

# Lab Validation Report

## **Dell EqualLogic and IBRIX** Simply Scalable Enterprise-class File Serving

**By Brian Garrett**  
With Terri McClure

**March, 2009**

# Table of Contents

<b>Table of Contents .....</b>	<b>i</b>
<b>Introduction .....</b>	<b>1</b>
<i>Background .....</i>	<i>1</i>
<i>Dell EqualLogic and IBRIX .....</i>	<i>2</i>
<b>ESG Lab Validation.....</b>	<b>3</b>
<i>Getting Started .....</i>	<i>3</i>
<i>Scalability .....</i>	<i>5</i>
<i>Proven Performance and Scalability .....</i>	<i>5</i>
<i>Introducing the IBRIX Fusion Scale-Out NAS Platform.....</i>	<i>6</i>
<i>Enterprise-class NAS .....</i>	<i>8</i>
<i>Snapshots.....</i>	<i>8</i>
<i>Continuous File Replication.....</i>	<i>10</i>
<i>Tiered Migration.....</i>	<i>11</i>
<i>High Availability .....</i>	<i>13</i>
<b>ESG Lab Validation Highlights.....</b>	<b>14</b>
<b>Issues to Consider.....</b>	<b>14</b>
<b>ESG Lab's View.....</b>	<b>15</b>
<b>Appendix .....</b>	<b>16</b>

## 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 IBRIX.

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 ever-increasing volume of unstructured file data that needs to be shared, managed, and protected has become a real problem for IT managers—particularly those who need to provide online access to shared files for revenue generating workflows and processes. This report explores how Dell EqualLogic PS5000 Series iSCSI arrays and PowerEdge servers, along with IBRIX Fusion—a purpose-built platform for scale-out NAS (network-attached storage)—have been combined to create a fast and scalable networked file serving solution that is easy to deploy and manage.

## 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. IT managers must store, deliver, and manage large volumes of unstructured data while meeting ever-increasing capacity requirements. ESG research indicates that the vast majority of corporate digital assets are stored in unstructured files. Unstructured content—which includes 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.<sup>1</sup>

**FIGURE 1. DIGITAL ARCHIVE FORECAST**



Scale-out and consolidation initiatives are driving a wave of large-capacity networked file system deployments in data centers around the world. Data center managers are struggling to feed the ever-expanding capacity and performance requirements of applications and workflows that require shared access to a large network-attached file system. Data intensive applications—including rich media, backup and archive, scientific computing, content distribution, and clustered computing—are driving the need for cost-effective scale-out. Mixed applications sharing a consolidated pool of storage are driving the need for scale-out NAS solutions that are fast for a wide

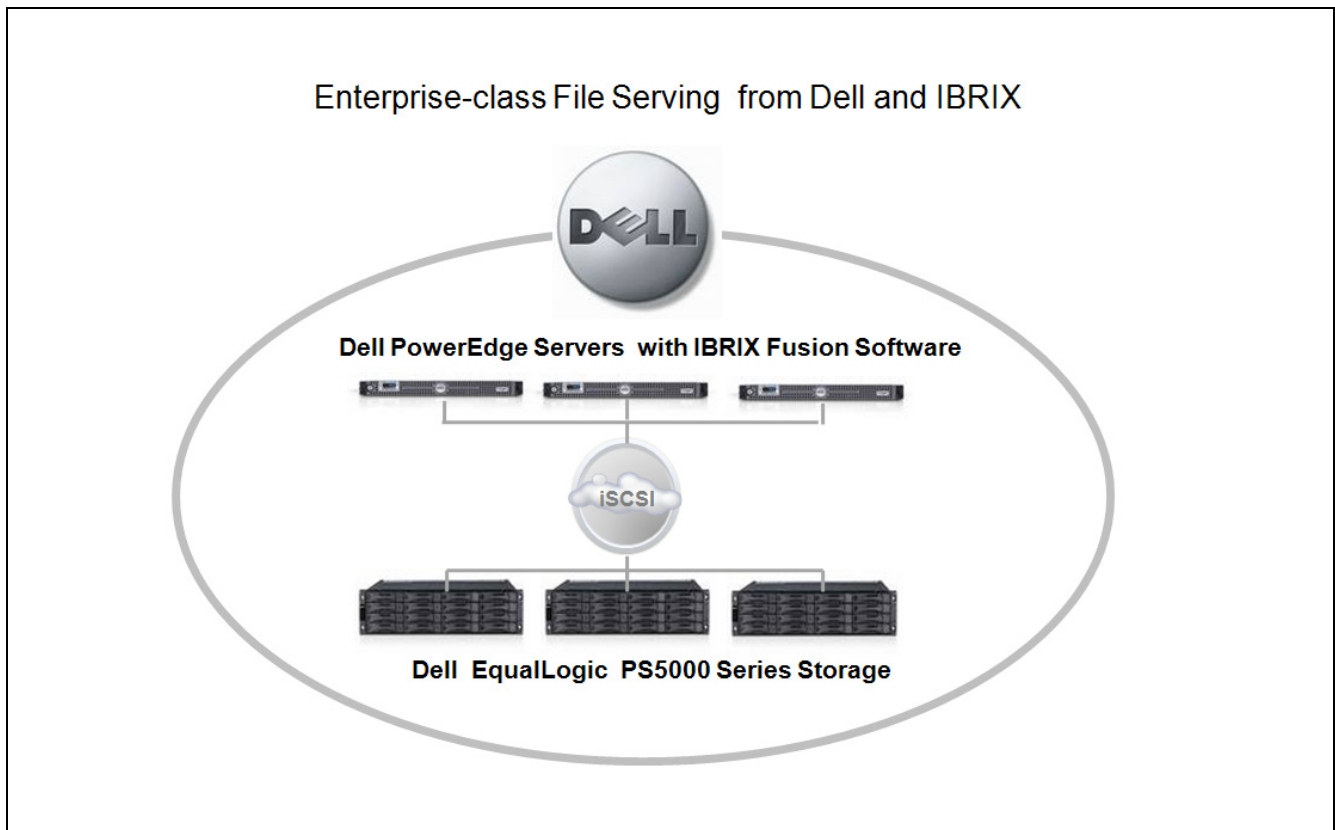
<sup>1</sup> Source: ESG Research Report, *Digital Archiving Survey*, November 2007

variety of application workloads—from the sequential, bandwidth intensive needs of large files to the random, response-time sensitive needs of metadata management and general purpose applications.

## **Dell EqualLogic and IBRIX**

IBRIX recently introduced a fully integrated, enterprise-class bundled solution designed to meet the needs of organizations requiring tens of terabytes—or more—of shared network-attached storage capacity. As shown in Figure 2, IBRIX Fusion software running on Dell PowerEdge servers harnesses the CPU, memory, and bandwidth of Dell servers to present iSCSI-attached Dell EqualLogic PS5000 Series storage capacity as a single pool of network-attached storage. The combination of Dell hardware and IBRIX Fusion software leverages industry standard server, Ethernet, and iSCSI technology to cost effectively meet the growing capacity and performance needs of organizations with scale-out NAS requirements.

**FIGURE 2. ENTERPRISE-CLASS FILE SERVING**



Scale-out NAS solutions from IBRIX are preconfigured for quick and easy installation and deployment. Enterprise-class NAS capabilities—including snapshots, continuous remote file replication, and automated migration of files between tiers of storage—are supported. Field-proven scale-out NAS solutions from IBRIX are predictably fast for a wide variety of applications, are easy to upgrade, and are fault tolerant.

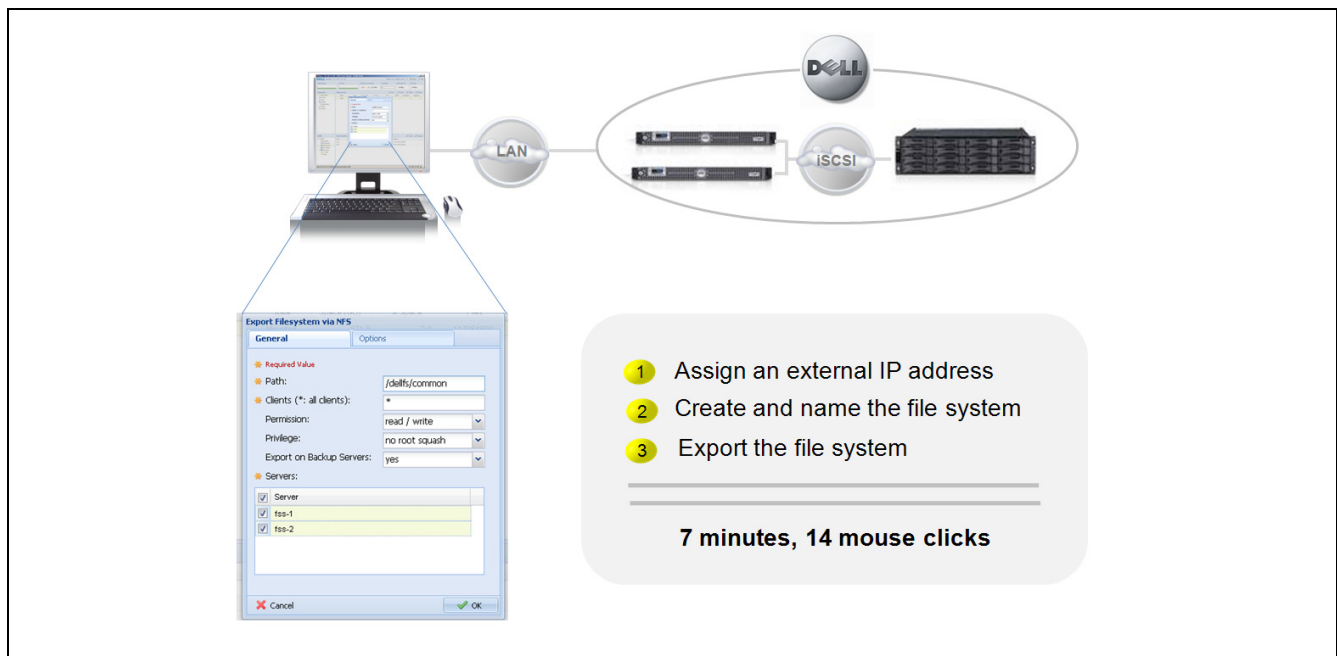
# ESG Lab Validation

ESG Lab performed hands-on testing at IBRIX corporate headquarters in Massachusetts. Areas of focus during this ESG Lab Validation include ease of deployment, scalability, enterprise-class NAS capabilities, and fault tolerance.

## Getting Started

ESG Lab testing began with a hands-on evaluation of the initial configuration process from an IT administrator's perspective. As shown in Figure 3, a LAN-attached workstation was used to configure IBRIX Fusion software pre-installed on a pair of Dell PowerEdge Servers connected to a Dell EqualLogic PS500 Series storage system.<sup>2</sup> A Web browser was used to access the centralized management console. A three step process was used to assign an external IP address for the system, create and name a file system, and export that file system for access as a shared network mount point. Seven minutes after starting, files were being copied over an Ethernet LAN to an NFS-mounted network-attached file system.

**FIGURE 3. QUICK AND EASY DEPLOYMENT**



ESG Lab noted that the time that it takes to configure IBRIX Fusion is the same, regardless of the amount of capacity deployed. This is due to the fact that the IBRIX Fusion file system is pre-formatted and ready to use when it arrives at the customers site.

Through a web browser aimed at the IP address of the management console, ESG Lab browsed the configuration of the system. This single management console was used to manage the entire solution including the Dell PowerEdge servers, Dell EqualLogic iSCSI disk arrays, and the IBRIX Fusion file system.

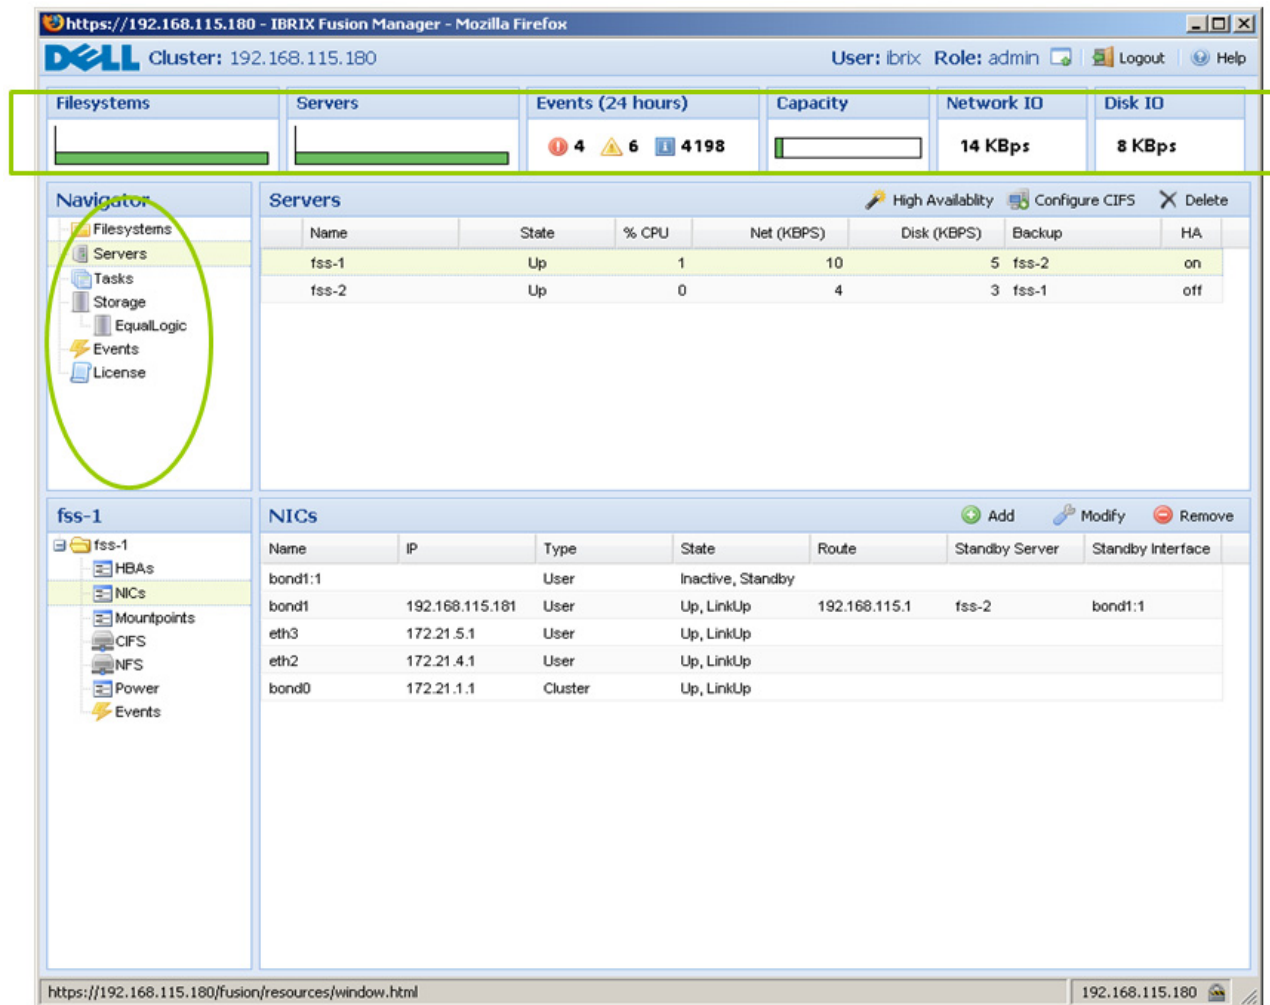
ESG Lab has tested IBRIX in the past and was impressed with the new look and feel of the management GUI. As shown in Figure 4, the top region of every management interface display panel can be used to determine the

<sup>2</sup> Configuration details are listed in the Appendix.

overall health and status of the system. The health and overall capacity utilization of the system are displayed graphically alongside a network and disk performance summary.

Navigating the management GUI was straightforward and intuitive. A familiar, expandable topology view on the left was used for navigation and control. Drilling down into any of the components of the system expanded the view and brought new levels of detail into view in the bottom panel. In the example shown in Figure 4, the configuration of the network interface cards (NICs) within one of the server (*fss1*) was displayed.

**FIGURE 4. INTUITIVE CENTRALIZED MANAGEMENT INTERFACE**



## Why This Matters

Whether it's a clustered file system in front of SAN-attached storage or a number of traditional NAS systems glued together with virtualization software, time and money can be wasted trying to deploy and grow legacy file systems for large-scale applications.

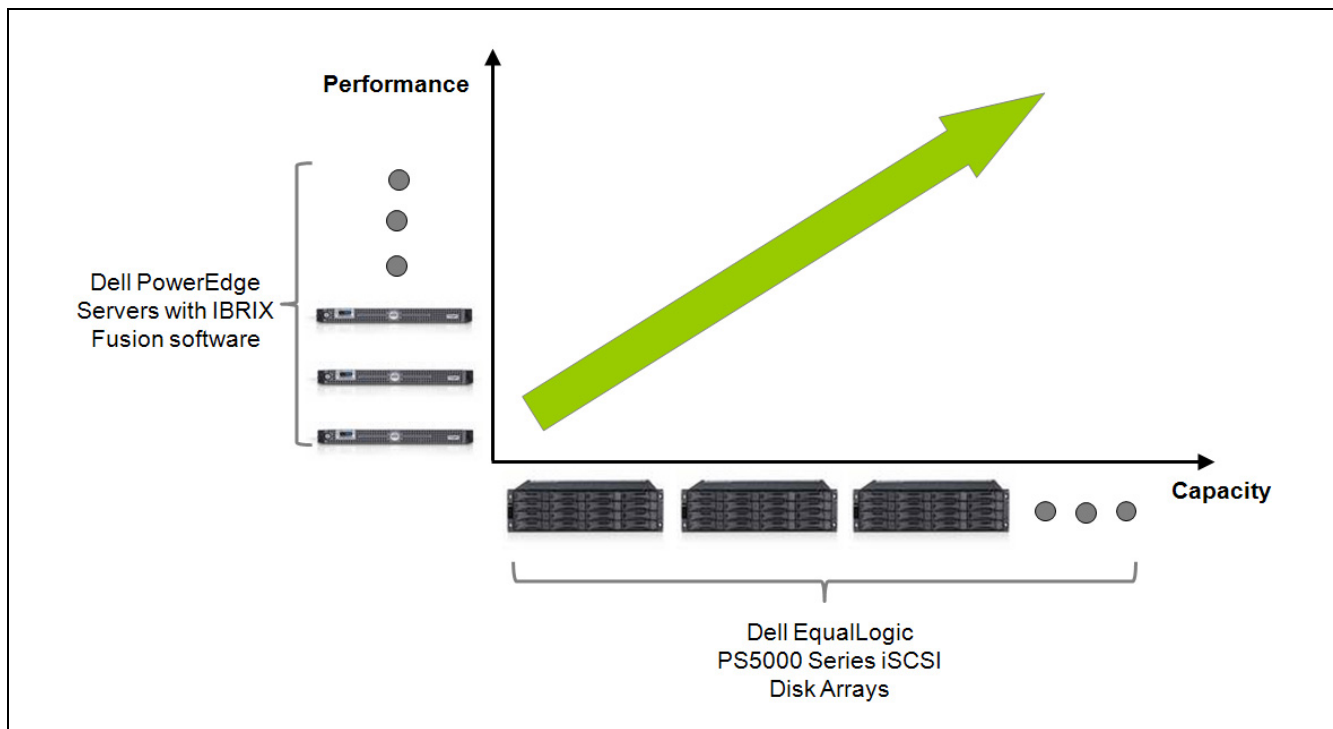
A Scale-Out NAS solution from IBRIX is extremely easy to install, deploy, and manage. Seven minutes and fourteen mouse clicks after beginning a configuration from scratch, ESG Lab was accessing files over an Ethernet network from Linux and Windows clients.

## Scalability

One of the key benefits of IBRIX is its ability to provide a high-level of aggregate performance and throughput that scales in near-linear fashion as servers and storage are added. With its practically limitless file system size (16 million gigabytes), IBRIX can scale to meet the needs of applications that require a boost in performance, a boost in capacity or a combination of both. As additional Dell PowerEdge servers running IBRIX Fusion software are added, performance increases due to the additional processing power, memory, and bandwidth of each server. As additional iSCSI-attached Dell EqualLogic PS5000 Series disk arrays are added, the capacity—and performance—of the solution increases as well.

Different models of Dell EqualLogic PS5000 Series storage systems can be used to meet a variety of application performance needs. Fast SAS drives are supported for performance-sensitive applications and workflows, including clustered scientific high performance computing and rich media editing and delivery. Affordably dense SATA drives are supported for capacity-intensive applications including deep archives, backup to disk, and large scale consolidation of general purpose legacy file systems.

**FIGURE 5. FLEXIBLE SCALABILITY**



Unlike a number of competitive scale-out NAS solutions, which are tuned exclusively for bandwidth intensive workloads, a scale-out NAS solution from IBRIX performs well for a wide variety of workloads including response time-sensitive small random workloads (e.g., general purpose file sharing) and bandwidth-intensive large sequential workloads (e.g., high definition video editing).

## Proven Performance and Scalability

ESG Lab has been tracking the continuously improving performance scalability of the IBRIX and EqualLogic product lines for a number of years. Both solutions use a clustered architecture to cost effectively support extreme levels of performance to meet a wide variety of application needs. Large-scale production deployments and hands-on testing by ESG Lab have confirmed that each can scale to deliver gigabytes per second of aggregate bandwidth and tens of thousands of small I/O operations per second. As evidence, a few of the highlights from previous ESG Lab Validation Reports are summarized:

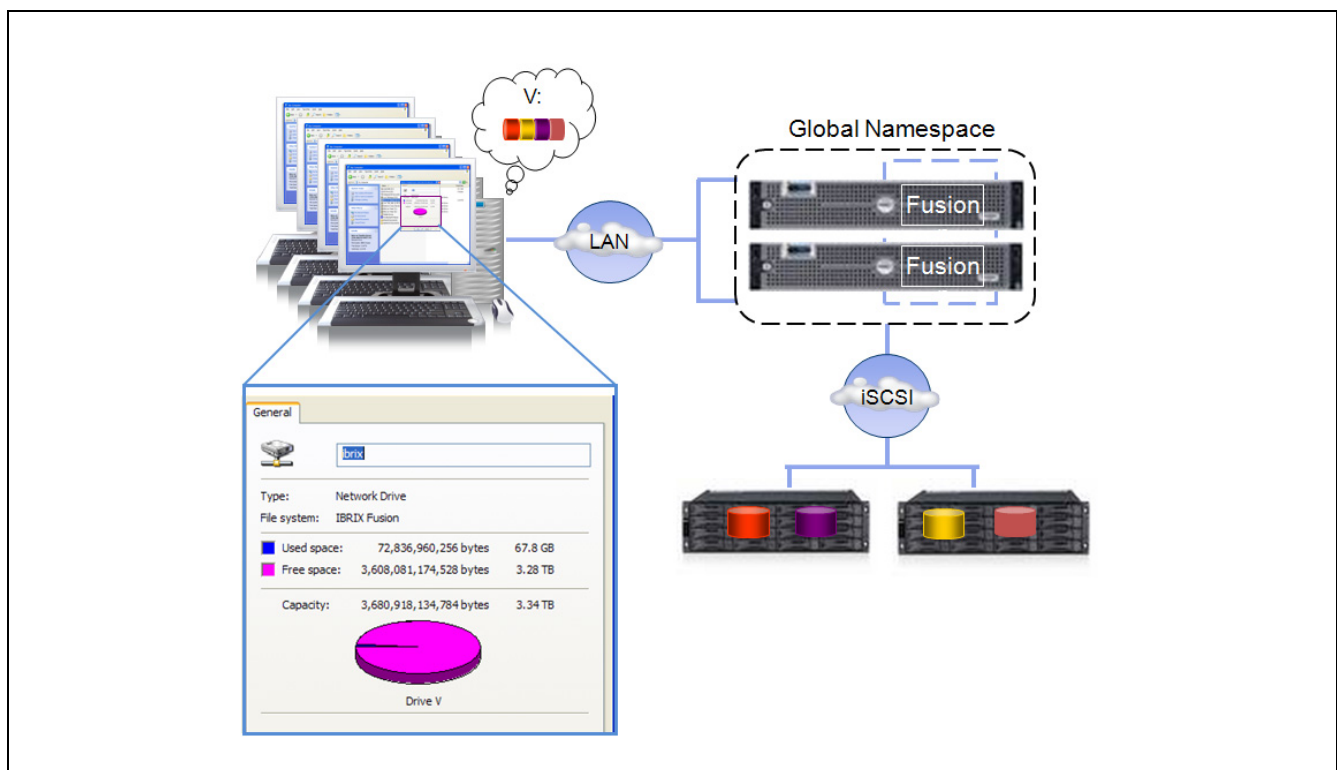


- A single server connected to a modest three array EqualLogic system supported 9,000 typical corporate Microsoft Exchange E-mail users (loadsimm MMB2).<sup>3</sup>
- A single EqualLogic array supporting 50,000 I/Os per second (IOPS) from cache scaled up to 1.2 million IOPS on a 25 array system.
- Three servers running IBRIX Fusion software sustained 1.5 GB/sec in aggregate read performance, achieving linear scalability as nodes were added to the IBRIX cluster.<sup>4</sup>
- A three dual core server IBRIX Fusion solution achieved 75,000 NFS file operations per second. Performance scaled in perfect linear fashion as IBRIX Fusion servers were added.

### Introducing the IBRIX Fusion Scale-Out NAS Platform

IBRIX Fusion is built using storage elements called segments. Segments are made out of hard drive capacity presented as iSCSI-attached logical drive units (LUNs) by a Dell EqualLogic PS5000 Series disk array. In the example shown in Figure 6, a clustered IBRIX Fusion file system has been created using four internal hard drives within two Dell EqualLogic arrays. The clustered file system is presented as a single global namespace. Files and directories are striped over all of the segments in the file system (in this example, a network-attached V: drive accessed by a Windows XP client).

**FIGURE 6. THE IBRIX FUSION ARCHITECTURE**



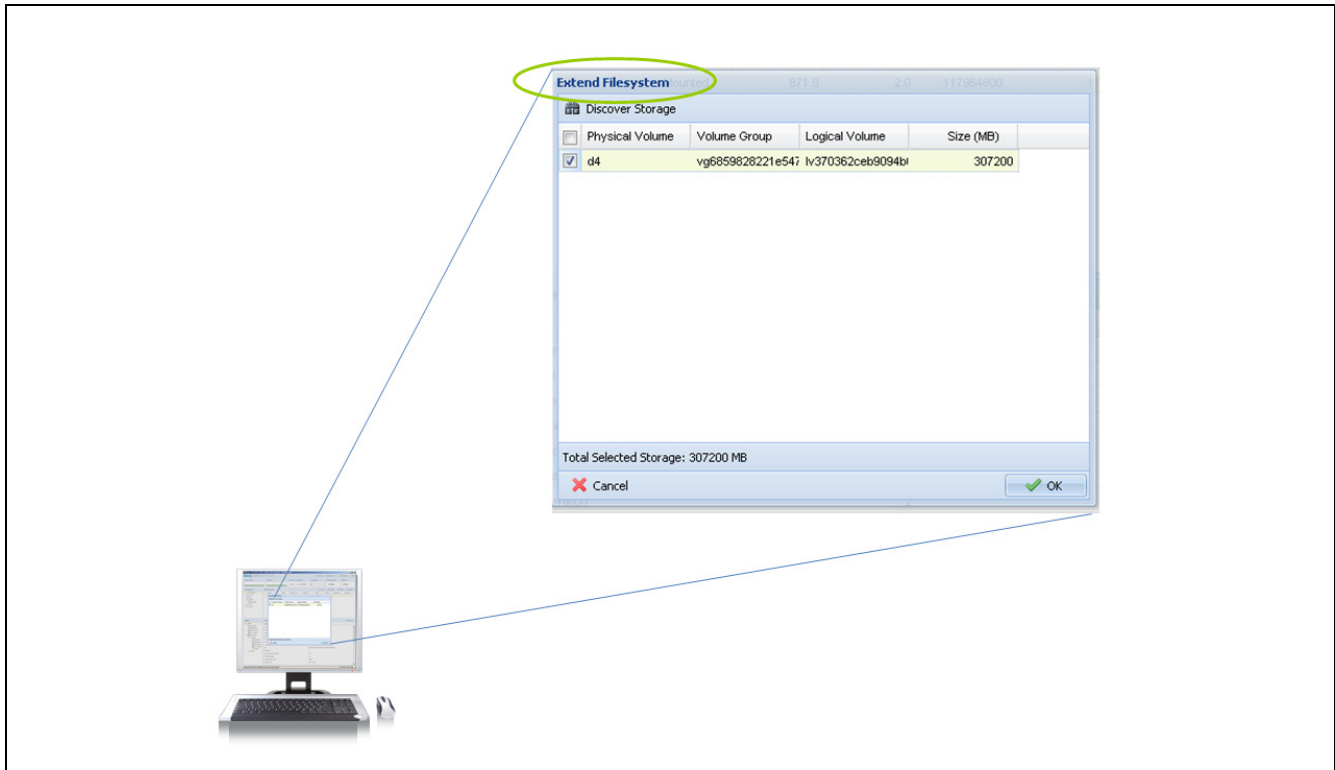
<sup>3</sup> Source: ESG Lab Validation Report, *EqualLogic PEER Storage*, September 2004

<sup>4</sup> Source: ESG Lab Validation Report, *IBRIX Fusion*, April 2008



Server horsepower and storage capacity can be added on-the-fly to an existing IBRIX environment. ESG Lab confirmed this capability by expanding a file system while files were actively being copied to an existing file system. As shown in Figure 7, expanding an existing file system is a simple web management console operation. An intuitive right mouse click on an existing file system was used to access the panel shown. Less than three minutes after clicking OK, the additional capacity was online and available. The file copy operation continued without error or interruption throughout the upgrade.

**FIGURE 7. EXPANDING THE FILE SYSTEM**



## Why This Matters

As organizations struggle to meet demand for increased capacity and performance, the reconfiguration of legacy storage and file systems can lead to downtime, lost productivity, and increased IT expenditures. IBRIX aggregates industry standard server and storage technology into a single file system that can expand online to an extremely large single pool of capacity (up to 16 petabytes). ESG Lab grew an existing scale-out NAS solution from IBRIX in less than three minutes while applications remained online and available.

A scale-out NAS solution that performs well for responsive time-sensitive workloads can be used to reduce the cost and complexity required to serve a growing number of applications and users. A scale-out NAS solution that performs well when processing large files, or a large number of files, can be used to speed workflows—saving time and money. The performance of IBRIX scales in a linear fashion to meet the transaction, response time, and throughput requirements of a wide variety of applications.

## Enterprise-class NAS

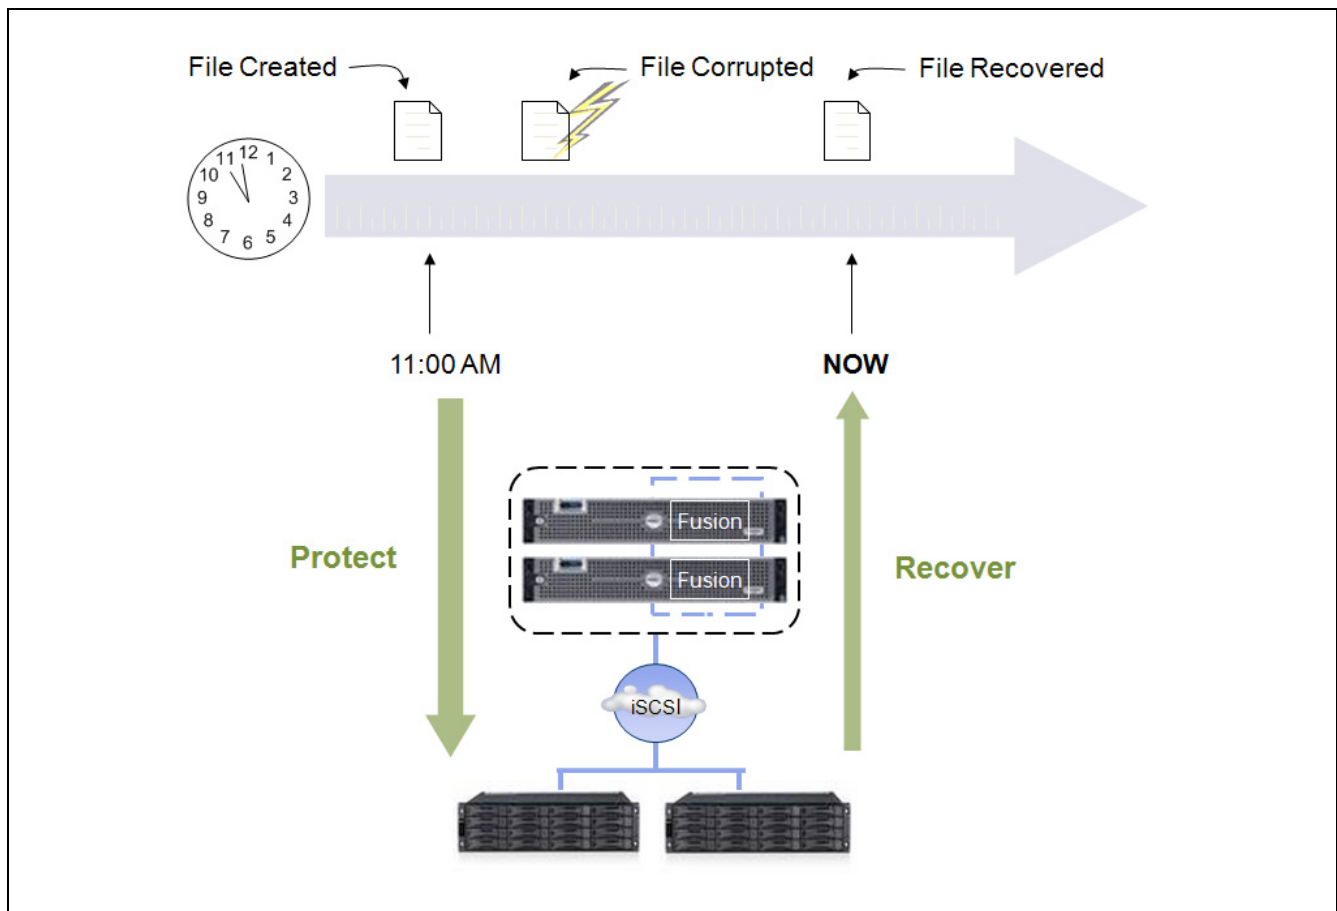
IBRIX supports a number of features that are typically associated with enterprise-class NAS solutions. These valuable features, including snapshots, continuous remote replication, and automated online migration between tiers of storage are often absent in competitive scale-out NAS solutions.

### Snapshots

Snapshots are used to create space-efficient point in time images of a file system on disk. Disk-based snapshots provide quick and easy recovery from corrupted or accidentally deleted files. Snapshots can also be used to quickly create a read-only version of a file system for backup jobs. Snapshot technology has been available for more than a decade in general purpose scale-up NAS solutions. Scale-out NAS solutions often lack this valuable capability. The scale-out NAS solution from IBRIX uses the field proven snapshot capabilities of EqualLogic PS5000 Series iSCSI disk arrays to support snapshots within the IBRIX Fusion file system.

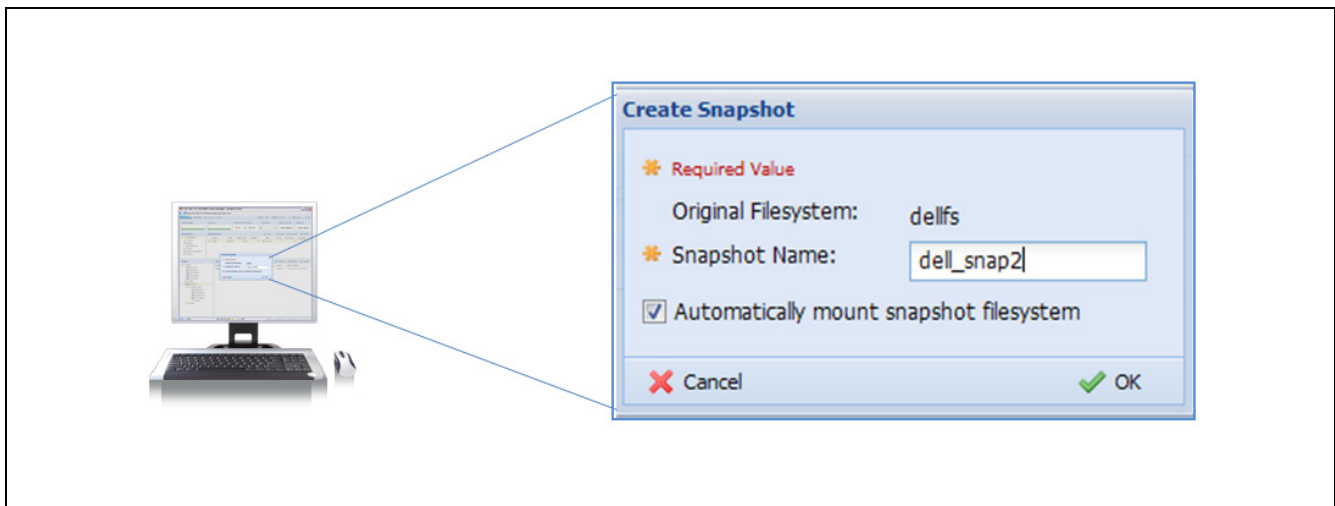
ESG Lab used a snapshot to simulate the recovery of a corrupted file during validation testing as shown in Figure 8. The management GUI was used to create a snapshot of a file created at 11:00 AM. The file was corrupted and recovered using the snapshot. Note that snapshot commands are managed by the IBRIX Fusion file system and executed by iSCSI-attached Dell EqualLogic PS5000 disk arrays.

**FIGURE 8. USING A SNAPSHOT TO RECOVER A CORRUPTED FILE**



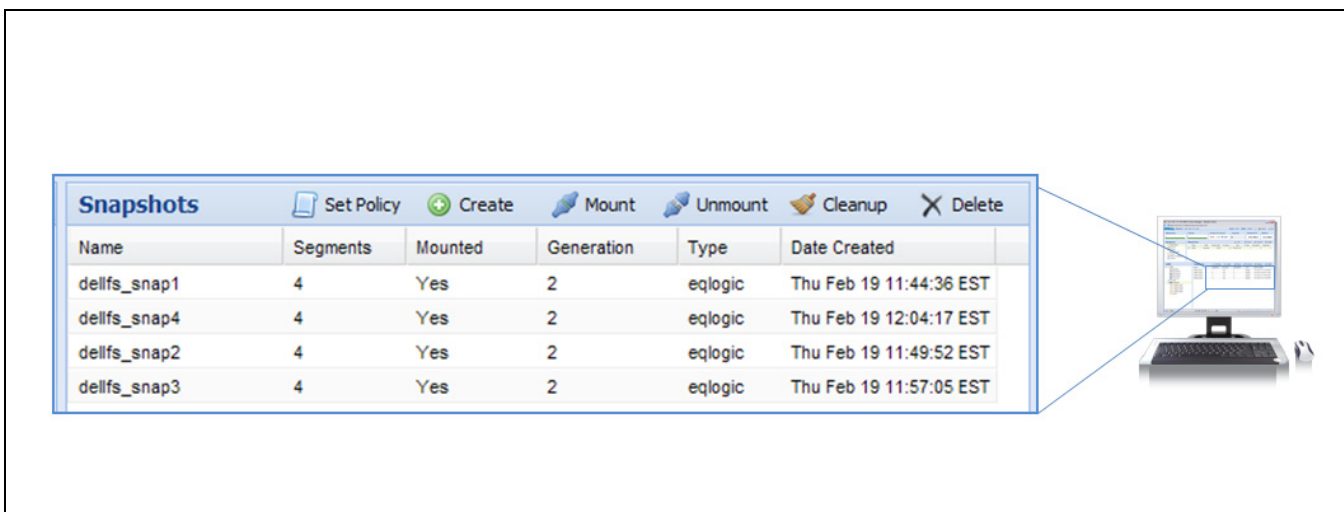
The GUI used to create a snapshot during ESG Lab testing is shown in Figure 9.

**FIGURE 9. CREATING A SNAPSHOT**



A series of four snapshots was created. Snapshots can be created, mounted, un-mounted, or deleted using the intuitive GUI shown in Figure 10. ESG Lab noted that each of these routine snapshot management tasks can also be performed using a well documented command line interface.

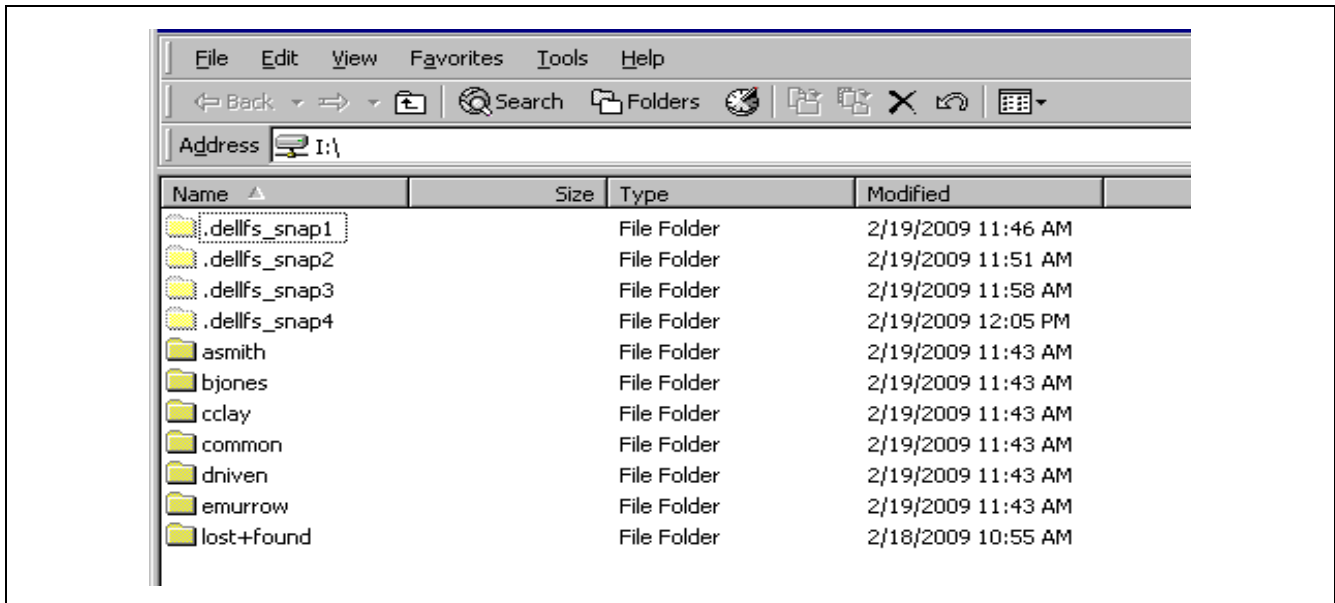
**FIGURE 10. MANAGING SNAPSHTS**



Next, ESG Lab edited a file in the *dellfs* file system to simulate file corruption. The file was recovered from Linux and Windows clients using the familiar interface shown in Figure 11.

Note that the contents of the file system at the time of each snapshot can be browsed by navigating into a snapshot directory that begins with a "." (e.g., *dellfs, snap1*). The file system contents, frozen at the point in time that the snapshot was taken, are available as a mounted file system that looks and feels exactly like the production file system. Recovering the corrupted file was performed in a matter of seconds by copying the file from the snapshot directory to the production file system. ESG Lab noted that the IBRIX command line utility can be used to perform snapshots on a scheduled basis.

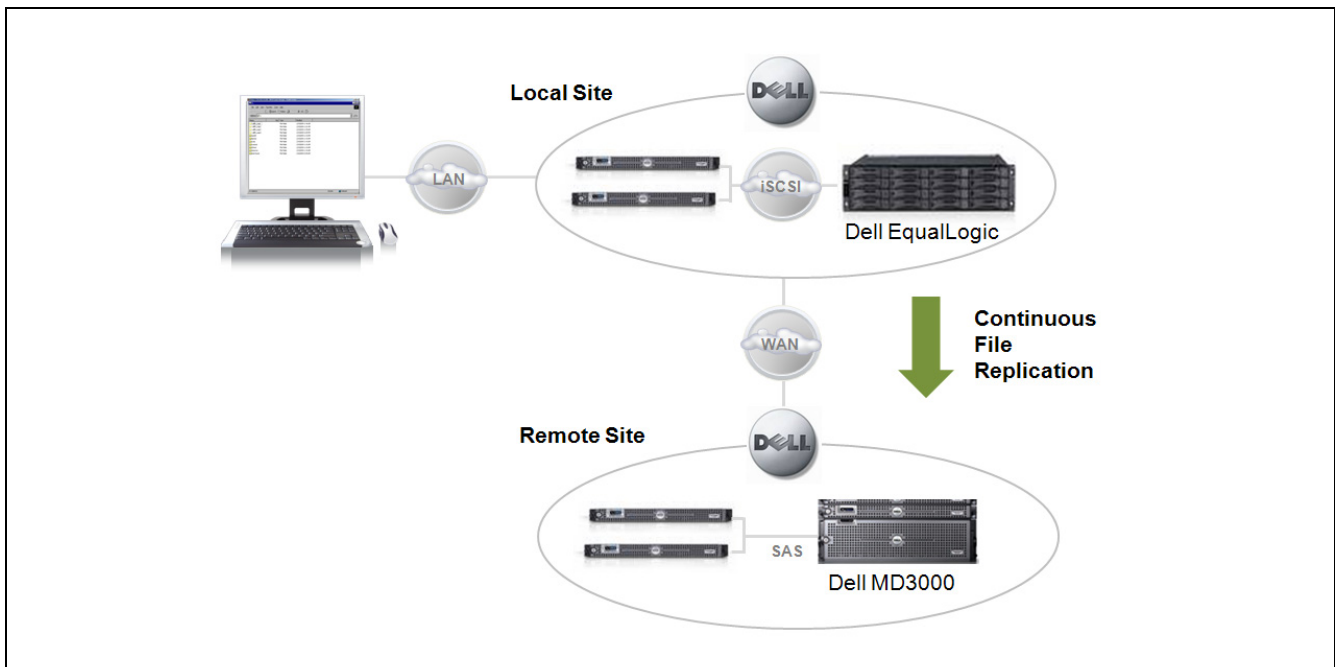
**FIGURE 11.** A WINDOWS EXPLORER VIEW OF SNAPSHOT RECOVERY IMAGES



### Continuous File Replication

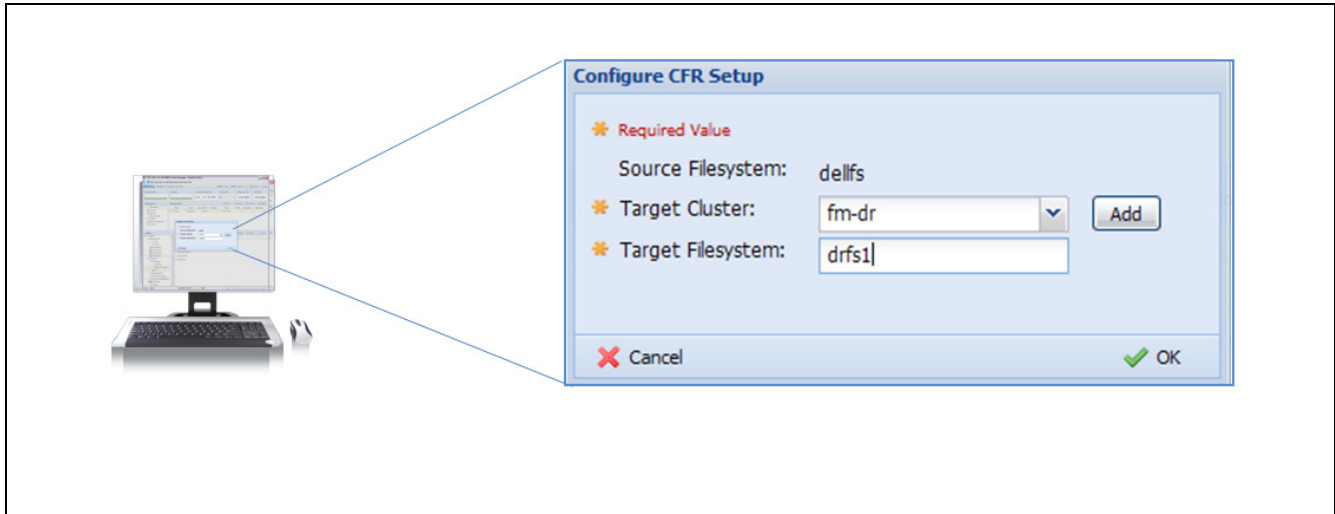
IBRIX can be configured to continuously replicate changed files to a remote site for disaster recovery. This valuable enterprise-class NAS capability was tested by ESG Lab using the configuration shown in Figure 12.

**FIGURE 12.** THE CONTINUOUS REMOTE REPLICATION TEST BED



Files residing in a local data center and stored on a Dell EqualLogic array were replicated to a simulated remote site which used a SAS-attached Dell MD3000 disk array for affordable disaster avoidance. Replication between the IBRIX cluster in the local site and the IBRIX cluster in a simulated remote site was configured using the GUI shown in Figure 13.

**FIGURE 13. CONFIGURING CONTINUOUS FILE REPLICATION**




With the relationship between the source and target system configured, the management console was used to start replication and a replication status of *Running* was observed. A 300 MB directory tree was copied from a local hard drive to the IBRIX file system at the local site. Ten seconds after the file copy had completed on the source side, the remote site was used to access a complete copy of the 300 MB directory tree at the simulated remote site. ESG Lab noted that the continuous replication process ran in parallel with the file copy. In other words, as each new file was saved, it was automatically replicated to the remote site.

### Tiered Migration

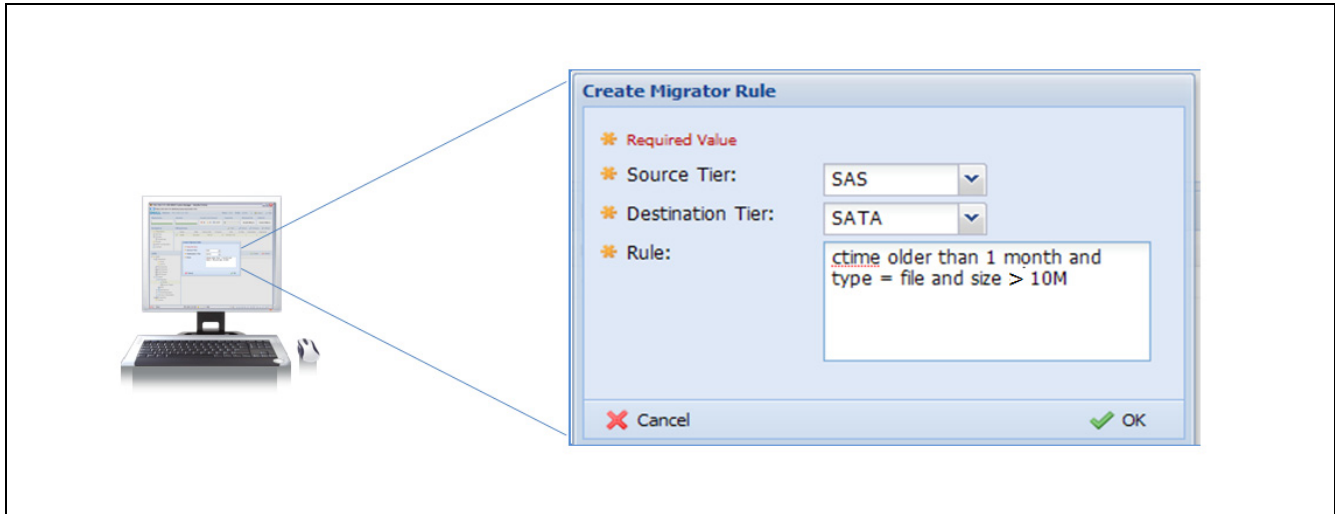
Tiered migration is used to transparently move files between different types of disk drives. This is most often used to move infrequently used files to a more cost effective tier of storage. In the example shown in Figure 14, files are being migrated from fast SAS drives within one EqualLogic PS5000 Series disk array to affordable dense SATA drives in a second disk array.

**FIGURE 14. TIERED MIGRATION**



ESG Lab defined two pools of disk drives within the same EqualLogic array to test this capability. A migration rule was defined using an industry standard regular expression. In the example shown in Figure 15, the rule is being used to move files that are greater than 10 MB in size and have not been accessed in the last month from the SAS pool to the SATA pool.

**FIGURE 15. CONFIGURING TIERED MIGRATION**



The rule was enabled and monitored from the graphical user interface. After the migration job had completed (about five minutes later), a command line utility was used to list the files stored on IBRIX Fusion segments built over the SATA pool. As expected, all files matching the migration rule had been moved automatically by the IBRIX Fusion file system. The IBRIX Fusion file system remained online and available throughout the migration.

## Why This Matters

Organizations are susceptible to a wide range of data management hazards, including viruses, human error, hardware failures, power outages, and natural disasters. The downtime associated with these hazards can lead to lost revenue and lost productivity. As a matter of fact, ESG research indicates that 63% of organizations cannot tolerate more than four hours of downtime for their most critical applications before experiencing significant revenue loss or other adverse business impact.<sup>5</sup>

While snapshot and remote replication technologies have become available in recent years as a quick and reliable alternative to recovering from a tape backup, a number of emerging scale-out NAS solutions are missing these valuable capabilities due to the challenges of implementing these features while meeting extreme scalability and performance requirements. The extreme capacity needs of scale-out NAS solutions are also driving the need for the automated, cost-optimized movement of files between different tiers of storage.

ESG Lab has verified that scale-out NAS solutions from IBRIX support valuable enterprise-class NAS features including snapshots, continuous remote file replication, and tiered storage migration.

<sup>5</sup> Source: ESG Research Report, *Data Protection Market Trends*, 2007

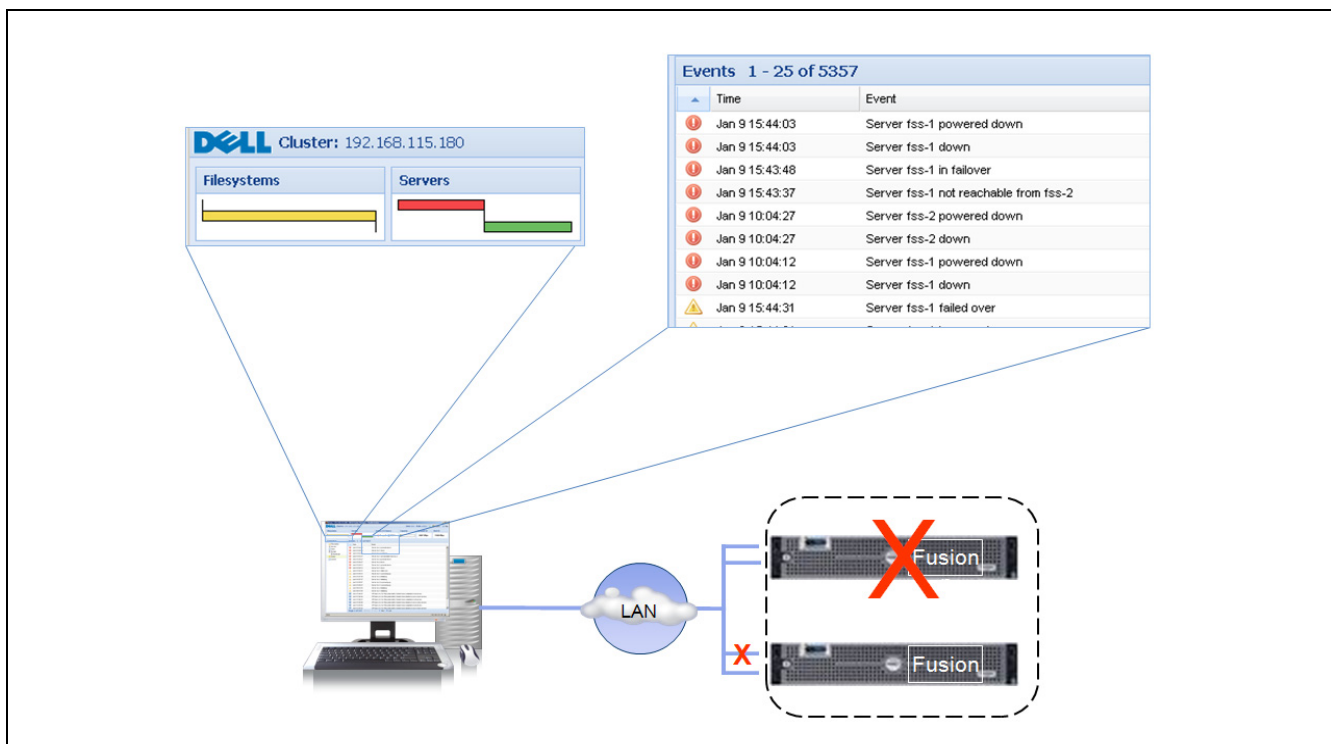
## High Availability

IBRIX uses redundant components and clustering technologies to ensure that the file system remains functional at all times—even in the unlikely event of a software or hardware failure. Clustering technologies are used to eliminate unscheduled downtime and protect data from:

- Server failures (e.g., failed processor or power supply)
- Storage failures (e.g., failed drive or EqualLogic array)
- Network failures (e.g., failed network cable or NIC)

ESG Lab validated the failover capabilities of a Scale-Out NAS solution from IBRIX by injecting faults into a two-node cluster. The first scenario ESG Lab tested was a node failure. The test started by writing some files to the IBRIX Fusion system. In order to simulate a node failure, a segment server (*fss1*) was manually powered down while data was being accessed from a Linux client. ESG Lab observed a short 30 second pause in I/O access to the mounted file system before I/O resumed through server *fss2*. Figure 16 shows the system status from the management console.

**FIGURE 16. FAILOVER TESTING**



Next, a LAN network cable was pulled while a client was accessing data from a file system mounted on a Windows platform. Files remained accessible—the scale-out NAS solution kept running smoothly. Once the server failover was completed, ESG Lab performed a failback operation from server *fss2* back to server *fss1*. There was no noticeable impact from an end-user standpoint.

## Why This Matters

Data unavailability can be extremely costly in terms of lost productivity and, in many cases, lost revenue. ESG Lab has confirmed that the combination of Dell hardware and IBRIX Fusion software provides transparent online recovery from hardware and cabling failures.



## ESG Lab Validation Highlights

- ☑ ESG Lab deployed a scale-out NAS solution from IBRIX in seven minutes and 14 mouse clicks.
- ☑ Clients were able to mount and access a multi-terabyte shared file system less than ten minutes after getting started.
- ☑ ESG Lab noted that the new and improved management GUI has an intuitive look and feel with expandable tree navigation and helpful graphical representations of system health and utilization.
- ☑ ESG Lab reviewed a number of successful production IBRIX scale-out NAS deployments in order to verify IBRIX's ability to support exceptionally large volume or fast performance environments.
- ☑ Additional EqualLogic storage capacity was added to an existing file system in ten seconds with no interruption.
- ☑ Advanced data protection and recovery capabilities typically associated with enterprise-class NAS appliances were tested. Local and continuously replicated remote snapshots were used to recover files after common errors, including accidental file deletion, file corruption, and power failure.
- ☑ The file system remained online and available after server and cable failures were introduced. File system failover was found to be fast and transparent to clients.

## Issues to Consider

- ☑ While initial configuration of IBRIX is a simple three-step process that can be completed in a matter of minutes, ESG Lab believes that a wizard wrapped around those three steps would further simplify and accelerate initial deployment.
- ☑ ESG Lab believes that extending the existing graphical representation of performance and capacity utilization statistics to include historical trending would be a useful enhancement. A graphical interface for the scheduling of routine activities (e.g., snapshots) is also recommended.

## ESG Lab's View

The massive growth of file data flooding data centers today—and the burgeoning wave that will be generated as more and more Web 2.0 and rich media applications come online—can easily overwhelm traditional scale-up NAS solutions. A scale-out NAS solution from IBRIX offers scale beyond that which can be attained with traditional NAS solutions. Users can start out small with relatively little capital investment and then grow to a massively parallel system. The performance ceiling is raised by adding more processors and capacity is increased by adding more storage, enabling “just-in-time” scalability. And management is simple because a scale-out NAS solution from IBRIX is managed as a single entity—no matter how large it gets.

Legacy scale-up NAS solutions—often growing to tens, hundreds, or even thousands of terabytes—face a number of challenges. First, scale-up systems typically have capacity limits in the tens of terabytes range. As capacity is scaled and limits are hit, more discrete systems are needed—and those systems need to be managed. Second, scale-up systems have fixed performance ratios; there is a fixed number of NAS heads that can be included in a single file system, typically one or two. Third, scale-up NAS has a relatively expensive price/performance ratio compared to scale-out.

A scale-out NAS solution from IBRIX combines the field-proven performance and scalability of IBRIX Fusion software, Dell PowerEdge servers, and Dell EqualLogic PS5000 Series iSCSI storage to cost-effectively address legacy scale-up NAS challenges. What's more, it's surprisingly simple to deploy. ESG Lab was accessing files seven minutes after getting started. It was also easy to manage via an intuitive graphical user interface. Additionally, it supports enterprise-class NAS features that are often missing in scale-out NAS solutions; ESG Lab tested snapshots, continuous remote replication, and online migration between different tiers of storage. Last but not least, it's fault tolerant and fast.

The IBRIX Fusion file system is fast. However, unlike many other scale-out NAS solutions, it's fast with just about any type of data. It can handle large files, small files, and lots of files. The IBRIX Fusion file system is good at servicing a large number of users concurrently accessing lots of different data, as well as a large number of users accessing the same data.

IBRIX has a heritage of success in the scale-out NAS market, with dozens of satisfied customers in the demanding high performance computing, online services, and video production markets. Pixar Animation Studios and DreamWorks are notable examples of the types of customers that rely on IBRIX for fast performance and high capacity storage. Dell EqualLogic storage systems have a similar heritage in the block-based storage market. Known for ease of use and cost effective clustering and Ethernet technology, thousands of customers have deployed petabytes of Dell EqualLogic storage.

Using an appliance-based approach that is fault tolerant and centrally managed, IBRIX has harnessed the field-proven power of its IBRIX Fusion software and Dell hardware to create an enterprise-class scale-out NAS solution that is extremely scalable, extremely fast for a wide variety of applications, and extremely easy to order, deploy, and manage.

# Appendix

## TEST CONFIGURATION

Local Site	
IBRIX File System Servers	2 Dell PowerEdge 1950, 4 GB RAM, 2.6 GHz Xeon
IBRIX Software version	Fusion Version 4.3 Beta
Storage	Dell EqualLogic PS5000 Series
Clients	Red Hat 5.2, Windows 2000, SP4
Network	Cisco Catalyst IP 3750
Remote Site	
IBRIX File System Servers	2 Dell PowerEdge 1950, 4 GB RAM, dual core 2.6 GHz Xeon
IBRIX software version	IBRIX Software Version 4.3 Beta
Storage	Dell PowerVault MD3000
Clients	Red Hat 5.2, Windows 2000, SP4
Network	Cisco Catalyst IP 3750



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

[www.enterprisestrategygroup.com](http://www.enterprisestrategygroup.com)