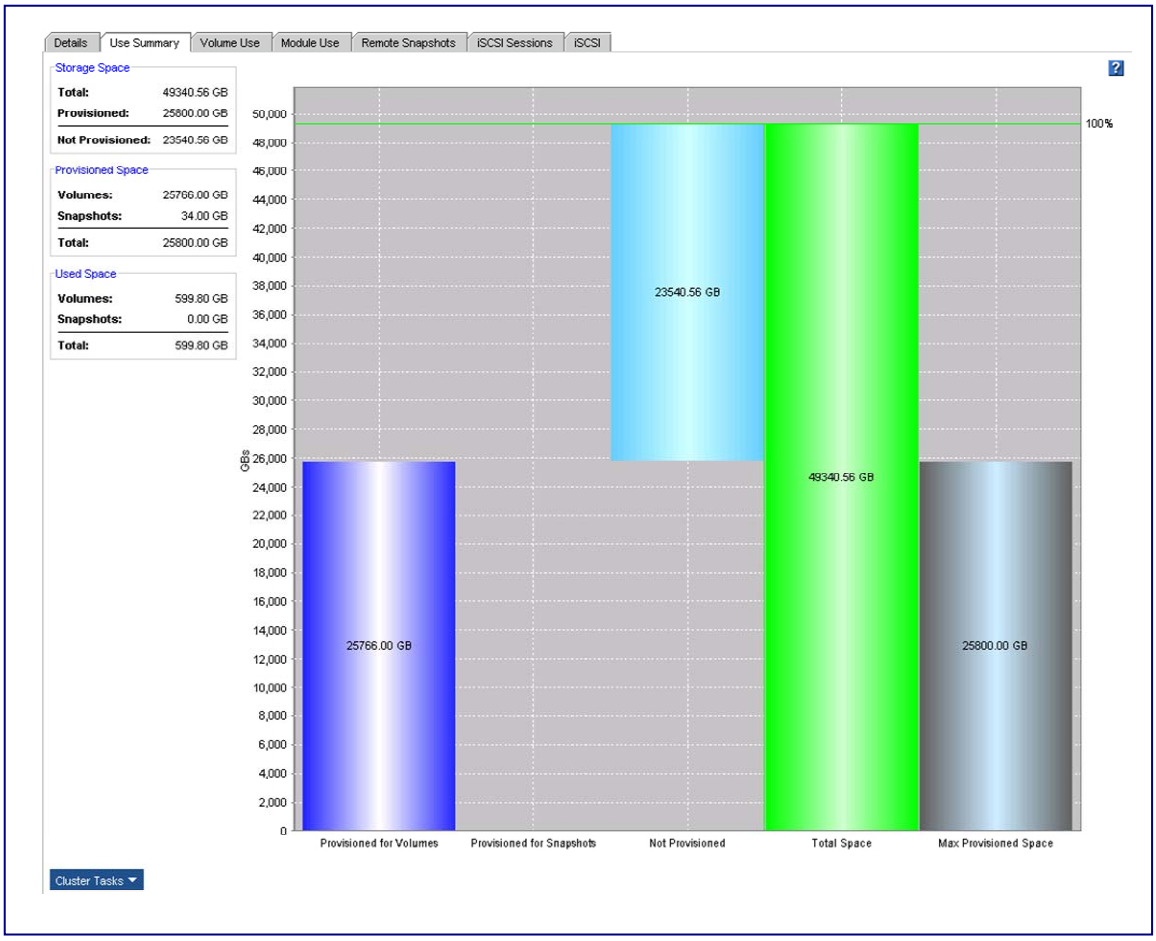

The utilization summary screen provides a simple, complete view into the utilization of the SAN/iQ cluster. In Figure Fifteen, each column represents a different aspect of storage allocation. The dark blue bar on the left represents the total capacity provisioned for volumes, including Full and Thin provisioned volumes. The empty space between the dark and light colored bars shows how much space is provisioned for snapshots. ESG Lab created only one snapshot, so this column is almost empty. The tall green bar represents the total usable capacity in the cluster and the gray bar at the right represents the total provisioned space, including volumes and snapshots.

Why This Matters

The biggest problem with managing storage environments of all sizes is the complexity of the storage systems themselves. Most leading storage systems are complicated machines that require great expertise and ongoing manual administration. Administrators with this level of experience are hard to come by and therefore costly. With no host bus adapters or specialized switches required, configuring and managing an iSCSI storage area network is inherently easier than configuring a Fibre Channel storage network. LeftHand Networks has built upon this simplicity with a user interface that can be used to quickly and easily perform complex storage management tasks with little or no impact to existing applications. ESG Lab believes that the ease of management of an iSCSI-attached SAN/iQ system can be used to reduce the total cost of storage ownership substantially compared to legacy Fibre Channel solutions.

ESG Lab Validation Highlights

- ☑ ESG Lab audit and testing confirmed that the 100 TB LeftHand SAN/iQ system was quite capable of handling 50,000 Microsoft Exchange 2007 users while synchronously replicating.
- ☑ SAN/iQ performance scaled in a nearly linear fashion as the number of nodes in the cluster was increased from five to 30 nodes for common business application workloads, including e-mail and on-line transaction processing.
- ☑ A performance boost of up to 100% was observed on a 30 node cluster due to the use of LeftHand's innovative DSM Plug-in for the Microsoft iSCSI MPIO driver.
- ☑ ESG Lab found that the new tree style GUI is powerful, intuitive and easy to navigate,
- ☑ ESG Lab re-confirmed the reliability of the system as presented in the ESG Lab Validation report published in May 2006[4] by forcibly powering down a node while the system was servicing a simulated Microsoft Exchange workload. As expected, the system kept running with no issues.

Issues to Consider

- ☑ The welcome page of the SAN/iQ GUI is rather plain and would be greatly improved with the addition of system wide statistics similar to those seen in the Utilization Summary screen.
- ☑ Trending and utilization graphs showing utilization for the entire system over time would be a useful addition to the SAN/iQ GUI.
- ☑ While ESG observed a SAN/iQ cluster connected to servers using a 10 Gigabit Ethernet network through a Force10 Networks switch in the lab at LeftHand Networks, these systems were not yet ready for stress testing by ESG Lab. While LeftHand assured ESG that the SAN/iQ architecture is 10 GigE ready, customers who may be considering 10 GigE iSCSI for performance hungry applications are advised to proceed with caution as the 10 GigE hardware ecosystems that SAN/iQ software relies upon approach maturity.

ESG Lab's View

ESG Lab believes the LeftHand solution is still one of the best SAN-based storage systems we've ever validated. LeftHand has a highly scalable clustered architecture that simplifies management and allows medium to large Enterprise customers to start at the level of capacity and performance required today and grow their environments on demand. Additionally, it is easy to use and manage, provides advanced features such as snapshots and remote replication and offers innovative and valuable technology such as thin provisioning. Currently, customers can stretch their clusters to create Multi Site SANs and a future upgrade will offer support for more than two sites in addition to advanced performance optimization and management features. We have seen storage systems that scale in this fashion with NAS and CAS products, but in our opinion, LeftHand is a leader in SAN attached true N-way clustered storage. ESG has long been a proponent of scalable clustered storage and we believe it will become the dominant approach due to the compelling value it brings.

LeftHand N-way clustering, thin provisioning, virtualized pools of storage and snapshots are all examples of powerful and valuable internal storage virtualization technologies that result in greater ease of management. Adding iSCSI support, intuitive management software and Windows integration, the result is an efficient and easy to use system that requires minimal manual support. IT shops are not burdened by LeftHand as with other complex apparatus.

ESG Lab was very impressed with LeftHand's performance, as it was not scalable with only one type of workload, but multiple I/O types. ESG Lab found that LeftHand performed well with e-mail, database and streaming media. Additionally, we were impressed with the near linear growth in performance as we added more nodes to the cluster. On top of this, LeftHand's performance was improved even further through leveraging its DSM plug-in for the Microsoft iSCSI MPIO driver, where we validated a 100% performance improvement with bandwidth intensive workloads. ESG Lab believes that LeftHand has one of the best price/performance SAN storage systems that we've tested first hand.

[4] http://www.enterprisestrategygroup.com/ESGPublications/ReportDetail.asp?ReportID=629

Appendix

ESG Lab Test Configuration

LeftHand SAN/iQ Software	Version 7.0
LeftHand Network Storage Module	HP DL320, 2.67 GHz Intel Dual-Core Xeon Processor, 2GB RAM, Dual GbE configured in Adaptive Load Balancing (ALB)
Number of Network Storage Modules	Thirty
Hard Drives per Network Storage Module	Twelve 300GB 15,000 RPM SAS drives
RAID configuration	RAID 10
Replicated configuration	2-way replicated volumes (Network RAID 10)
IP SAN Switch	Force10 E600
Gigabit Ethernet Modules	Two 48 port GbE Line cards, Jumbo Frames not enabled.
IOMETER Workload Generator Clients	Dell 1850 Dual 2.8GHz Zeon 2GB RAM Dual GbE NICs with ANS Teaming
Number of IOMETER Workload Generator Clients	Fifteen
Microsoft iSCSI initiator – 64 bit	2.03