

LAB VALIDATION REPORT

HP StorageWorks ExDS9100 **Extremely Scalable, Affordable, and Easy to Manage File Storage**

By Tony Palmer
With Terri McClure

March, 2009

Table of Contents

| | |
|--|-----------|
| Table of Contents | i |
| Introduction | 1 |
| <i>Background.....</i> | <i>1</i> |
| ESG Lab Validation..... | 4 |
| <i>Simplicity and Manageability.....</i> | <i>4</i> |
| <i>Scaling Capacity and Performance.....</i> | <i>8</i> |
| <i>Availability.....</i> | <i>15</i> |
| ESG Lab Validation Highlights..... | 17 |
| Issues to Consider | 17 |
| ESG Lab's View | 18 |
| Appendix..... | 19 |

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of 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 sponsored by Hewlett Packard.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. Copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482.0188.

Introduction

The accelerating rate of growth in the volume of unstructured file data organizations need to share, manage, and protect has become a serious challenge for IT managers—particularly those who need to provide online access to shared files for revenue generating workflows and processes. In this Lab Validation, ESG Lab explores how the HP StorageWorks 9100 Extreme Data Storage System leverages standards-based servers and storage, along with clustered file services, to create an extremely scalable, easy to manage networked storage solution designed to reduce data center footprint, administration requirements, and acquisition costs.


Background

The management of file-based or “unstructured” content (i.e., multimedia files, web pages, office productivity documents, etc.) has become one of the most pressing and persistent challenges facing today’s IT organizations. In fact, ESG research indicates that the vast majority of corporate digital assets are stored in unstructured files. Unstructured content—including digital images, audio, and video files—accounted for 77% of global digital archive capacity in 2007 and is expected to constitute the bulk of digital assets for the foreseeable future.¹ This massive amount of file data is creating demand for a new type of file storage solution that can scale bandwidth, performance, and capacity under a single system, reducing management costs as the data under management continues to grow. The term “scale-out NAS” describes systems designed from the ground up to meet this need. While scale-out NAS has seen most of its early adoption in media and entertainment, Web 2.0, and high performance computing, out-of-control file data growth, larger file formats, and the commercial adoption of file intensive applications has brought the requirement for scale-out systems squarely into traditional data centers.²

The growth of file-based data, as well as the shifting nature of the files to richer formats, has left enterprise data centers bursting at the seams. These factors are leading data center managers to consider taking a new approach to storing and managing file-based data while simultaneously pushing NAS vendors to introduce entirely new architectures. For managing growth and meeting the performance characteristics required by richer file data, scale-out NAS is gaining traction, as seen in Figure 1. Forty-nine percent of IT managers and professionals surveyed by ESG indicated that they had deployed or were planning to deploy scale-out NAS in their environments.

FIGURE 1. SCALE-OUT NAS INTEREST

Please describe your organization’s usage of or interest in scale-out NAS systems. (Percent of respondents, N=338)



¹ Source: ESG Research Report, *Digital Archiving Survey*, November 2007


² Source: ESG Research, *2008 Enterprise Storage Survey*, November 2008

It is no surprise that we often hear data center managers say that they love their first NAS appliance—and curse their tenth, or worse yet, their hundredth. Traditional file storage solutions are designed to be monolithic—all storage sits behind one or two file server heads and capacity is scaled by adding storage behind those heads. When the storage limit is hit or additional performance is required, a new monolithic system is installed. There is no way to share the workload between the systems and migrating directories or files between systems means remapping and remounting for each and every client requiring access. There is no economical way to scale performance without significant additional cost and complexity.

HP StorageWorks 9100 Extreme Data Storage

The HP StorageWorks 9100 Extreme Data Storage System (ExDS9100) is a scale-out NAS solution aimed at addressing concerns about massive file data growth and management challenges with a capacity-optimized implementation of the PolyServe Clustered file system. This solution is deployed on a massively scalable platform built from HP BladeSystem c-Class blade servers and multiple large SAS arrays.

FIGURE 2. THE HP EXDS9100



ExDS9100 systems are built upon storage elements called capacity blocks and server elements called performance blocks. A single capacity block is composed of a dual redundant SAS controller and two capacity drawers containing a total of eighty-two 1 TB midline SAS drives with RAID-6 data protection. Capacity blocks are divided into eight 10 disk RAID groups with two hot spares for a total usable capacity of 64 TB per capacity block. Each performance block consists of an HP BL460c blade server with a SAS mezzanine card running ExDS9100 file serving and management software on a Linux operating system. A single ExDS9100 can house up to ten capacity blocks for 820 TB of raw storage in just two floor tiles. Up to 16 performance blocks use that storage capacity to present clustered file systems via standard NFS and HTTP protocols. CIFS support is available as an optional protocol enabled by HP.

Since the ExDS9100 runs a Linux-based operating system, many applications—such as those for streaming media—can be run directly on the performance block. HP calls this unique feature ‘application co-hosting.’ Not all applications are able to take advantage of this powerful feature, but if careful attention is paid to specific application requirements, users can unlock huge potential for decreased complexity, increased performance, and reduced cost. By layering the application tier directly on the storage system, administrators can potentially eliminate server tiers from their application architecture. HP recommends that users discuss their application requirements with their account team to understand the technical requirements.

Additional features and benefits of the ExDS9100 solution include:

- **Extreme Scalability:** The ExDS9100 scales to 820 TB in just 13 square feet of floor space while driving up to 3.2 GB/sec of performance using independent, very dense capacity blocks and high performance blade server-based performance blocks. Performance and capacity scale independently, online, and seamlessly.
- **Unified Management:** The ExDS9100’s simple management interface provides a single point for configuration, control, and monitoring of the entire system, enabling a single administrator to manage petabytes of storage.
- **Extreme Availability:** Rolling upgrades to servers and storage mean the system will not have to be taken down for maintenance. RAID-6 provides an extra layer of data protection.
- **Extreme Affordability:** The ExDS9100 comes fully integrated and configured. All hardware and software required to get up and running quickly is included; no additional software or services are required. Online maintenance and upgrades lower operational costs, eliminating offline windows. The ExDS9100 also provides the ability to host applications directly on the system, simplifying an organization’s overall application architecture.

This ESG Lab Validation report examines the ease of implementation and management of the HP ExDS9100 in addition to its performance and capacity scalability as well as the acquisition and operational cost benefits of the platform.

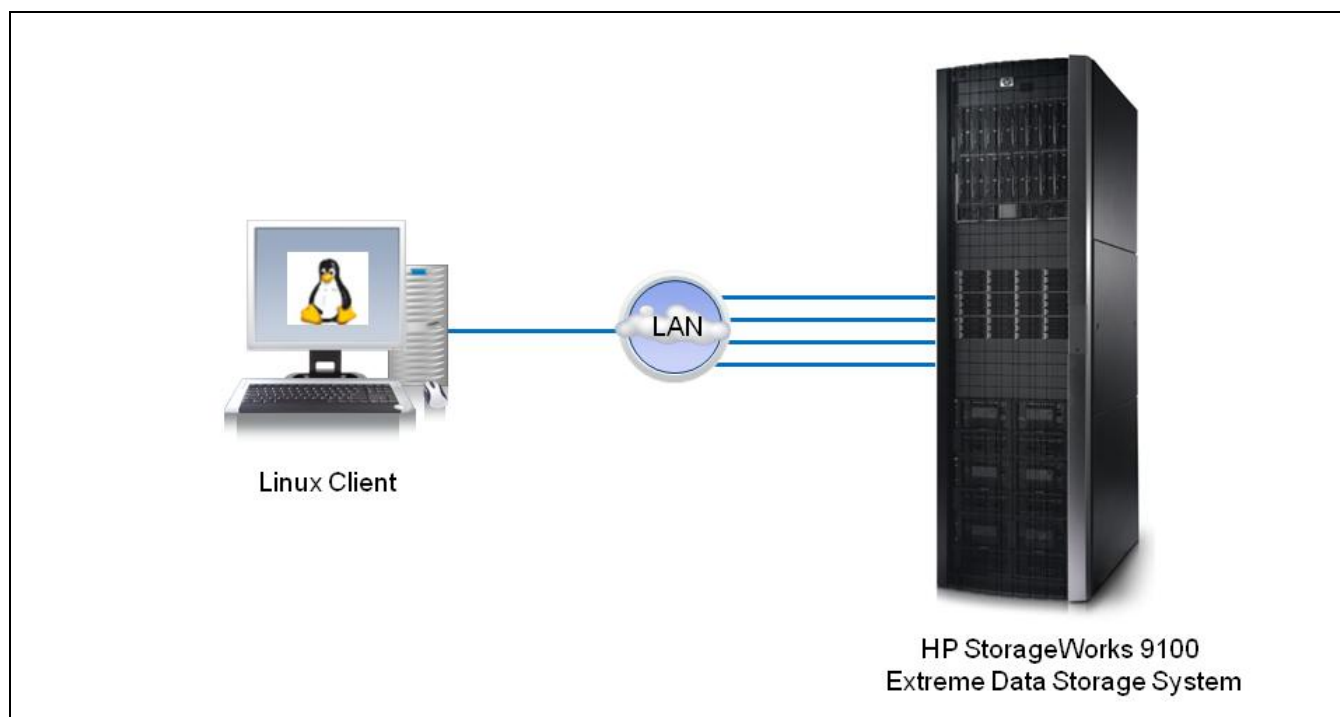
ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of the HP ExDS9100 at HP's Executive Briefing Center in Cupertino, CA. Testing began with a physically installed and powered up HP ExDS9100. ESG lab walked through the initial configuration process a typical administrator would experience after an ExDS9100 had been delivered and installed by HP professional services.

Simplicity and Manageability

Figure 3 shows a high level view of the test bed used by ESG Lab. An ExDS9100 configuration containing 246 TB of raw capacity and a client machine running Linux were attached to a Gigabit Ethernet LAN.³

FIGURE 3. THE EXDS9100 TEST BED



While the underlying components of the system each have their own management and monitoring interfaces, the ExDS9100 is configured and managed with a simple command line interface. Many configuration activities are carried out with just a single command. To further simplify configuration and management, the ExDS9100 offers a single capacity-optimized storage layout that enables HP to pre-format the back-end storage at the factory and offload storage management tasks from ExDS administrators.

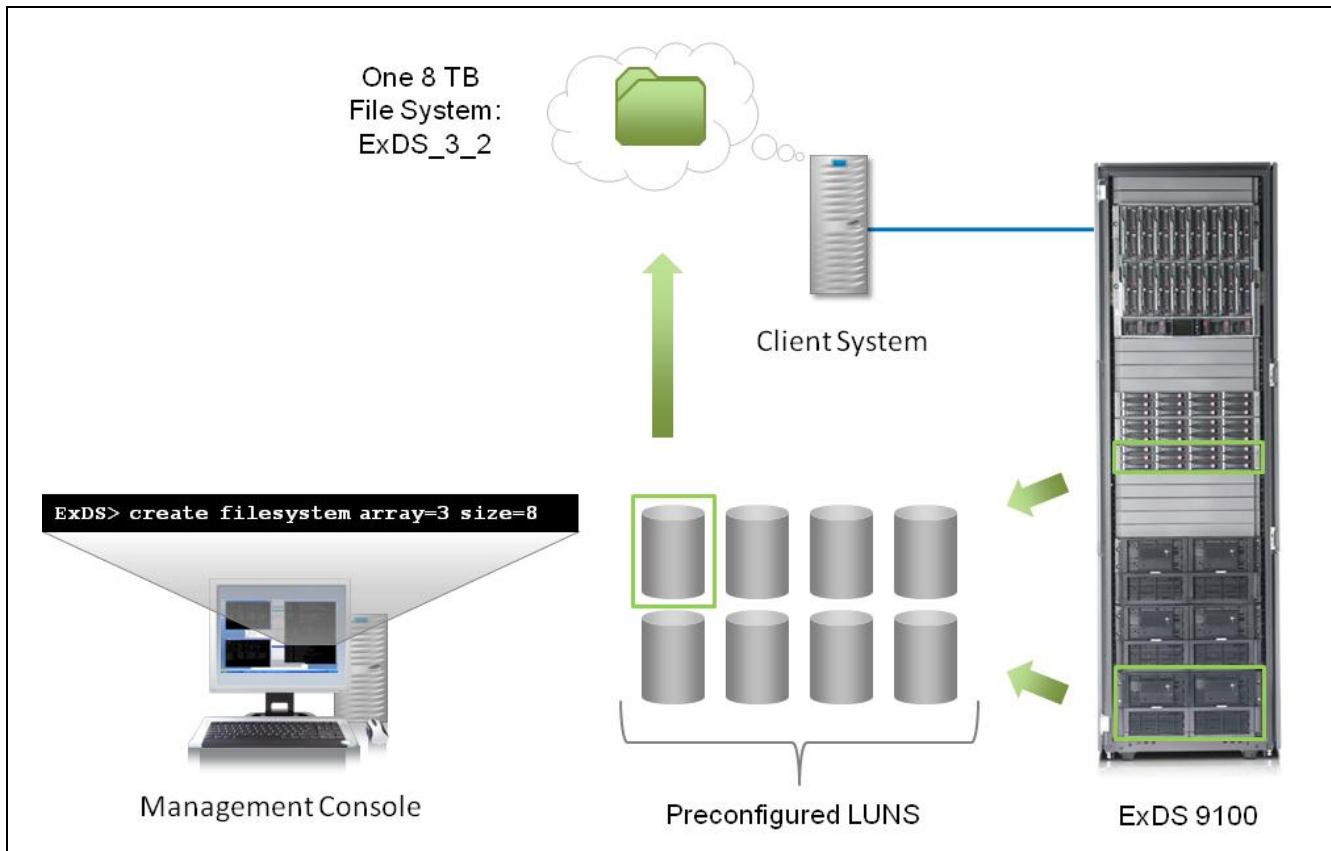
ESG Lab used the HP CFS (Clustered File System) GUI and the HP Onboard Administrator for the purpose of illustration in this report, but all configuration and management activity was executed from the ExDS Manager Command Line Interface.

³ Configuration details are listed in the Appendix.

ESG Lab Testing

The ExDS was powered on and configured with IP addresses by HP professional services, just as it would be for a new customer. ESG lab used the PuTTY ssh client to establish a secure connection to ExDS1, a blade server in the system. Once connected, the ExDS manager CLI was launched with the 'exdsmgr' command. Any blade server may be used to access the ExDS manager console. As shown in Figure 4, ESG lab first created an 8 TB file system using the 'create filesystem' command in the ExDS Manager.

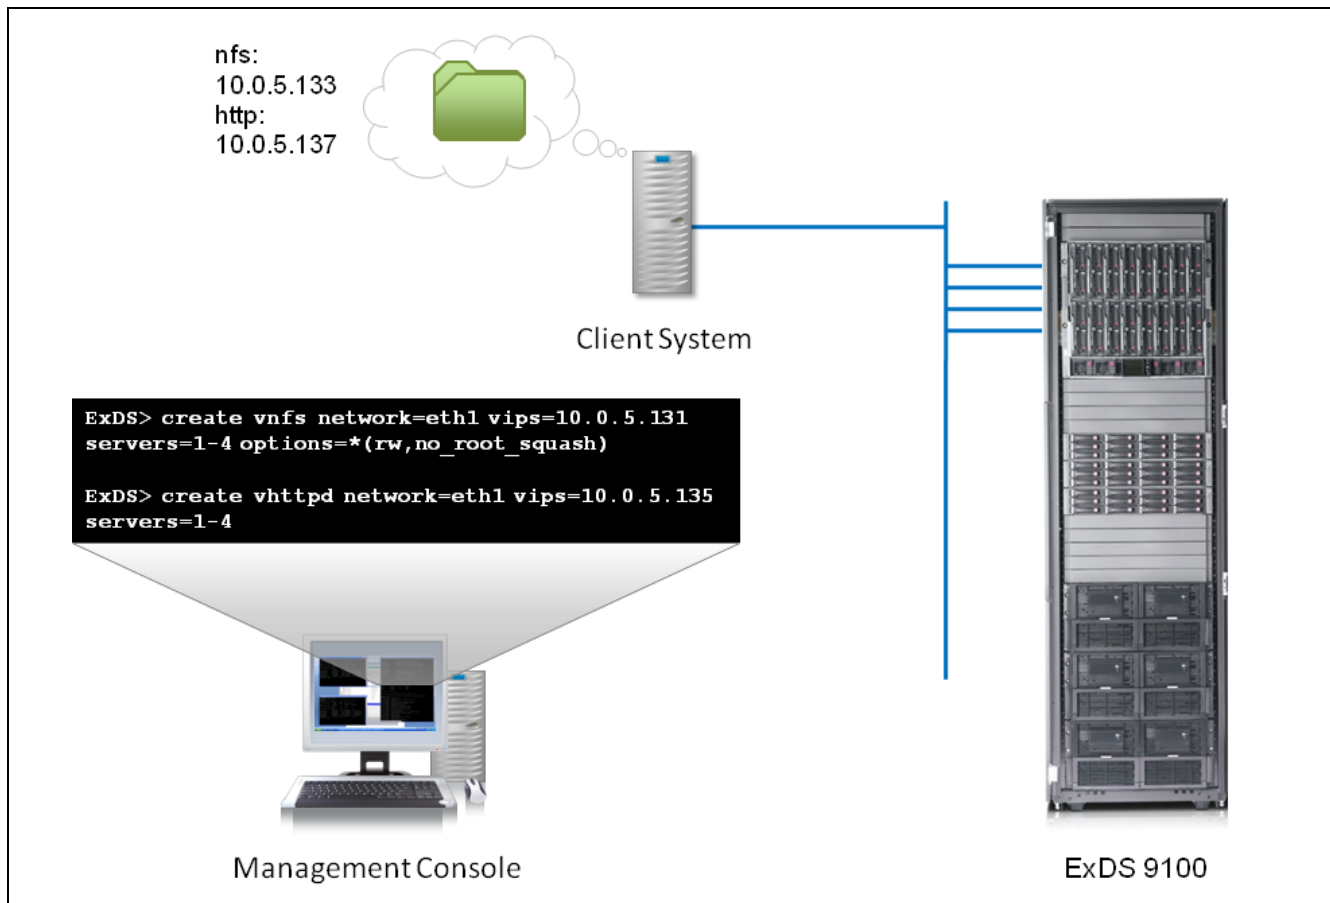
FIGURE 4. CREATING A FILE SYSTEM



The 'array' command line parameter was used to define which capacity block should be used to create the file system and the 'size' parameter was used to specify the capacity of the new file system. In this example, an 8 TB file system on array three was created. The command completed in just over one minute.

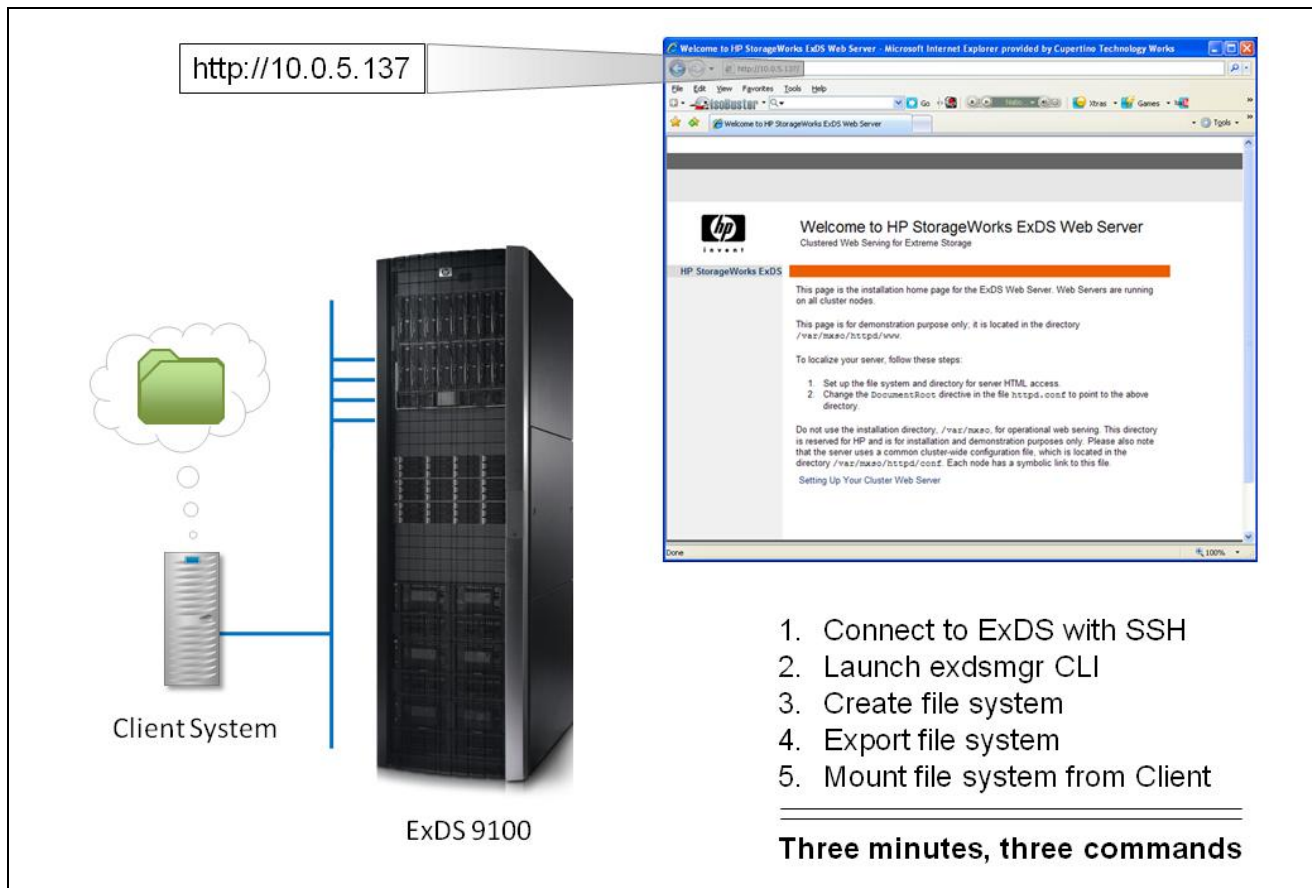
Next, the filesystem was exported via NFS and HTTP protocols using the 'create vnfs' and 'create vhttpd' commands, shown in Figure 5. These commands are very similar; the only mandatory parameters are the virtual IP addresses of the share and the list of servers that will export this file system.

FIGURE 5. EXPORTING A FILE SYSTEM



The commands executed immediately and with no delay, the newly created file system was exported by all servers in the ExDS over NFS and HTTP. Finally, ESG Lab mounted the file system over NFS and pointed a web browser at the virtual IP created in the previous step.

FIGURE 6. CONFIGURING THE EXDS9100



With just three commands and in less than three minutes, ESG Lab was able to configure, export, and mount an 8 TB file system from the ExDS9100.

Why This Matters

ESG research found that among organizations using scale-out NAS solutions, more than half indicated ease of management and the need for faster provisioning times as among the most important factors driving them to a scale-out NAS solution.⁴ Installing and configuring traditional large-scale file systems can be extremely complex and time consuming. Whether it's a clustered file system in front of SAN-attached storage or a number of traditional dual controller NAS systems glued together with virtualization software, time and money can be wasted trying to deploy legacy systems for large-scale, data-intensive applications. The HP ExDS9100 is a massively scalable, self-contained, pre-configured file services appliance. Once the system was powered up and attached to the network, ESG Lab simply sat down, logged in, and the ExDS was providing scale-out file services over multiple protocols three minutes later.

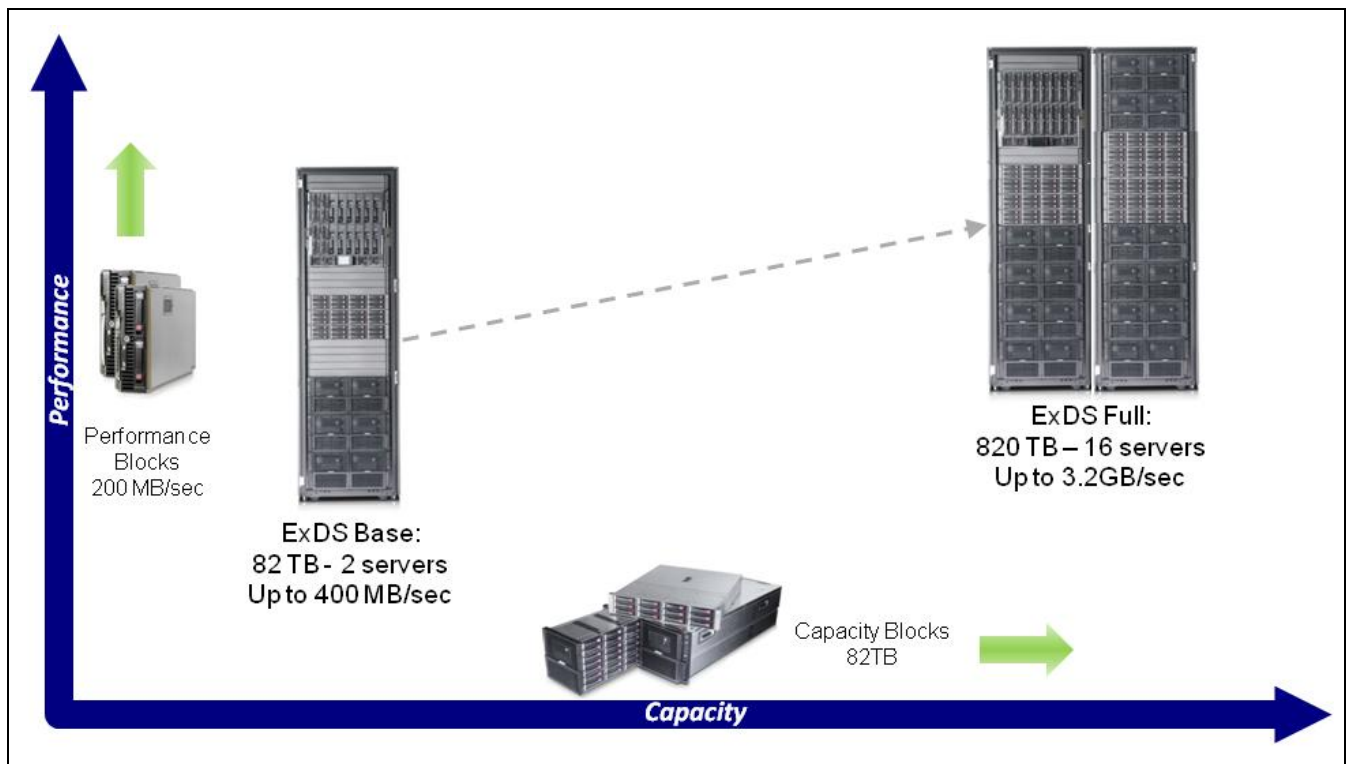
⁴ Source: ESG Research, 2008 Enterprise Storage Survey, November 2008

Scaling Capacity and Performance

The ExDS is a file services platform optimized for capacity, with a pre-defined configuration and layout that offers high system throughput and near linear scaling. Costs are kept low by utilizing 1 TB midline SAS drives, which are SATA disk platters with SAS electronics and connectivity. Midline SAS drives offer the best of both worlds for a capacity optimized system: the high capacity and lower cost of SATA combined with the connectivity, availability, and functionality of SAS.

An HP ExDS9100 file system can scale capacity and performance independently as shown in Figure 7. Performance blocks or capacity blocks can be added hot and online to meet changing business requirements. Blocks are installed in the system with a single command, eliminating the management of disk drives and RAID groups.

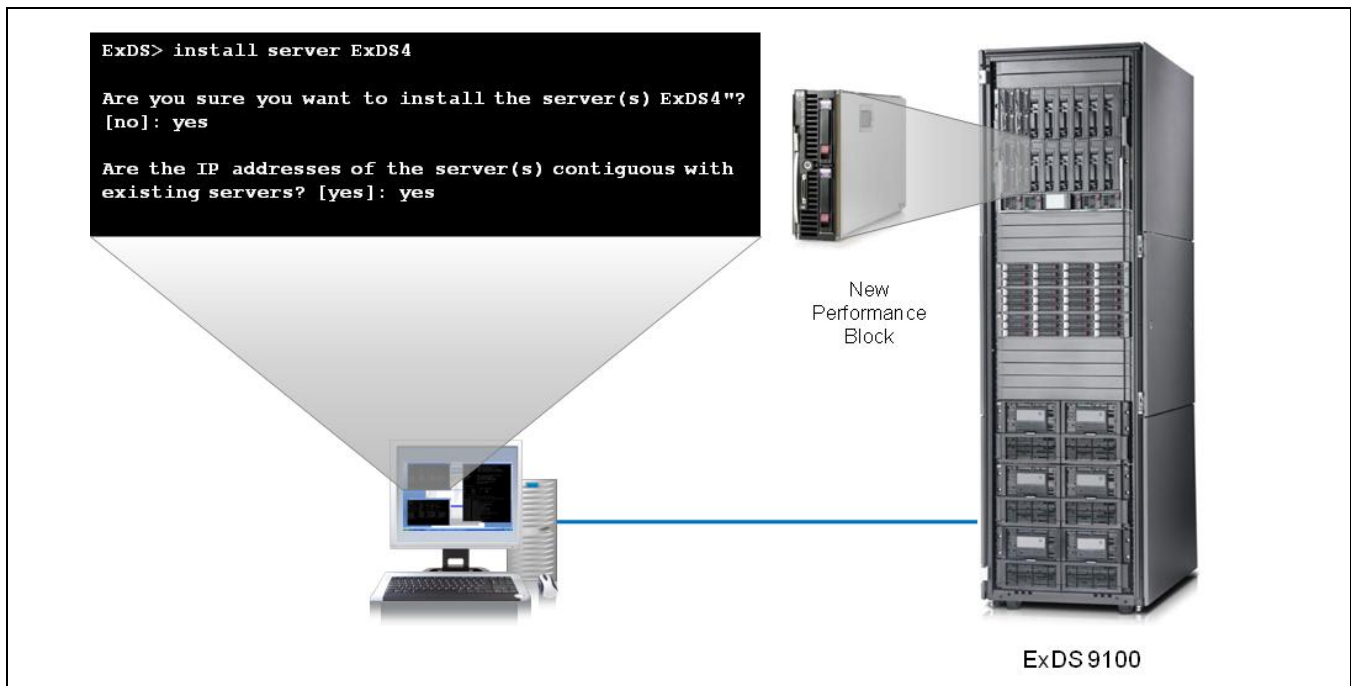
FIGURE 7. SCALING THE EXDS9100



ESG Lab Testing

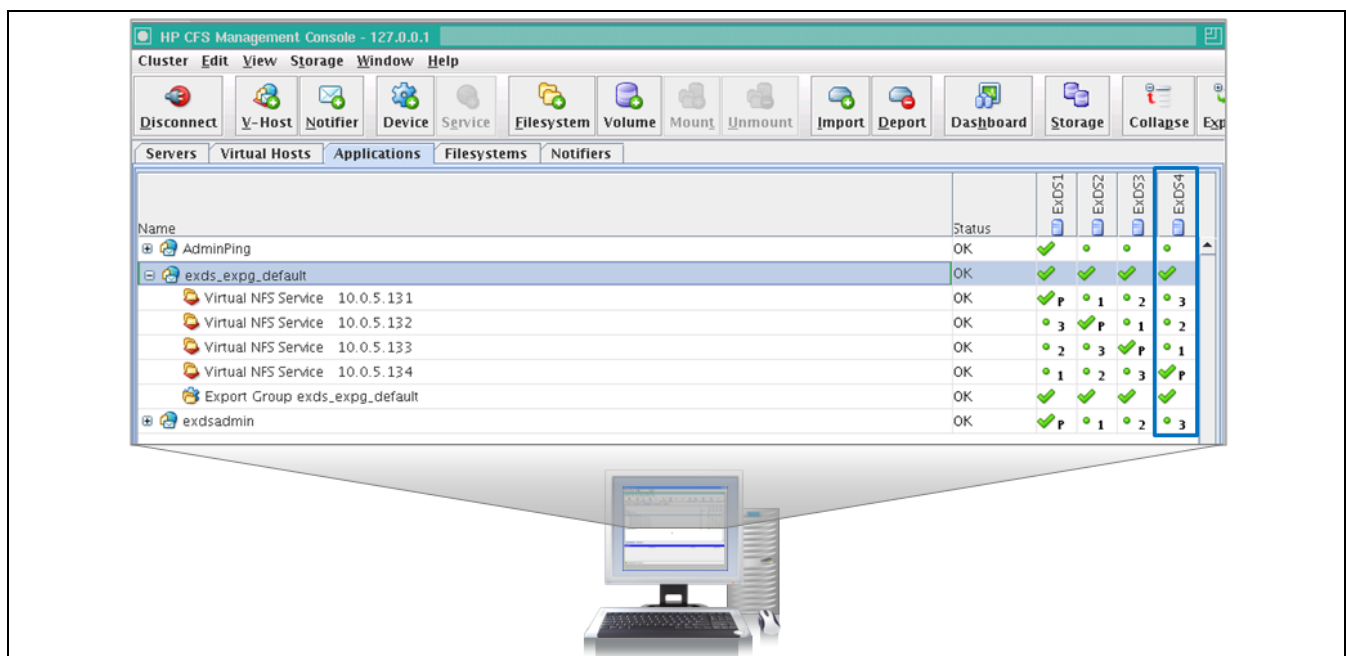
ESG Lab tested the independent online scalability of the ExDS9100 by installing a new performance block and auditing the procedure for installing a new capacity block in a running system. The process could not have been simpler. After inserting the new performance block (a BL460c blade server) into an empty slot in the integrated HP c7000 chassis of the ExDS, ESG Lab confirmed that it was receiving power and issued the 'install server' command, as seen in Figure 8.

FIGURE 8. INSTALLING A PERFORMANCE BLOCK




This kicked off a completely automated process that installed and configured both the operating system and HP CFS software, then joined the new performance block to the cluster with no further user interaction required. The entire process took about 15 minutes. When exdsmgr indicated that the installation was complete, ESG Lab verified that the server was a live member of the cluster using the HP CFS GUI, seen in Figure 9, and by mounting a file system from the newly installed server, ExDS4.

FIGURE 9. A PERFORMANCE BLOCK IS ADDED TO THE CLUSTER



Next, ESG Lab examined the procedure for upgrading the system's capacity. Storage upgrades are performed by adding additional capacity blocks to an existing system while it is online. Capacity blocks are pre-provisioned at the factory with a standard LUN layout; initialization on all LUNs is also completed at the factory. All an administrator needs to do is create file systems as required. One or more storage blocks can be added at a time. New capacity blocks are shipped with a five page glossy install guide, shown in Figure 10, which clearly lays out the procedure used to install and configure a new capacity block.

FIGURE 10. INSTALLING A CAPACITY BLOCK



Once the new block is racked, cabled, and powered up, a single command discovers the storage in the capacity block and adds it to the pool of available storage in the cluster. At this point, the storage is available for the administrator to create and export file systems.


ESG Lab next examined the performance scalability of the ExDS9100 in a content delivery environment using the IOzone file system benchmark tool.⁵ The term 'content delivery' is used to define a broad category of systems which deliver digital files, video, audio, and images over a delivery medium such as a broadcasting network, Intranet, or the Internet. The delivery of large files over a content delivery medium requires a storage solution with lots of capacity and high aggregate read throughput.

IOzone was used to simulate the delivery of content using a number of threads performing 1 MB sequential reads. Multiple threads were directed at multiple file systems on each of four performance blocks spread over two capacity blocks in an ExDS9100.⁶ These results were compared to scale out tests run by HP using the same test parameters in a fully loaded system with sixteen performance blocks and ten capacity blocks. The results, illustrated in Figure 11, show the ExDS9100's maximum throughput scaling linearly as performance and capacity blocks are added.

⁵ <http://www.iozone.org>

⁶ Test configuration details can be found in the Appendix.

FIGURE 11. PERFORMANCE SCALING: LOCAL FILE SYSTEM 1 MB STREAMING READS



Why This Matters

ESG research has confirmed that the number one challenge IT managers report facing in their storage environments is keeping pace with the runaway capacity demands of growing volumes of data.⁷ As the size and number of files that need to be kept online continues to grow, capital equipment and operating budgets are being stretched to their limits. Reconfiguring an application to recognize capacity in the next new storage array can lead to downtime, lost productivity, and, in some cases, lost revenue as legacy storage arrays are filled to capacity. ESG Lab has confirmed that adding performance and storage capacity to an HP ExDS9100 can be done online with no disruption to users or applications.

⁷ Source: ESG Research, 2008 Enterprise Storage Survey, 2008

The maximum throughput recorded (3,430 MB/sec) was used to calculate the number of streams that could be delivered for a number of well known content types.⁸ For example, a bit stream rate of 128 Kbps for MP3 files was used to determine that the ExDS9100 has the bandwidth required to support up to 219,520 song downloads (3,430 MB/sec divided by 128 Kbps) as shown in Figure 12.

FIGURE 12. CONTENT DELIVERY: MAXIMUM OBJECTS SERVED FROM AN EXDS9100





TABLE 1. CONTENT DELIVERY RESULTS

| Content Type | Objects Served |
|------------------------------------|----------------|
| Internet Audio (MP3) | 219,520 |
| Internet Images (JPG) | 140,493 |
| Internet Video (MPEG) | 93,662 |
| CD Audio (ISO) | 19,600 |
| Broadcast TV (Standard Definition) | 7,317 |

ESG Lab also tested and audited network file system throughput to validate the ExDS9100's performance in a scale-out NAS environment.

⁸ The bit rates for each content type are listed in the Appendix.

FIGURE 13. NETWORK FILE SYSTEM THROUGHPUT: 32 K TRANSFER BLOCK SIZE



As can be seen in Figure 13, the ExDS9100 was able to sustain impressive throughput that scaled linearly as performance and capacity blocks were added. Results were very good for transfer block sizes from 32 k down to 4 k, which speaks well to the ExDS9100's operational efficiency. Many NAS systems generally perform significantly better with larger block sizes, which place lighter demands on a storage system's resources.

What the Numbers Mean

- A single ExDS9100 has the raw bandwidth required to service 93,662 concurrent Internet video (MPEG) downloads.
- The raw throughput results are based on a content delivery system that accesses content stored on disk using a large block size of 1 MB. With many clustered file system solutions, a smaller block size (e.g., 32 KB) might serve fewer objects. ESG Lab validated that the ExDS9100 delivers nearly identical throughput for a wide range of block sizes.
- The ExDS9100 showed excellent NFS throughput, driving 2.6 GB/Sec—more than 80% of the raw throughput available from the SAS back-end.
- Running small block IO (4 KB) over NFS was equally impressive, sustaining 1.9 GB/Sec or nearly half a million IOs/Sec

Why This Matters

ESG research asked enterprise IT managers to name the most important considerations in their decision to deploy scale-out NAS.⁹ Scalability and performance both ranked in the top five responses. Meeting the bandwidth-intensive performance demands of file-oriented applications using traditional storage architectures can lead to over-provisioning, increased complexity, and wasted capital costs.

Leveraging the cost-effective processing power, bandwidth, and memory of HP blade servers which are factory integrated into an HP ExDS9100, a single pool of HP Extreme Data Storage can scale to meet the needs of applications with extreme aggregate capacity and throughput requirements. An additional benefit of the standards-based architecture is the potential to host the application tier directly on the storage system, potentially eliminating server tiers from the application architecture. Taking advantage of spare CPU horsepower within the storage system can not only make file intensive operations run quicker, it can also be used to drastically reduce the amount of data moving over a network. Applications running directly in an open architecture storage system eliminate the need to move content between digital islands and reduce cost by shrinking server hardware requirements.

ESG Lab was extremely impressed by the ExDS9100's consistent, near linear-scalable performance for both local and networked file systems. A fully configured ExDS9100 was able to drive 3.34 GB/second of streaming IO.

⁹ Source: ESG Research, *Enterprise Storage Survey*, 2008

Availability

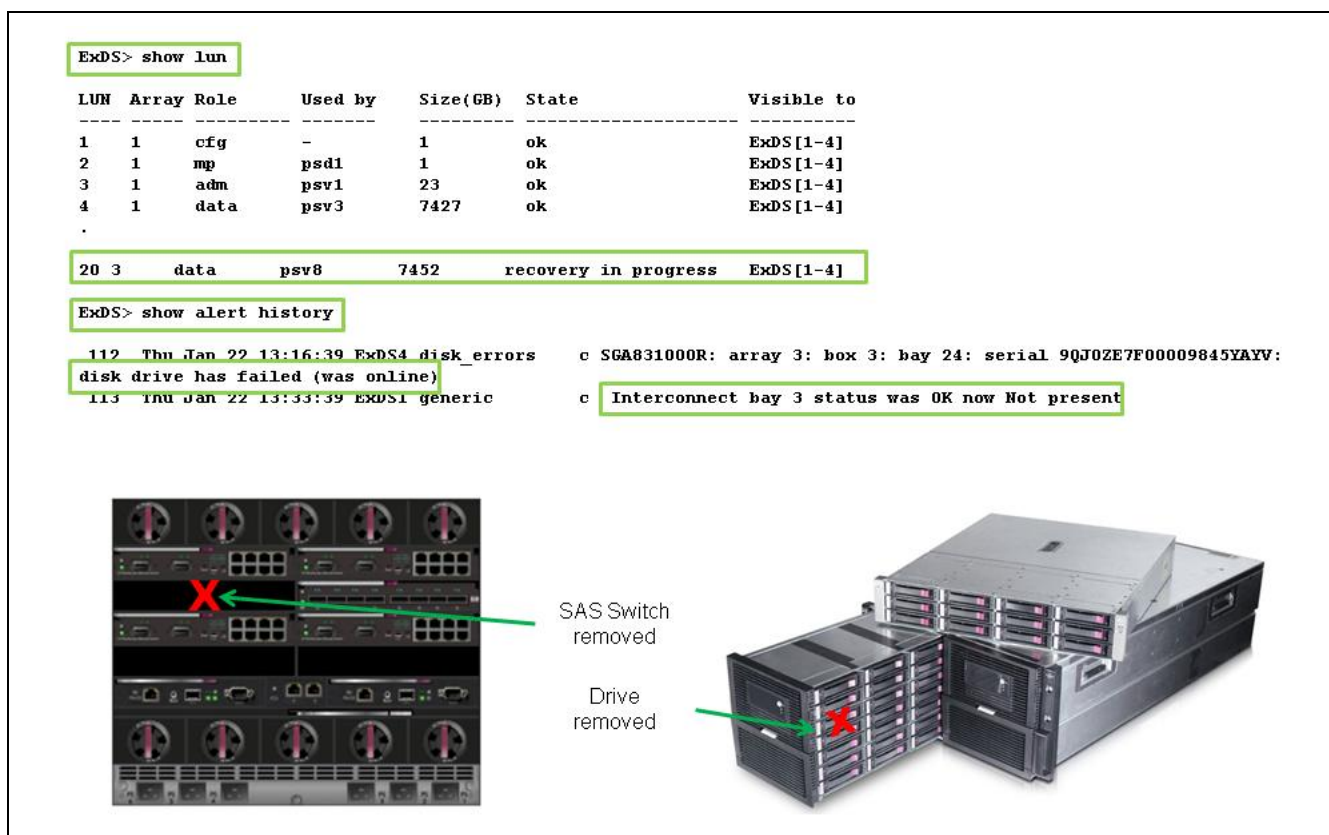
High availability (HA) and fault tolerance are major concerns for today's IT administrators. They have come to expect the same enterprise-class reliability in a scale-out Tier 2 solution that they have with their primary storage systems. The ExDS9100 is architected as an integrated appliance, meaning that disks and RAID protection are managed automatically. This allows tight control and better response to hardware failures with less administrator interaction.

ESG Lab Testing

ESG Lab simulated disk, SAS switch, and server failures while the system was running and actively servicing clients. First, ESG Lab generated writes with the Linux client by running the linux dd command to copy 500 MB of data to a file on the target file system in a continuous loop. A disk drive was pulled from the capacity block housing the file system while data was being written. An administrator's e-mail account being monitored by ESG Lab received a message stating that a drive had failed and indicating the specific array, box, bay, and slot the drive was in. In a live customer system, an e-mail would also be sent to HP customer service.

As seen in Figure 14, the ExDS manager 'show lun' command displayed a notification that LUN 20 in Array 3 was undergoing a RAID rebuild. While the rebuild was taking place, ESG lab pulled one of the SAS switches from the HP c7000 chassis. These SAS switches provide connectivity between the performance blocks and the capacity blocks.

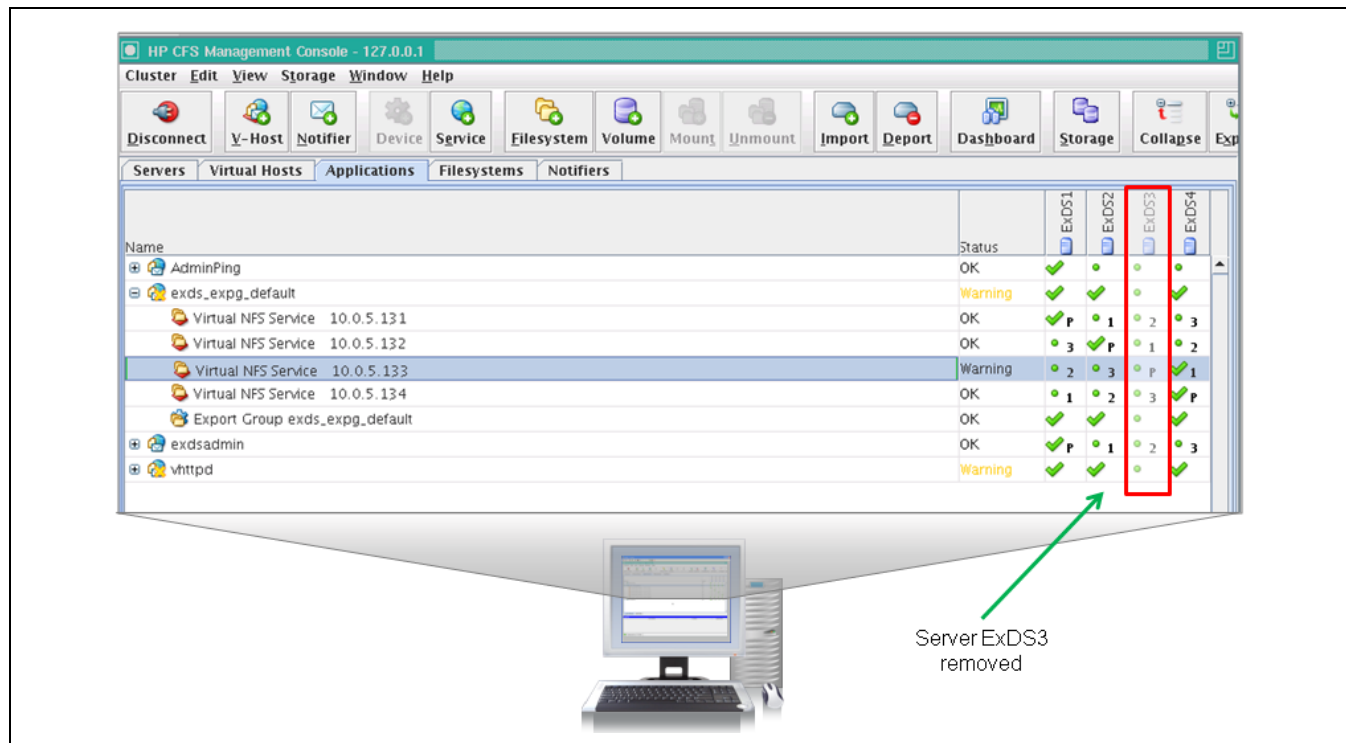
FIGURE 14. SIMULATING COMPONENT FAILURES



Again, the administrator's mailbox almost immediately received a message indicating that a failure had occurred. In this case, interconnect bay 3 had disappeared. ESG Lab verified both errors using the 'show alert history' command, which gave detailed data on each failure.

Finally, ESG Lab tested file system availability by shutting down a performance block while a client was mounted to a file system and writing data to it. The net effect of this procedure was exactly the same as a server failure or crash. It is also similar to what takes place when rolling upgrades, which require reboots of the performance blocks, are performed. ESG Lab again generated writes with the Linux client by running the linux dd command to copy 500 MB of data to a file on the target file system in a continuous loop. Next, ESG Lab issued the shutdown server command and observed the transition of performance block ExDS3's virtual IP (and all file systems exported on it) to server four, as seen in Figure 15.

FIGURE 15. TRANSITIONING FILE SYSTEMS BETWEEN BLADES



As the virtual IP transitioned from ExDS3 to ExDS4, write activity paused but the write operation did not fail. The file system remained mounted on the client and the dd operation continued without error after the transition completed in a few seconds. The ExDS9100's clustered file system's global file handles ensure that file system access can transition between performance blocks with no disruption and no re-mount required.

Why This Matters

As storage systems grow in size and complexity, so too do the chances of things going wrong. More than half of IT managers surveyed by ESG indicated data availability as a major driver in choosing to deploy scale-out NAS.¹⁰ Regardless of the number and types of hardware failures that may occur during the life of digital files on disk, managers, employees, and customers expect their data to be available. ESG Lab has confirmed that the ExDS9100 can continue to operate through hardware and software faults thanks to a robust, integrated architecture combined with bulletproof clustered file services. HP's clustered file system provides continuous access to file data, minimizing both planned and unplanned outages.

¹⁰ Source: ESG Research, 2008 Enterprise Storage Survey, 2008

ESG Lab Validation Highlights

- ☑ Within three minutes of sitting down in front of the console of a factory default ExDS9100, ESG Lab had completed configuration and the ExDS was providing file services over NFS and HTTP.
- ☑ With just one command, ESG Lab installed a new performance block. The new server was automatically installed, configured, and joined to the cluster with no interaction required.
- ☑ The ExDS was able to sustain 3.2 GB/sec of raw performance and 2.6 GB/sec of network file performance—enough to serve hundreds of thousands of individual streams of internet content.
- ☑ The ExDS continued to serve file systems through hard drive, back-end switch, and server failures seamlessly and without interruption.
- ☑ The ExDS clustered file system allows file systems to move between servers hot and online without disrupting client access.

Issues to Consider

- ☑ The ExDS provides access to CIFS clients through the use of Samba. CIFS access does not come pre-configured like NFS and HTTP, but must be added to the ExDS solution by HP. The Samba solution for CIFS provides high availability for CIFS connections where a single blade manages specific file systems. The exponential growth of unstructured data will soon make systems like the ExDS attractive for large scale CIFS consolidation and a pre-configured CIFS option will make sense.
- ☑ While the ExDS CLI is exceptionally easy to use, users have access to the disparate GUIs that are part and parcel of the various components. An ExDS-specific consolidated GUI that brings together all the functionality of the CLI while enabling management of multiple systems from one pane of glass would be a welcome addition.

ESG Lab's View

Market drivers for early adopters of scale-out NAS include faster provisioning, improved scalability and performance, easier management, and the need to support specific, fast-growing applications. The management efficiencies of scale-out NAS are so compelling that, in a recent ESG survey of 504 North American and European IT professionals, lower cost of infrastructure was literally *last* on the list of buying criteria for early scale-out adopters.¹¹ Planned and potential users have vaulted lower cost into the top tier of purchasing criteria, but cost remains second to the improved scalability that scale-out NAS brings to the table. Even during the current economic crisis, the manageability of scale-out NAS continues to be a major reason users are considering deploying it and is on par with cost as a driver for future adoption.

The ExDS9100 is a highly scalable, affordable, and easy to manage file storage system designed to enable enterprises to store hundreds of terabytes of file-based data in a highly available repository providing storage services to users and customers. The ExDS9100 delivers excellent performance and availability and its ability to independently scale both capacity and performance provides users and applications with 3.2 GB/sec of raw performance and up to 820 TB of storage in a single system.

ESG Lab found the ExDS9100 to be easy to set up and manage, providing impressive performance in a package that took just three commands to configure. ESG Lab was able to scale capacity and performance independently with simple commands. A fully configured ExDS9100 was confirmed to sustain up to 3.2 GB/sec of raw disk performance, while providing the ability to host applications on the server blades in the enclosure.

The ExDS9100 also demonstrated enterprise class availability automatically and transparently. File systems transitioned between servers with no disruption to access, eliminating both planned and unplanned outages. Disk and component failures were handled transparently as well, alerting the administrator that an issue had occurred while automatically taking non-disruptive, corrective action.

In the current environment of economic uncertainty, the advanced features that the HP ExDS9100 brings to the table, such as scalability and ease of management, are more important than ever. With the HP ExDS9100, users can scale capacity without scaling headcount. Provisioning is quick and easy. The system has an "always on" architecture that can seamlessly withstand multiple failure modes. HP has done a good job leveraging IP from multiple business units to create a seamless, appliance-like experience. Unlike some systems that require multiple interfaces to manage the storage, blades, and file system, HP has an all-in-one management console. HP seems to have met its design goals with this system; the HP ExDS9100 is a massively scalable, highly affordable file storage appliance that is shockingly easy to deploy, manage, and upgrade.

¹¹ Source: ESG Research, 2008 Enterprise Storage Systems Survey, November 2008

Appendix

TABLE 2. TEST CONFIGURATION

| | |
|---|---|
| HP ExDS9100 Extreme Data Storage System 4 Performance blocks 3 Capacity blocks: 246 1 TB SAS HDs | |
| Local File System Clients | |
| 16x HP BL460 Blade servers <ul style="list-style-type: none"> Dual Quad Core Xeon CPUs 8GB RAM | OS: HP ExDS Linux Variant |
| NFS Clients | |
| 24 HP BL460C Blade servers <ul style="list-style-type: none"> Dual Quad Core Xeon CPUs 8GB RAM 24 HP HP DL 320 servers <ul style="list-style-type: none"> Single Dual Core Xeon CPU 4GB RAM | OS: RedHat Enterprise Linux 5 Update 2 (x86_64) |
| 1 HP Workstation: Management station and HTTP client | OS: Windows XP SP3 |
| IOzone benchmarking software | Version: 3.303 32 bit |
| 3x HP ProCurve 2900-48G GBE Switches with dual 10GbE uplinks | |

TABLE 3. CONTENT DELIVERY BIT RATES

| Content Type | Bit Rate |
|--|-----------|
| Internet Audio (MP3) | 128 Kbps |
| Internet Image (640x480 JPG, 4 second download/view) | 200 Kbps |
| Internet Video (MPEG) | 300 Kbps |
| CD Audio | 1.4 Mbps |
| Broadcast TV (SD) | 3.75 Mbps |



20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com