

# WHITE PAPER

## **Symantec Capitalizes on Veritas File System Expertise for Scale-Out NAS**

**By Terri McClure**

**June, 2009**

# Table of Contents

<b>Table of Contents</b> .....	<b>i</b>
<b>Keeping Pace with Data Growth Continues to be a Top Storage Challenge</b> .....	<b>1</b>
<b>An Efficient Approach to Dealing with Massive File Data Growth: Veritas SFS Scale-out NAS</b> .....	<b>2</b>
Users Examine More Cost Effective Storage Approaches .....	2
The Economics of Scale-Out NAS.....	4
<b>Summary</b> .....	<b>6</b>

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. This ESG White Paper was developed with the assistance and funding of Symantec.


# Keeping Pace with Data Growth Continues to be a Top Storage Challenge

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 in the face of increasing capacity requirements and storage growth. In fact, ESG data indicates that keeping pace with overall data growth is the top challenge IT managers face in their storage environment (see Figure 1).

ESG research also 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. TOP 10 STORAGE CHALLENGES FACING IT MANAGERS TODAY**

**In general, what are your organization’s greatest challenges with respect to its storage environment? (Percent of respondents, N=504, multiple responses accepted)**



Source: ESG Enterprise Storage Survey, 2008

The situation will only get worse thanks to advances in content capture, creation, and access devices like iPhones, HD video cameras, high resolution photography, and netbooks. Every endpoint device is now a content capture and sharing device that provides a new communication channel to potential customers. As a result, Web

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

2.0 is making its way into commercial data centers dealing with the next generation of end-users—the “Facebook Generation”—using rich media tools such as blogs, video, and online communities as a way to interact with customers and prospects. This will drive a whole new set of storage demands with new performance profiles for commercial data centers to deal with.

The massive amount of data generated by these new applications has a profound impact on already stressed IT infrastructures. The end result is often greater fragmentation in storage deployments, with stovepiped information infrastructures becoming extremely expensive to operate in terms of the number of employees required to manage the storage; excessive power, cooling, and floor space consumption; poor utilization rates; the list goes on. With large file sizes being driven by adoption of richer media formats, provisioning needs to be fast and easy as the demand for storage jumps by large chunks when even a single new file is added. The story isn't new, but the rate of data growth is accelerating to exacerbate the problem: Storage growth is quickly outpacing IT's ability to manage it. IT needs to find a way to consolidate and manage massive amounts of file data. Scale-out storage is the answer.

This paper explores how Veritas Storage Foundation Scalable File Server (SFS) provides an IT solution that enables users to efficiently manage more data with fewer people, saving money along both capital and operational budget lines. Scale-out NAS helps IT do more, with less.

## An Efficient Approach to Dealing with Massive File Data Growth: Veritas SFS Scale-out NAS

### Users Examine More Cost Effective Storage Approaches

Today, the challenge facing most enterprises is out of control file data growth. In a recent ESG survey, more than one in five medium-size businesses cited rapid growth in file-based content as one of their most pressing storage challenges.<sup>2</sup> As such, ESG expects customers to continue to make new NAS purchases to accommodate this growth and to drive capital and operational savings by consolidating sprawling file servers.


But consolidation and larger data files are not the only issues driving users to examine scale-out NAS solutions. The global financial crisis has driven IT to examine every new purchase with an increased focus on finding opportunities to reduce both capital and operational expenses. Technologies that reduce overall storage requirements or that drive higher levels of resource utilization are seeing a significant uptick in interest—and traditional ways of doing business are being re-examined to find ways to drive greater efficiencies. These activities are driving a corresponding interest in scale-out solutions in the enterprise thanks to scale-out's inherent scalability and manageability, which significantly reduces operational costs. During the current economic downturn, reducing operational costs has become the number one requirement for justifying IT spend, even ahead of reducing capital expenditures (see Figure 2), and scale-out NAS is a key enabler of operational savings.

---

<sup>2</sup> Source: ESG Research Report, *Medium-Size Business Server and Storage Priorities*, June 2008.

**FIGURE 2. PRIMARY CRITERIA TO JUSTIFY IT SPEND**

Which of the following considerations do you believe will be most important in justifying IT investments to your organization's business management team over the next 12-24 months? (Percent of respondents, multiple responses accepted)



Source: ESG Research Report, 2009 Data Center Spending Intentions Survey, March 2009

Symantec has quietly been making inroads into the scale-out NAS space since early 2008 with its Veritas Scalable File Server (SFS) solution. SFS builds on the firm's widely adopted Veritas Storage Foundation Cluster File System. It is a software-based scale-out NAS platform that provides a highly available, multi-tiered file storage environment capable of scaling into the multi-petabyte range. Moreover, it can leverage storage capacity already in the data center so that IT organizations do not need to buy more—as they would with a stand-alone NAS filer solution.

Scale-out NAS, which gives users the ability to independently scale and tune bandwidth, processing, and storage capacity on the fly—all while managing the file system within a single global namespace—is designed to overcome the management and performance challenges that arise when handling large amounts of file data. It also offers extremely granular scalability; users can start with as few as two nodes and scale up from there. Scaling granularly provides a price/performance advantage as it allows users to start small and grow where and as needed. And, since scale-out systems scale well into the multi-petabyte range and are managed as a single entity under a global namespace, the systems can meet most users' needs without incurring the management penalty often associated with deploying tens or hundreds of scale-up systems.

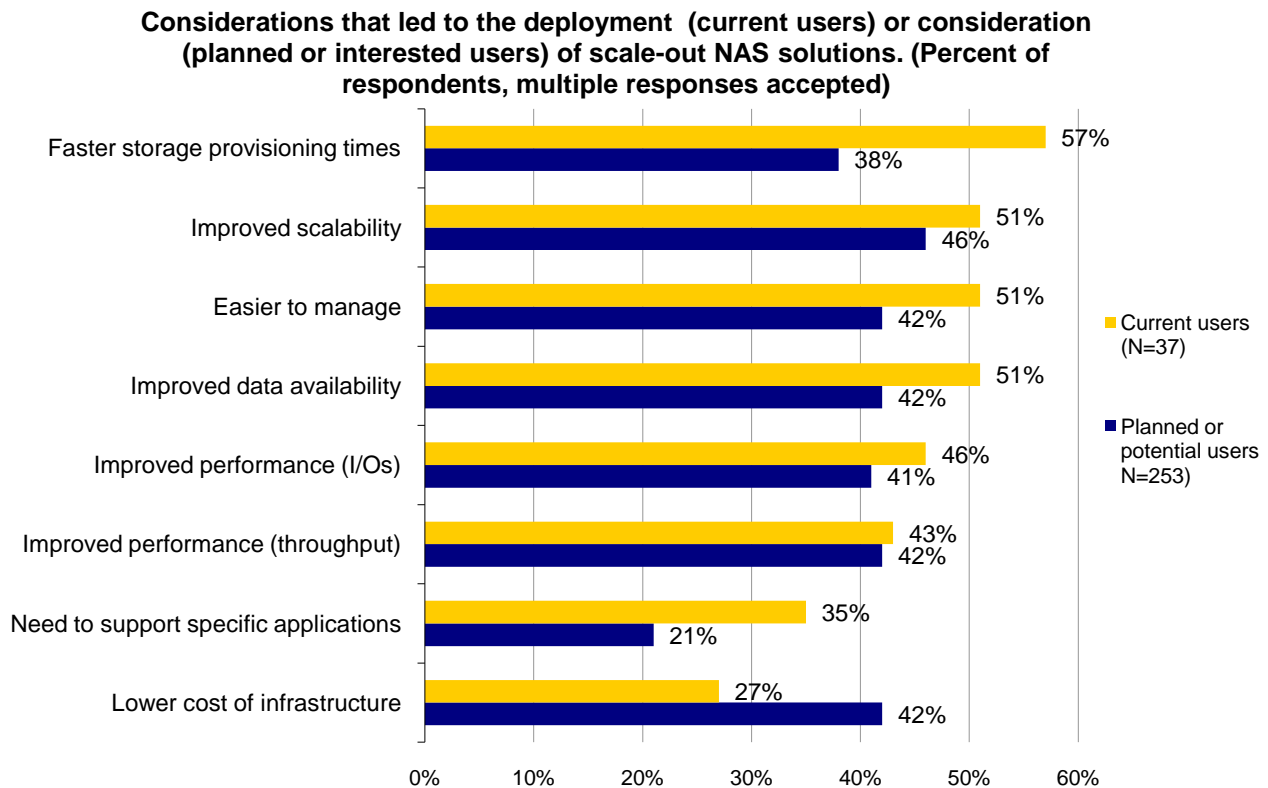
Scale-out storage architectures differ significantly from the scale-up storage architectures developed to meet distributed computing needs. Scale-up storage is just what it sounds like: it is designed to be monolithic, with lots of storage sitting behind one or two file server heads, and it is designed to scale into the multi-TB range behind those heads. Once the limit on storage is hit, a new system is installed with a new file system to manage. Scale-up systems can't scale bandwidth without some price penalty; high sequential performance in today's scale-up systems is often expanded by adding a storage rack and more spindles to increase throughput and reduce latency (and, as a byproduct, reduce storage utilization)—an expensive proposition.

## The Economics of Scale-Out NAS

Scale-out NAS is not a new concept. Users in high performance computing (HPC) and the media and entertainment markets were early adopters of scale-out NAS because of its bandwidth. As commercial enterprises adopt Web 2.0 applications and commercial HPC, scale-out is becoming a core data center requirement.

Initial cost has become a higher priority for users evaluating new NAS solutions than the value over time, though scale-out systems provide cost advantages in both areas. ESG recently conducted a survey of 504 North American and Western European IT professionals to assess data storage environments, including the adoption of scale-out NAS. Market drivers for early adopters included faster provisioning, improved scalability and performance, easier management, and the need to support specific, fast-growing applications. Lower cost of infrastructure was literally *last* on the list of buying criteria. However, planned and potential users have vaulted lower cost into the top tier of purchasing criteria—second only to improved scalability, which is the crux of the technology (see Figure 3).

**FIGURE 3. SCALE-OUT NAS ADOPTION DRIVERS**



Source: ESG Research Brief, *Scale-Out NAS Adoption and Market Drivers*, February 2009

Symantec brings a number of economic benefits associated with scale-out NAS platforms to the table with its SFS solution. Scale-out NAS architectures have a number of cost advantages over scale-up solutions, ranging from start up costs to managing technology refreshes—and many steps in between. Scale-out NAS carries a lower overall cost compared to scale-up systems for a number of reasons:

- **Its ability to scale capacity without scaling headcount:** With SFS, it is just as easy to manage a clustered storage system with sixteen nodes as it is to manage one with two nodes. Scale-out file storage systems enable this through a global namespace, which provides a single point of management for massive amounts of file data.

## Symantec Capitalizes on Veritas File System Expertise for Scale-Out NAS

- **Low entry cost:** The entry cost for scale-out systems varies depending on the minimum configurations supported. Most systems start as small as two nodes and scale-up from there. With clustered scale-out systems, you can add resources and scale as needed.
- **Just-in-time scalability:** As previously stated, because of the modular nature of scale-out systems, there is no need to buy (and power or cool) frames, power supplies, and mostly empty cabinets in advance of storage capacity.
- **Riding the commodity curve:** SFS is a software-based solution that can run on commodity servers and storage. Riding the processor and disk commodity curves as prices decline over time can add up to significant cost savings as the system grows, especially at the scale seen in these types of environments.
- **Higher utilization rates:** Better utilization means deferred purchases of new capacity (and better costs on the commodity curve). Because all of the NAS heads in SFS cluster can address the entire pool of usable capacity, no capacity is locked away behind underutilized NAS heads—a common problem in scale-up systems. It is not unusual to see utilization rates of 30% or less in scale-up systems and 60% or more in scale-out systems.
- **Reduced change management planning cycles:** When every endpoint device is a content capture device, conventional three or six month change management planning cycles are no longer effective. Requirements are unpredictable and time-to-provision is more important than ever. The modularity and scalability of scale-out NAS allow for extremely fast provisioning. SFS is plug-and-play; add a storage or processor node and the system self-discovers and expands the file system or incorporates it into load balancing algorithms on the fly. There is typically no disruption of service, nor is there a requirement to plan data layouts, create LUNS, or migrate data.
- **Non-disruptive technology refresh:** As with most scale-out systems, the process of managing technology refreshes with SFS is faster and easier than with monolithic NAS because the global namespace maps logical mount points to physical mount points in a virtualized manner, allowing back-end technology changes to be made with little or no disruption to client access.

Veritas provides extra features with SFS when compared to most offerings in the scale-out NAS market. First, SFS is delivered as a software-based solution, meaning that it stays true to the Symantec vision of being hardware agnostic. Users of SFS can leverage existing investments in servers and/or block storage devices. Second, SFS offers advanced features still fairly new to scale-out NAS, with built-in support for storage tiering and support for solid state disk. SFS integrated storage tiering allows for policy-based movement of files across storage tiers. Based on activity or desired performance levels, data can automatically be demoted to lower performing bulk-storage tiers, including MAID, or promoted to primary tiers if the data suddenly becomes very active. Consider the benefits this capability brings to environments with tiers that include solid state disk (SSD). For example, an SFS customer can initially choose to store data in a secondary tier (SATA) and then promote it to a primary tier based on IO access. SFS can then push the data back to secondary tier based on IO inactivity. This eliminates the need to peg data in valuable SSD real estate and then manually monitor and move it. All of this storage tiering and data movement occurs completely in the background, transparently to the end-user. The location of the files, and even their node numbers, does not change.

Recent ESG research indicates that users are applying scale-out NAS systems to new use cases. While most scale-out systems are tuned to perform well for high bandwidth applications, some can also support the smaller transaction-oriented file serving requirements of today's distributed computing environments. In fact, 43% of scale-out NAS users recently surveyed by ESG indicated that the technology is used to support database and OLTP transactions. SFS is one of the few scale-out NAS systems optimized for general purpose file sharing, with both high-levels of performance available in both throughput heavy operations (such as media streaming) and random IO operations (such as in data mining operations). Furthermore, SFS utilizes the same clustered file system used in thousands of mission-critical IT environments worldwide.

## Summary

Facing the possibility of a prolonged economic slowdown, users are looking for a number of critical qualities in a storage vendor. They want vendors they can trust; vendors that market proven solutions and products that offer real value in terms of cost savings and business agility. Keeping up with data growth driven by new types of applications, richer media types, and the ubiquity of content capture devices requires a new approach to storage spending. New rich media content is being created for everything from research and development, to training, to marketing, and is becoming a mandatory component of everyday business. Whether it's blogs, video, or HD imaging, content is easier than ever to create—and management will become harder than ever without significant changes.

Not every scale-out NAS solution is the same; there are significant differences. Most scale-out NAS solutions available today are designed to solve a single problem: supporting high bandwidth environments, like oil and gas or media and entertainment, where files can stretch to the multi-terabyte range. Many came from HPC and niche markets and they lack enterprise-class features that were not core requirements in those markets. Symantec SFS has a head start thanks to its Veritas Storage Foundation DNA—it is already an all-purpose file system built for the enterprise. Symantec is better rounded than most with a robust, high availability clustered file system platform; tiered storage support; and snapshot capability. Symantec has confirmed that in the next release of SFS—due late 2009 or early 2010—will provide two anticipated features: replication and a web-based GUI. Symantec has enterprise-class experience—it's already in the data center, not on an enterprise learning curve like so many other scale-out vendors trying to make enterprise inroads today.

Enterprises that deploy scale-out NAS solutions can get more value, dollar-for-dollar, from their infrastructure investments than they can with traditional scale-up solutions. Scale-out NAS has a compelling value proposition relative to scale-up systems. Its lower infrastructure costs, power efficiency, and management efficiencies should put scale-out solutions on the short list for anyone deploying new NAS capacity. Solutions like SFS, that have enterprise-class features, can easily be deployed into existing infrastructure; have a simple, intuitive set of commands; and can shrink and grow on demand without compromising file system availability are fast becoming a core IT requirement.

Three things in life are guaranteed: death, taxes, and information growth. Information growth will continue even in a down economy. This deluge of file-based information must be dealt with. Scale-out storage is the wave of the future. It is a path for IT managers to meet their number one storage challenge: keeping pace with overall data growth.

Because of these considerations, more and more enterprises are taking a serious look at scale-out NAS solutions like SFS—clustered scale-out solutions are going mainstream. But commercial enterprises are not just interested in the increased bandwidth scale-out solutions bring to the table: users are expanding use cases for scale-out NAS thanks to the higher scalability and manageability of these systems. In short, scale-out makes economic sense.



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

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