

*ESG Lab Review*

# Veritas Storage Foundation from Symantec: Performance & Efficiency Advantages for Linux

**Date:** October 2012 **Author:** Vinny Choinski, Senior Lab Analyst, and Kerry Dolan, Research Analyst

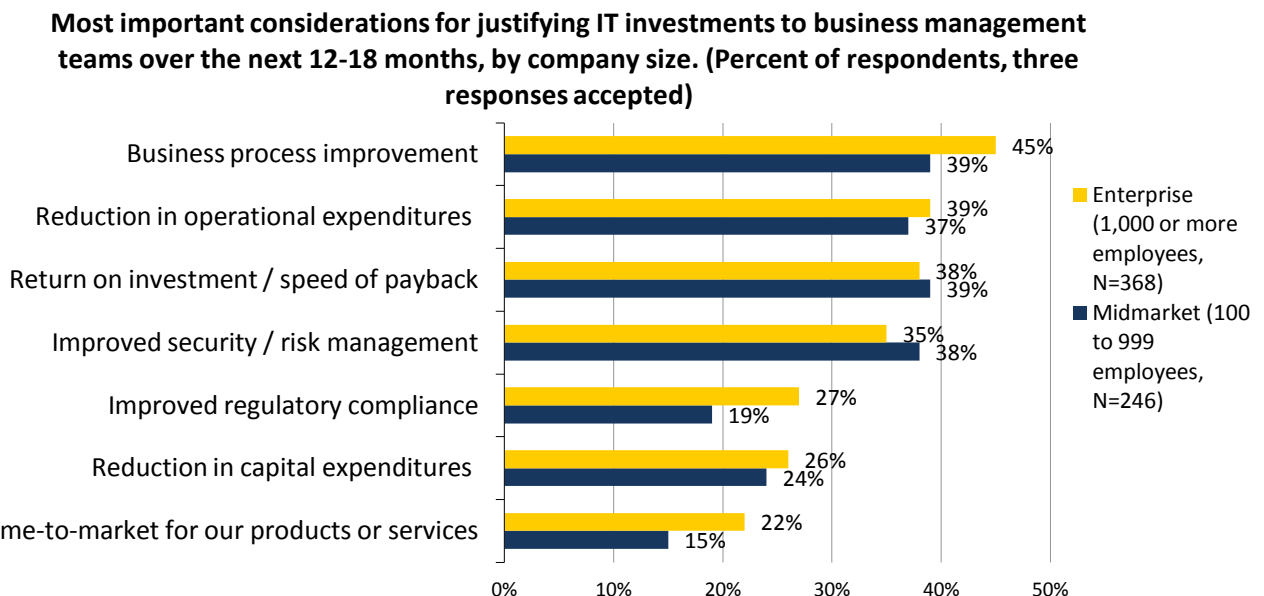
**Abstract:** This ESG Lab review documents hands-on testing and performance test audits of the Linux version of Veritas Storage Foundation from Symantec. Topics include performance, ease of management, and storage efficiency.

## The Challenges

Data centers are faced with daunting challenges in today’s business environment. Applications and data are needed around the clock to handle global enterprises and cloud computing, and the “consumerization” of IT results in end-user demands for high performance and always-on availability. Providing these levels of service across multiple operating systems and storage devices in physical and virtual environments is a juggling act for IT—one that is complicated by diverse, separately managed silos of infrastructure.

ESG research bears this situation out. In our 2012 IT Spending Intentions Survey, respondents indicated that while cost reduction remains important, improving business processes has risen to the top of the list. When asked which were the most important considerations for justifying IT investments to business management teams, 45% of enterprise organizations and 39% of midmarket organizations cited business process improvement, making it the number one response, followed by cost reduction and ROI (see Figure 1).<sup>1</sup>

Figure 1. Most Important Considerations in Justifying IT Investments



Source: Enterprise Strategy Group, 2012.

<sup>1</sup> Source: ESG Research Report, [2012 IT Spending Intentions Survey](#), January 2012.

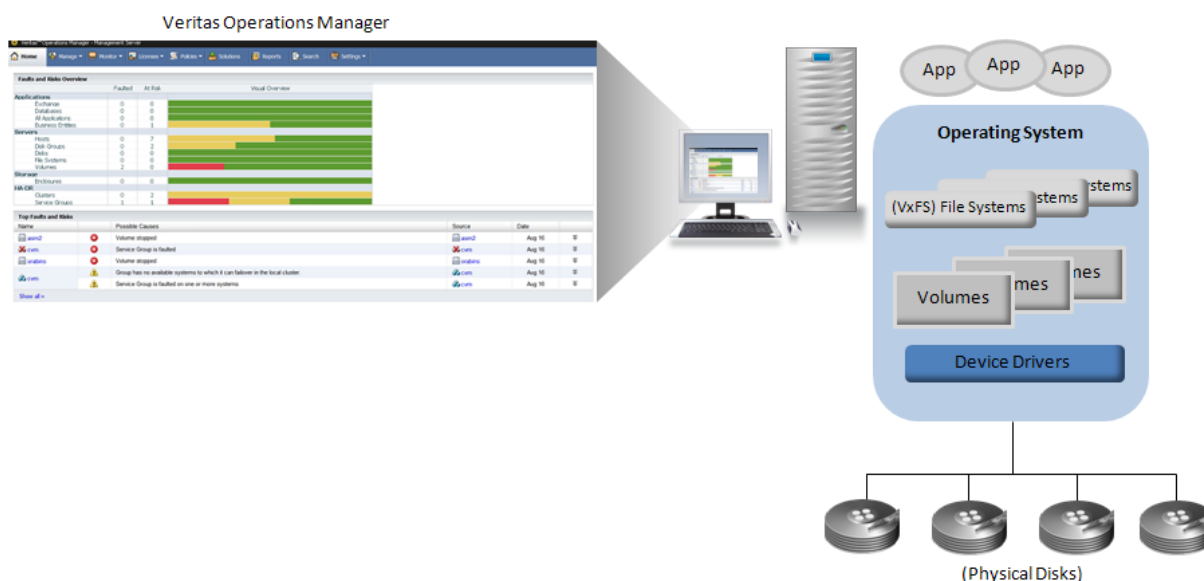
The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these 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 Symantec.

## The Solution: Veritas Storage Foundation from Symantec

Veritas Storage Foundation from [Symantec](#) includes the Veritas File System (VxFS) and Veritas Volume Manager (VxVM). These programs, combined with Veritas Operations Manager (VOM), which can be downloaded at no charge, provide extensive storage and storage management features for heterogeneous arrays.

- **Veritas File System**, a high-performance, journaling file system, provides simple management and fast application recovery; it delivers scalable performance, continuous availability, increased I/O throughput, and structural integrity.
- **Veritas Volume Manager** eliminates physical storage limitations, enabling administrators to configure, share, manage, and optimize storage I/O across disk and vendor boundaries. Online performance and storage management tools let you optimize your environment with minimal downtime.
- **Veritas Operations Manager** provides centralized monitoring, reporting, and management of storage resources.

Figure 2. Veritas Storage Foundation for Linux



## Features

Storage Foundation provides an extensive set of software capabilities that improve efficiency, performance, and availability for heterogeneous server and storage environments. Advanced features include the following:

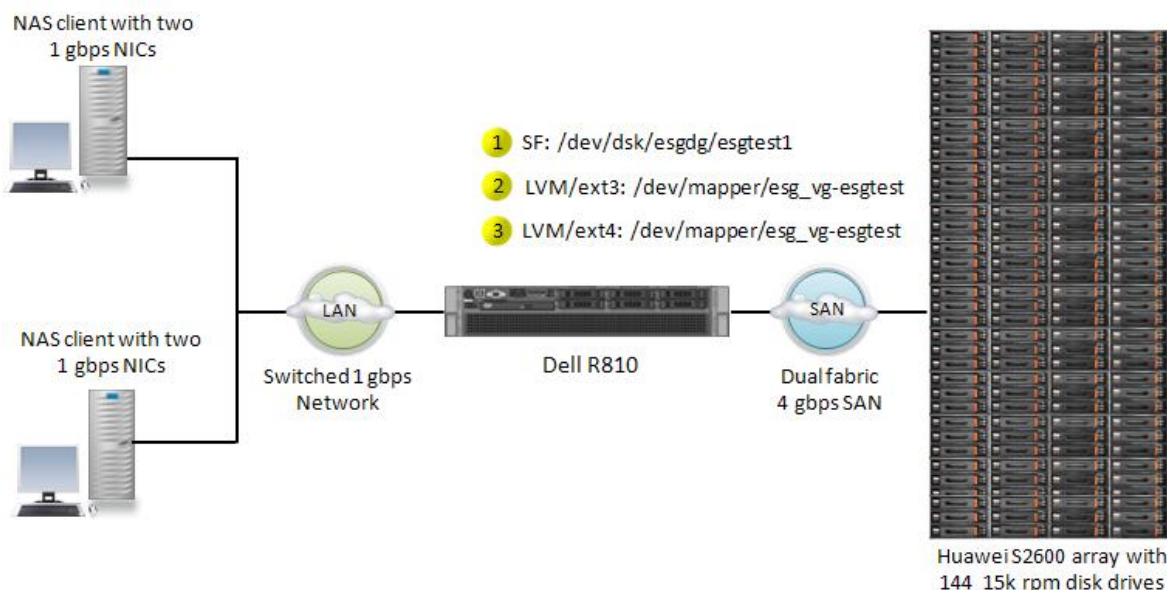
- A comprehensive view of storage enables better storage capacity management and efficiency across operating systems (including UNIX, Linux, and Windows) and storage hardware (including EMC, HDS, IBM, NetApp, HP, Fusion-IO, Violin Memory, and more). Dynamic, online provisioning improves flexibility, and automation of repetitive storage tasks (such as RAID reconfiguration, defragmentation, file system and volume resizing, and storage migrations across heterogeneous arrays) helps to keep applications online.
- Centralized application, server, and storage management with Veritas Operations Manager. A common management interface with visibility across the environment helps to speed deployment, improve service level delivery, and reduce errors.
- Deduplication and compression at the file-system level help to reduce the storage footprint despite continual data growth. Deduplication identifies and eliminates redundant data blocks to reduce storage needs by up to 80%, while compression can reduce storage requirements by up to 70% and free up bandwidth between servers and SANs.
- Dynamic Multipathing improves storage I/O performance and availability, balancing I/O loads across HBAs and rerouting I/O in case of path failure.

- SmartTier provides automated storage tiering, matching the importance of data with the appropriate storage availability and performance characteristics.
- The “thin-friendly” file system enables not only thin provisioning to speed storage provisioning as needed, but also thin reclamation. Instead of stranding storage resources when they are no longer needed, VxFS reclaims unused capacity and makes it available again. In addition, SmartMove enables efficient, host-based data migration across heterogeneous arrays, including moving from traditional to thin-provisioned storage with automatic reclamation, all while applications remain online. These capabilities can be managed using VOM.
- Portable Data Containers make it easy to move application workloads between UNIX and Linux operating systems, as well as supporting migration of legacy UNIX applications into virtualized Linux servers. Both tasks enable administrators to optimize performance and cost without downtime.
- Storage Foundation also simplifies array migration, which is a common and yet painful task required for rebalancing loads across storage arrays or when moving to new hardware when arrays come off lease. Storage Foundation simplifies the data movement, even across arrays from multiple vendors.
- Database accelerators for Oracle, DB2, and Sybase provide the performance of raw disk partitions with the manageability of a file system.
- Tasks such as RAID reconfigurations, volume and file system resizing, and snapshots can be managed through VOM or the command line interface (CLI), with automation of some tasks to reduce errors and free up staff resources.
- Local and remote data protection across heterogeneous operating systems can be done from a single tool. FlashSnap technology handles full or space-optimized point-in-time copies, accessible by multiple hosts, for jobs such as backup, testing, or decision support. Veritas Replicator, an optional add-on license, provides host-based remote replication between multi-vendor arrays for flexible, cost-effective protection.

Storage Foundation centralizes management while delivering high performance and availability across heterogeneous environments, simplifying IT service delivery and improving business processes.

To validate the benefits of this Symantec solution, ESG Lab leveraged the test environment shown in Figure 3. The left side of the figure shows a pair of NAS clients LAN-attached to a Dell R810 server. The Dell R810 was SAN-attached to a Huawei S2600 storage array. Both Storage Foundation (version 6.0) and the native Linux solution (RHEL 6.1) were exported from the Dell R810 server for performance testing, while general features were validated leveraging VOM in a separate lab environment. Volumes and file systems were re-created on the same disks to ensure consistency between test runs.

Figure 3. ESG Lab Test Bed



## Performance Advantage

The performance advantage of Storage Foundation for Linux is shown by comparing an audit and analysis of performance metrics for the Symantec solution compared with the native Logical Volume Manager (LVM) with Ext3 and Ext4 on Linux. ESG Lab conducted the detailed analysis by auditing performance test results of the three configurations running on the same hardware infrastructure using a network file serving workload.

### ESG Lab Tested

ESG Lab began the performance analysis by auditing the output results of an industry standard workload designed to measure file server throughput and response that was run multiple times against each configuration. Included with each run were detailed system trace logs that were also analyzed. As shown in Figure 4, each configuration was run to demonstrate the maximum number of transactions that could be produced for the workload on the same hardware infrastructure.

Figure 4. Maximum Number of Transactions per Second




Table 1. Detailed Results for Figure 5

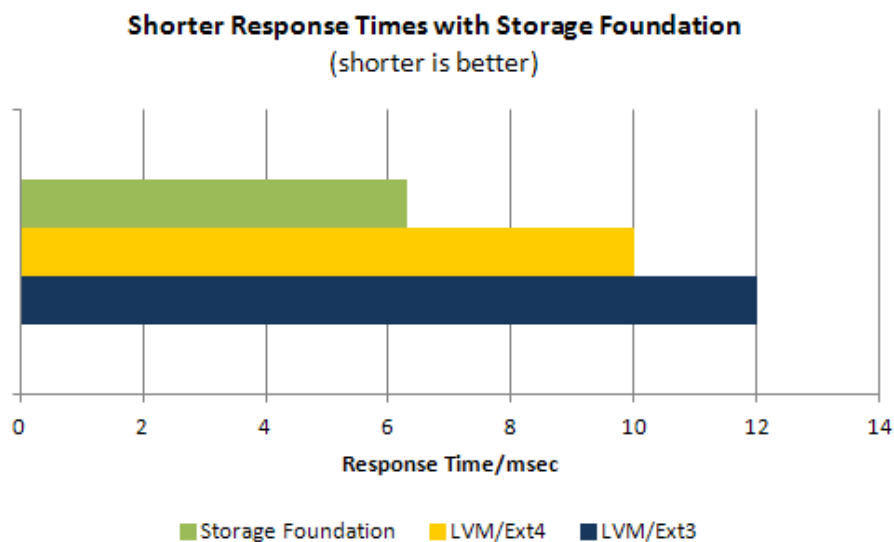
Performance Metric	LVM/Ext3	LVM/Ext4	Storage Foundation
Transactions/Sec	11,042	13,428	21,428
Response Times	12 ms	10 ms	6.3
CPU Utilization	53%	63%	91%

### What the Numbers Mean

- Storage Foundation was able to handle 92% more transactions/sec than Ext3 with the NFS file serving workload.
- Storage Foundation was able to handle 60% more transactions/sec than Ext4 with the NFS file serving workload.
- With Storage Foundation the test environment was able to better leverage the multi-core CPUs.

ESG Lab also explored the responsiveness of each of the three configurations. Figure 5 shows the response times in milliseconds for each configuration under the same workload on the same test bed hardware. The response times displayed in the figure are with each configuration running at its maximum number of transactions per second.

Figure 5. Response Time



**What the Numbers Mean**

- Storage Foundation was 90% faster in responding to requests than Ext3 with the file serving workload.
- Storage Foundation was 60% faster in responding to requests than Ext4 with the file serving workload.
- The response times showed an even greater advantage for Storage Foundation over the native Linux solution when running in the middle range of transactions per second. This means that with real world I/O loads, users would likely see even better system responsiveness with Storage Foundation.

**Why This Matters**

Storage performance is a major contributor to application performance—end-users may not see the storage bottlenecks, but they definitely feel them when applications drag. As a result, IT organizations are actively seeking solutions that improve performance, including increasing the use of solid-state drives (SSDs). In addition, as virtualization deployments expand and workloads become more and more aggregated, storage devices get hammered, and delivering on performance SLAs becomes increasingly difficult. ESG survey respondents report that *increasing* their use of virtualization remains a top priority,<sup>2</sup> suggesting that storage environments will become even more stressed and SLAs will become harder to meet.

ESG validated that Veritas Storage Foundation for Linux delivered better results than the native Linux solutions for throughput and response time, which are critical for application performance. The keys to Veritas Storage Foundation’s ability to deliver these high performance levels are both its algorithms and its ability to optimize CPU utilization.

<sup>2</sup> Source: Ibid.

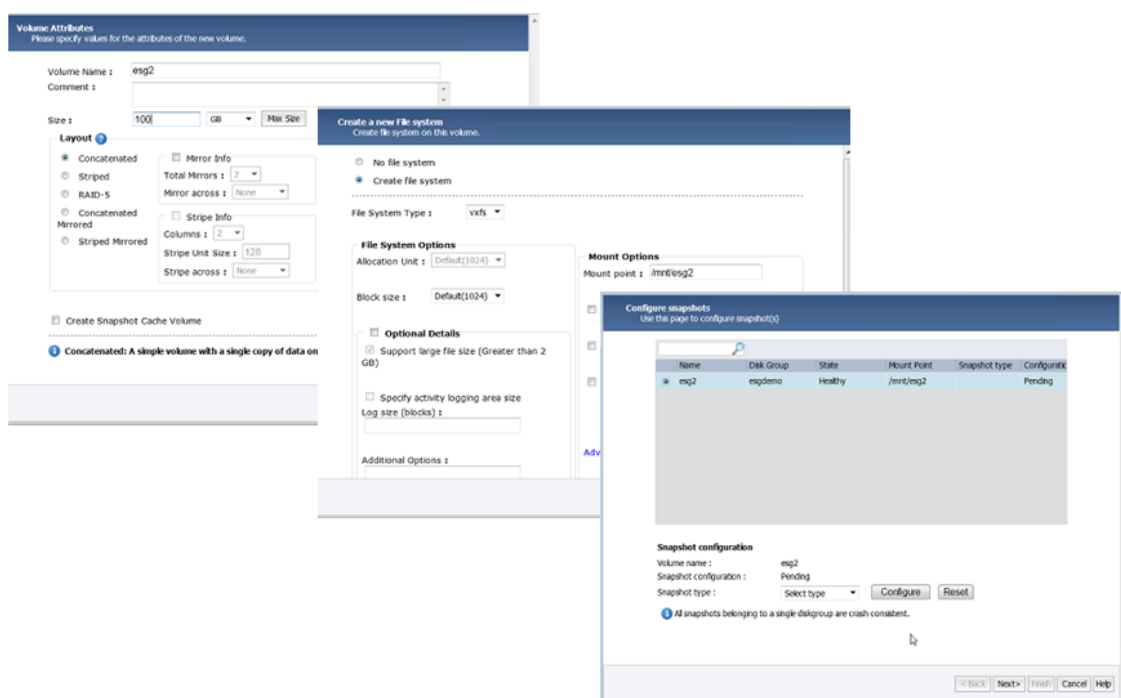
## Ease of Management

The graphical user interface (GUI) makes it easy to monitor and manage the environment. ESG Lab tested ease of management by creating a volume and file system, and running a multipathing test to identify any non-redundant paths between hosts and storage.

## ESG Lab Tested

With easy-to-use screen forms built into VOM and available from the command line, ESG lab created a concatenated volume called *esg2* on RAID-protected storage; mirroring, striping, and RAID options are also available. Next, ESG Lab created a new VxFS file system on that volume with a mount point of */mnt/esg2*, and configured a snapshot. These three screens are shown in Figure 6. Additional views allow administrators to see configuration and utilization details and manage storage resources. For example, storage templates can be used to create tiers, custom policies can be established, and numerous pre-built and custom reports are available. Storage Foundation can also be used for tasks such as shrinking a file system for better utilization.

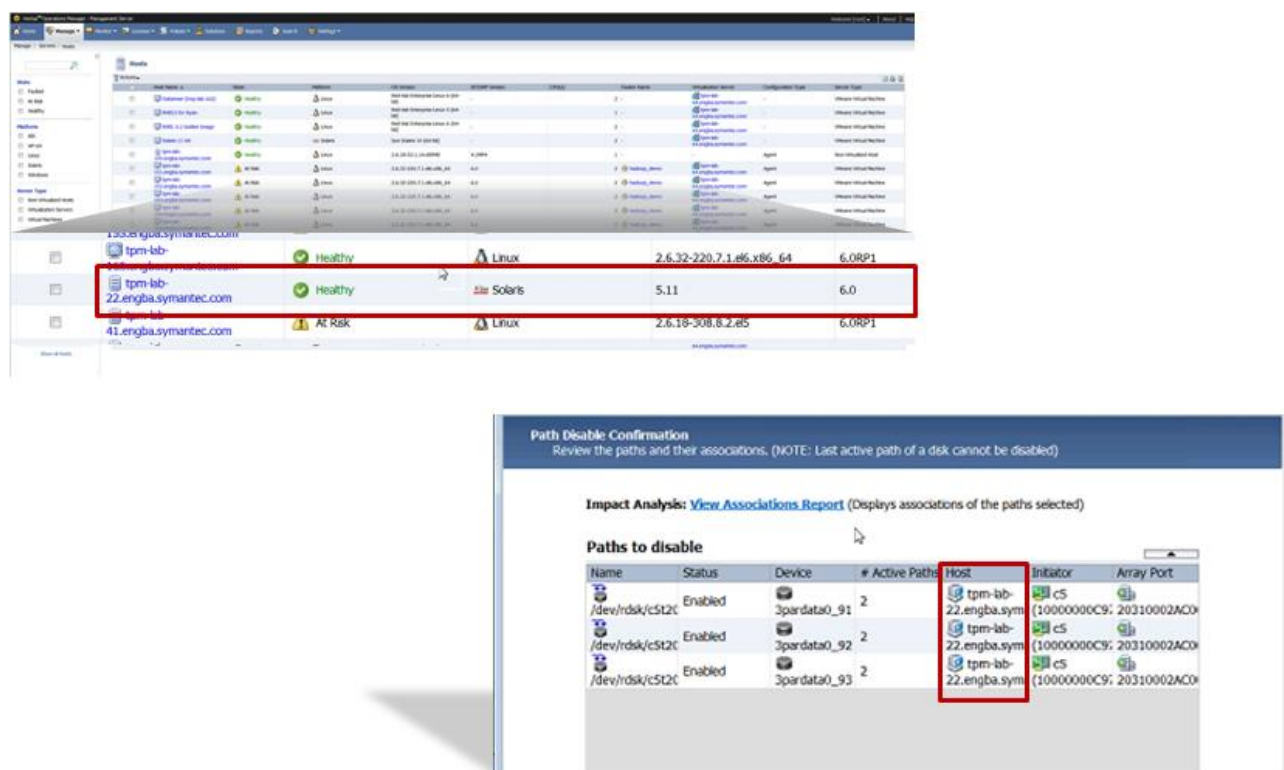
Figure 6. Creating a Volume, File System, and Snapshot



As shown in Figure 7, all systems can be managed from a single screen. In the top screenshot, all hosts are listed with their health status, platform, OS version, CPUs, cluster, server type, and configuration. The enlarged portion outlined in red shows that ESG Lab was able to drill down to a Solaris host (*tpm-lab-22*) and identify its status and configuration details.

Storage Foundation also provides powerful management tools that enable administrators to create policy and health checks to be run (scheduled or on demand) on all or selected hosts. These can help with tasks such as ensuring compliance, identifying patches and OS versions across a cluster, and viewing configurations prior to maintenance. While the system manages only its own toolset, it can monitor others so that administrators can view associations across the environment. ESG Lab tested this feature by running a dynamic multipathing policy test across the hosts to ensure path redundancy. The bottom of Figure 7 shows that two active paths were enabled between disks and host *tpm-lab-22*, so that one path could be disabled for controller maintenance without impacting data availability.

Figure 7. Dynamic Multipathing Policy Test



## Why This Matters

Manual management of storage tasks wreaks havoc with IT departments, particularly in environments with multiple operating systems, physical and virtual resources, and multi-vendor storage deployments. The task of gathering information using multiple spreadsheets and output from various components requires logging into individual servers, storage, and SAN components, collecting information, and combining it to understand the environment as a whole. Many organizations use different IT staff to manage parts of the infrastructure; this increases costs and results in errors as administrators struggle to keep all the parts in sync. Distributed management is often a barrier to growth due to the time required to keep silos of infrastructure operating effectively.

ESG validated that Storage Foundation provided extremely simple, primarily GUI-based management of storage tasks. Templates and wizards simplify common tasks, and the holistic view with drill-down capabilities delivers an overview of the complete environment. The ability to quickly identify configuration details eliminates time-consuming manual collection of information, and makes tasks faster, simpler, safer, and more reliable. For example, it helps to ensure that failover will occur without administrator intervention as needed should a problem arise, or that controller maintenance can be performed without application disruption.

## Efficiency

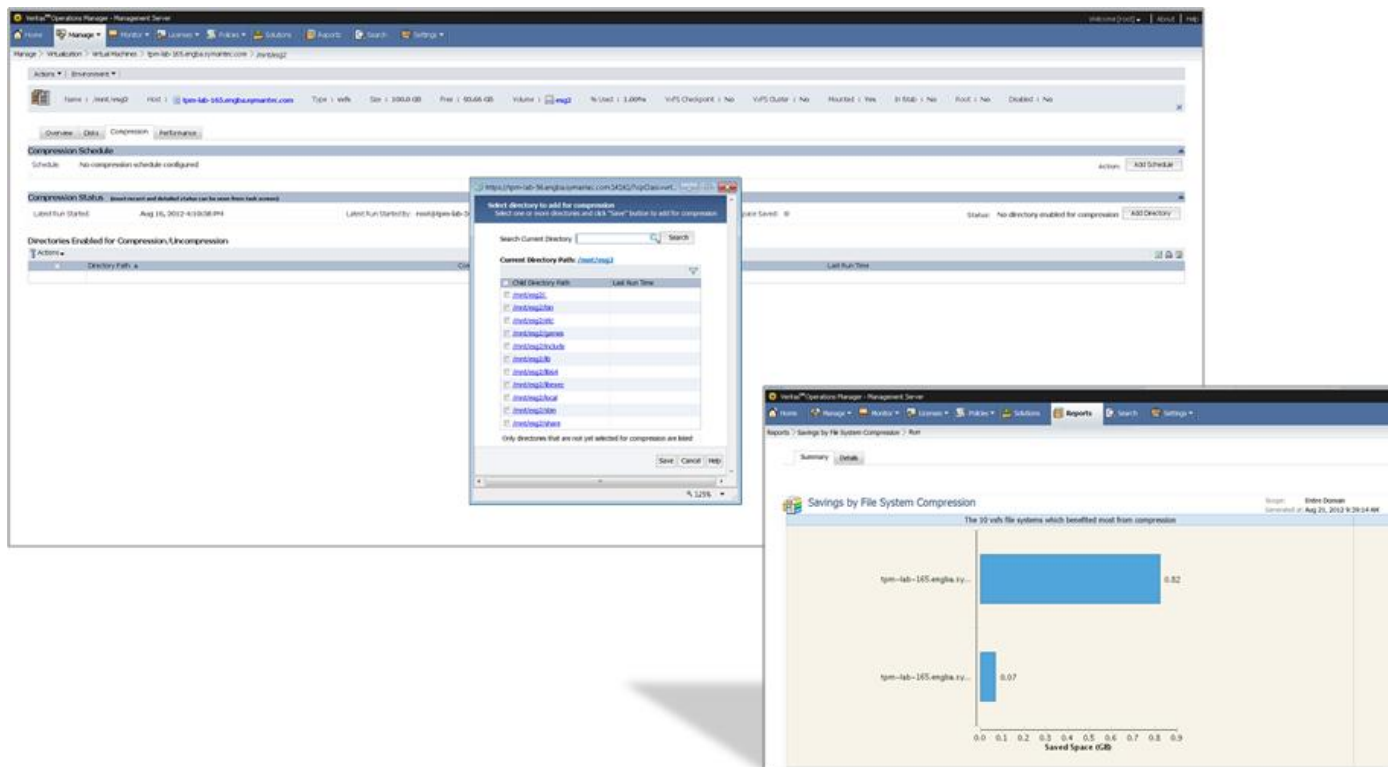
Deduplication, compression, and reclamation of thin-provisioned storage can improve storage efficiency. Thin provisioning enables easy scalability, but in addition Storage Foundation offers reclamation of thin provisioned storage so unused capacity does not remain allocated when it is no longer needed. Reclaiming stranded storage can be done on demand or automatically using the Thin Reclamation API, managed from VOM. This enables the storage environment to “get thin and stay thin.”

Storage Foundation provides periodic file system-level deduplication and compression outside of the write path, allowing writes to continue and minimizing application impact.

### ESG Lab Tested

ESG Lab used the Management tab to assign several directories in the *esg2* file system to be compressed, as shown in the top of Figure 8. Compression can be scheduled or completed on-demand. The Lab was able to compress a number of file systems and directories during validation testing. A summary screen in the bottom of Figure 8 graphically displays savings from the VxFS file systems benefiting most from compression.

Figure 8. Compression



Deduplication is CLI-based, and a “dry run” option can show the administrator how much space will be saved before executing the task. The chunk size can be tuned from 4Kb-128Kb depending on the data type. Figure 9 shows CLI output; the sections outlined in yellow show the initiation and completion of the tasks of scanning the file system for duplicate blocks, deleting redundant blocks, and showing space savings for that mount point.




Figure 9. Deduplication

```
[root@tpn-lab-165 test1]# fsdedupadm status /mnt/test1
Saving      Status      Node          Type          Filesystem
-----
04%        COMPLETED  tpn-lab-165.eng  MANUAL        /mnt/test1
          2012/08/21 11:25:41 Begin full scan
          2012/08/21 11:26:51 End detecting duplicates and filesystem changes.

[root@tpn-lab-165 test1]#
[root@tpn-lab-165 test1]#
[root@tpn-lab-165 test1]#
[root@tpn-lab-165 test1]# /opt/VRTS/bin/fsadm -t vxfs -HS shared /mnt/test1
Mountpoint  Size      Available    Used         Logical_Size Space_Saved
/mnt/test1  10 GB    7.99 GB     1.477 GB    1.544 GB    68.4 MB

[root@tpn-lab-165 test1]#
```



### Why This Matters

Data growth can have a significant impact on data centers, and that growth seems never-ending. Respondents to ESG research surveys chose managing data growth as the fourth most important on the list of IT priorities.<sup>3</sup> With laptops, tablets, and smartphones continually used to both create and consume content, the file data avalanche threatens to overwhelm IT. IT departments need more storage, more bandwidth, and more staff to manage the equipment—all costly propositions. In addition, as silos of infrastructure grow, so do management silos, complicating the landscape. Deduplication and compression can dramatically reduce data volumes, easing the equipment and management pressures.

ESG Lab validated that Storage Foundation provides easy-to-use deduplication and compression features. Of special note is the thin reclamation feature that enables administrators to reclaim unused storage. This feature increases efficiency and also allows IT to identify exactly how much space any department uses to simplify chargeback. The centralized management also makes management more efficient.

<sup>3</sup> Source: Ibid.

## The Bigger Truth

IT environments are increasing in complexity due to virtualization, cloud computing, and parallel infrastructure silos, while business managers want IT to do more with less. Combined with the onslaught of continual data growth, many IT infrastructures are becoming inefficient and costly to deploy and maintain. Technologies like virtualization have enabled new architectures and fast, easy provisioning—but their aggregated and mobile workloads can hinder performance and complicate management.

Storage Foundation pulls together a high-performance file system and volume manager with a centralized monitoring, reporting, and management interface. The extensive set of capabilities that this solution delivers can help IT administrators achieve the efficiency, performance, and availability requirements demanded of them much more easily and cost-effectively. Storage Foundation provides a centralized, holistic view of the environment and enables administrators to drill down into particular components for detailed information and troubleshooting. In addition, it simplifies a complex environment made up of diverse physical and virtual components while offering extensive storage management features.

ESG Lab tested the Linux version of the Storage Foundation, focusing on performance, ease of management, and efficiency capabilities. The solution outpaced native Linux file systems in transaction volume and response time performance, which are key to end-user satisfaction—and will become even more significant as virtualization deployments expand. In addition, ESG Lab came away very impressed with its ability to make a wide range of capabilities easy to use in a diverse environment. The centralized management feature can bring a sigh of relief to administrators struggling with multiple management tools, and efficiency features such as deduplication, compression, and thin reclamation will help them keep storage growth in check.