



Confidence in the connected world.

# **Veritas Storage Foundation Scalable File Server (SFS) from Symantec**

## Introduction to SFS and Use Cases

# Veritas Storage Foundation Scalable File Server (SFS) from Symantec

## Introduction to SFS and Use Cases

### Contents

<b>Introduction</b> .....	<b>3</b>
<b>Veritas Storage Foundation Scalable File Server Overview</b> .....	<b>3</b>
Gateway Model.....	3
Leveraging core strengths .....	4
<b>New SFS Components</b> .....	<b>5</b>
Installation .....	5
Administration .....	5
NFS Lock Management (NLM) .....	5
<b>SFS Key Benefits</b> .....	<b>6</b>
Consolidate and Reduce Costs of Storage .....	6
Increase Storage Utilization .....	6
Operational Cost Reduction.....	7
Dynamic Storage Tiering (DST).....	7
Consolidated Backup/Restore .....	8
Scaling and Seamless Growth .....	8
Modular Growth at the Processing Tier .....	8
Modular Growth at the Storage Tier .....	8
Near Linear Scaling.....	9
Availability, zero interruption of file services for company critical data .....	9
<b>Example Use Cases</b> .....	<b>10</b>
Enable Scale out compute clusters and Heterogeneous Sharing of Data .....	10
Enable virtual machine mobility among physical nodes.....	10
<b>Infrastructure checklist</b> .....	<b>11</b>
Server Hardware Requirements .....	11
Storage Hardware Requirements .....	11
<b>Summary</b> .....	<b>11</b>

## Introduction

This document outlines the various use cases for Veritas Storage Foundation Scalable File Server (SFS) product. SFS is essentially a specialized usage model of the Storage Foundation Cluster File System (SF CFS), hereby referred to in the rest of this paper as “core CFS”.

In most core CFS deployments the Network File System (NFS) components of application architecture are replaced by CFS. CFS in these cases provides more performant storage for latency sensitive applications via allowing a Fibre Channel (FC) transport for all application I/O to and from a shared file system. Examples of these are with applications such as SAP, TIBCO, Oracle RAC and others. CFS is a homogeneous architecture meaning that all nodes sharing the file system must be at the same operating system level.

SFS by comparison has a singular function, to share out NFS file systems in a simple to operate, highly scalable, and highly available manner. You cannot run applications on SFS nodes, thus it cannot perform the functions of core CFS in the scenarios described above. The applications run on the NFS client tier, utilizing the standard Operating System (OS) NFS client

While a customer could configure core CFS to only serve out NFS file systems, this is an unlikely use case due to operational heaviness associated with maintaining all the core CFS components, CVM, CFS, VCS, GAB/LLT, DMP, etc. SFS masks all of these components and provides a simple, appliance-like administration model for NFS file serving.

This document describes the scenarios where a CFS based NFS serving cluster is desirable. The document will not describe the technical details of core CFS or SFS.

## Veritas Storage Foundation Scalable File Server Overview


Storage Foundation Scalable File Server is a fixed function soft appliance based extensively on Storage Foundation Cluster File System. The solution provides a highly available and scalable NFS file serving platform with a simplified “appliance like” admin model. The applications consuming the storage are running a standard NFS client. The components of SFS include a SUSE Linux Enterprise Server (SLES) 9 Operating system, core CFS, and the new SFS components on a single DVD image.

## Gateway Model

SFS is built on a gateway architecture. SFS can be installed on an industry standard sever running Intel Xeon based processor architectures from IBM, HP or Dell. The servers can be blade or other form factors.

The SFS servers are attached to a SAN via 2 FC HBAs in each of the SFS nodes. Thus the solution is an open storage offering. Customers can attach different storage devices to the SFS cluster. SFS supports the same HCL as core CFS for the SLES 9 operating system.

Figure 1 shows the conceptual model of a clustered gateway model. The diagram shows the front end IP network to handle NFS client requests and the backend heterogeneous Fibre Channel based storage tier.



**Figure 1 Clustered Gateway Model**

### Leveraging core strengths

SFS leverages the capabilities and strengths of the core CFS stack. Some of these key features include Dynamic Multi-pathing (DMP), Cluster Volume manager (CVM), cluster file system (CFS), Veritas Cluster Server (VCS) and Dynamic Storage tiering (DST). Each of these layers represents industry-leading technology solutions and when combined into a file serving appliance provide an incredibly robust foundation on which to host NFS services.

DMP provides advanced FC HBA load balancing policies and tight integration with array vendors to provide in-depth failure detection and path failover logic as well as the broadest HCL in the industry. DMP itself is a large technical differentiator when evaluating SFS against similar offerings from other vendors as it contains the most advanced capabilities of any multi-pathing driver in the industry.

CVM provides a cluster wide consistent virtualization layer leveraging all the strengths of Veritas Volume Manager (VxVM) including on-line re-layout and resizing of volumes, and on-line array migrations. Many customers will choose to mirror their underlying SFS file systems across separate physical frames to ensure maximum availability on the storage tier. This is a common practice that has been employed by thousands of customers for years. This technique also facilitates seamless removal and additions of new arrays as they come off of lease cycles. Finally CVM scalability allows a single SFS cluster to manage thousands of LUNs which clearly differentiates it from anything else in the industry.

CFS provides full POSIX compliance; rich cache consistency and global lock management at a file or sub-file level. CFS also allows all nodes in the cluster to perform metadata “transactions” for the same file system thus allowing linear scalability in terms of NFS operations per second.

VCS provides cluster wide monitoring, communication, and failover for all nodes and their associated critical resources including virtual IP addressing failover for NFS client connections.

DST provides an industry unique intelligent file system capability which provides a rich, highly extensive policy engine to move blocks of files to lower cost storage. The placement and migration of files is completely transparent to any application accessing the file system.

SFS inherits all of these core capabilities and with the addition of a few key modules explained below, provides a comprehensive solution for NFS file serving.

## **New SFS Components**

The following outlines the new modules that SFS has added to the core CFS stack in order to provide a complete file serving solution.

### ***Installation***

A single node of the cluster is booted from a DVD containing the OS image, core CFS, and SFS modules. Once this node is up and running the rest of the nodes in the cluster (defined via IP addresses during the first node boot up) are automatically installed with all necessary components. The key services are then automatically started to allow the cluster to begin discovering storage and creating NFS shares.

### ***Administration***

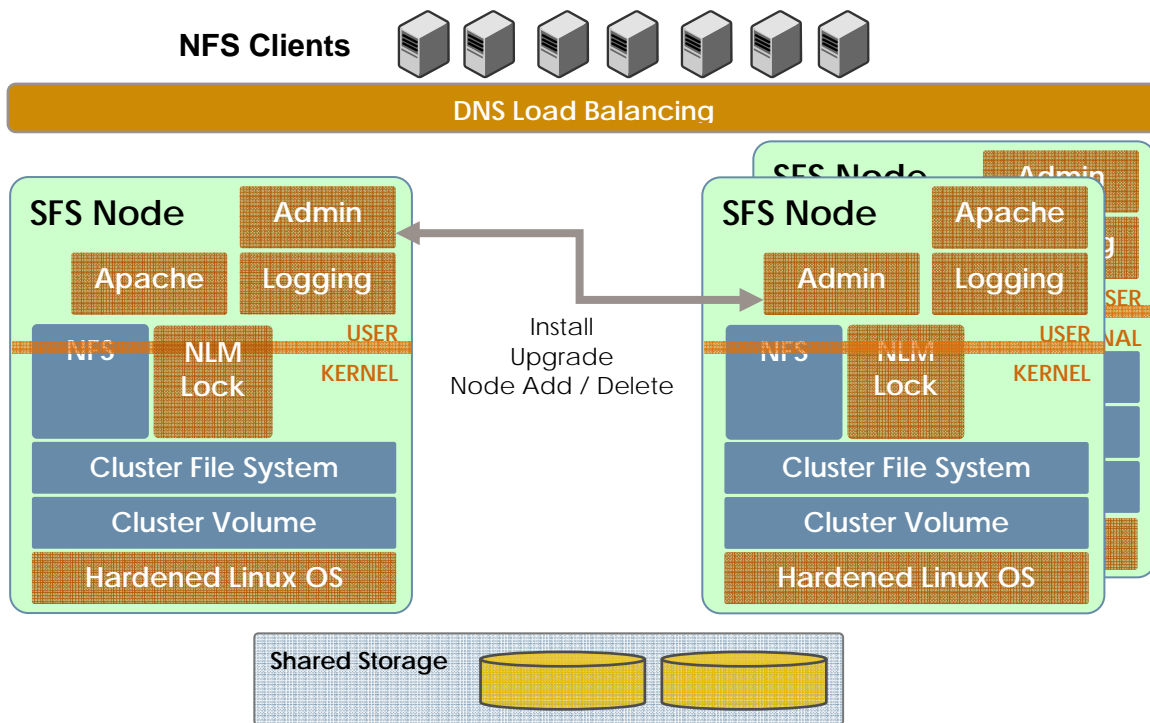
SFS contains a role-based administration model consisting of three key roles: Storage, Master, and Network. This delineation is consistent with the operational roles in many data centers. For each role the user uses a simple menu driven text interface. This interface masks any of the components and provides a single point of administration for the entire cluster. A user simply logs in as one of those roles on one of the nodes in the cluster and runs simple commands which perform tasks on all nodes in the cluster uniformly. There is no requirement to have any knowledge of Veritas Storage Foundation technology to install or administer an SFS cluster. However, it should be noted that users currently familiar with core CFS or SF in general will find familiarity with the basic management concepts.

Adding nodes is a simple non-disruptive operation and is highly automated following the installation model outlined above.

### ***NFS Lock Management (NLM)***

The NLM module allows a customer to use NFS advisory client locking in parallel with core CFS Global Lock management. The module consists of failing over the locks amongst SFS nodes as well as forwarding all NFS client lock requests to a single NFS lock master. The result is that no data corruption will occur if a user or application needs to use NFS client locking with an SFS cluster.

Figure 2 below depicts the architecture components of the SFS stack. The components in blue are elements leveraged from the core CFS stack while the brown components indicate the new SFS additions.



**Figure 2 SFS Component Architecture**

### SFS Key Benefits

The following represent the most common uses cases or likely deployment scenarios for the SFS product. This list is not exclusive or comprehensive but should provide the reader with a good understanding of how and when to use SFS to solve key business problems.

#### Consolidate and Reduce Costs of Storage

The value of consolidating several independent islands of NFS or NAS storage into fewer, larger shared pools has many cost benefits which are outlined below.

#### *Increased Storage Utilization*

Typical enterprises are running with storage utilization rates in the 30 – 40% range. This results in excessive spending on new storage when there is more than adequate free space in the data center. Ideally customers would group their storage assets into fewer, larger shared pools. This would naturally result in the increase of per backend LUN utilization. This increased efficiency can result in substantial cost savings from deferred purchases of new storage. SFS allows the creation of larger pools due to its scalable front processing capabilities. Eight nodes can provide up to 1.6 GBytes/sec of aggregate throughput which facilitates very large backend pools. This assumes each SFS node is configured with 2 GbE NICs for NFS client data traffic. More GbE NICs could be configured if necessary. NFS sharing allows this pool of storage to be securely shared amongst heterogeneous or homogeneous nodes. Thus SFS serving out NFS becomes an enabler for this common business strategy.

## ***Operational Cost Reduction***

The use of NAS heads from companies such as EMC and NetApp have resulted in the proliferation of many islands of NAS storage. As the number of independent islands increases the operational burden of managing the many NAS head-to-storage relationships becomes a significant time investment. Fewer, larger pools also reduces the ongoing management costs since there are a fewer relationships to manage now.

## ***Dynamic Storage Tiering (DST)***

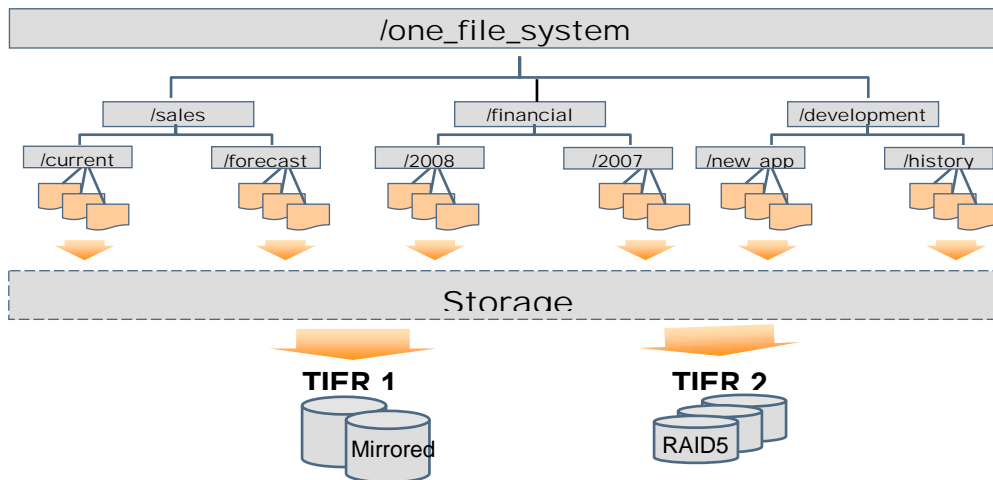
Few customers have adopted a storage tiering strategy even though the business benefits are clear. There are many limitations with the legacy or current storage tiering offerings in the market today that are causing this.

- Static storage tiering at project inception
  - Inflexible; Doesn't account for changes of data criticality over time
- Manual or Semi-automated Ad-hoc relocation
  - Applications and operating procedures must be modified to use relocated files
  - Apps and scripts must be adjusted/tested with every change in app/storage configuration
  - Operating costs are significant for managing storage consumption growth and other changes
- Hardware-based data relocation
  - Support for hardware (often only one model) from same vendor
  - Coarse granular tiering based on LUNs presented by disk arrays; no visibility of individual files and directories
  - Files to be relocated need to be organized so that they reside on LUNs scheduled for relocation. This is often impossible.
- Hierarchical Storage Management
  - Files need to be migrated back into the file system name space before the access request can be satisfied inevitably leading to delays and run-time overhead
  - File system must retain adequate free space to hold migrated files that are accessed by apps.

The DST feature introduced in the Veritas File System (VxFS), and brought over to core CFS and also SFS provides the unique ability to meet all of the key requirements customers need form a tiering solution.

- Create each file in its "optimal" storage tier
  - Based on pre-defined rules/policies
- Relocate files between storage tiers automatically as "optimal" changes
  - To take advantage of storage economies
- Retain original file access paths to minimize operational disruption
  - For applications, backup procedures, and other custom scripts
- Do all of this for **millions of files**
  - Typical in large data centers
- Quickly and without errors
  - Automate the above steps

Figure 3 below depicts the unique feature of SFS, the multi-volume file system, and how it maintains the application transparency. SFS can leverage this capability to further reduce costs of the backend FC storage by moving older data to lower cost storage.



**Figure 3 Multi-volume File System Namespace**

### ***Consolidated Backup/Restore***

A consolidated set of large storage pools reduces the management overhead of backing up and restoring disjointed file systems. SFS allows the collapsing of the file storage into fewer pools and thus reduces the backup interfaces and operations necessary. All critical file data can be backed up and restored through the SAN via integration with the market leading Veritas NetBackup (NBU) solution. Any of the nodes in the cluster can function as an NBU SAN media server and offload all the I/O for backups from the IP network. Additionally SFS can leverage CFS storage checkpoints to ensure a consistent backup is always taken.

### **Scaling and Seamless Growth**

SFS allows a customer to scale storage and processing independently and seamlessly (online). This provides great flexibility since a given application may need to scale in either dimension. The application may be primarily performance based and additional throughput may be needed for the same amount of storage; or perhaps the application is capacity bound and thus only needs to add backend storage without adding an SFS node. SFS provides seamless scaling in both dimensions in a very modular manner. You can add only the resources that the application(s) require without having to purchase a costly combined “set” of compute and storage resources.

### ***Modular Growth at the Processing Tier***

SFS automates the installation of a new node to the running cluster, configures it, and adds its capacity into the processing tier. The product scales from 1- 8 nodes (more nodes in the future). A user can add processing in units of single nodes, providing modular scaling that leverages cost savings compared to buying a large new and costly independent appliance, and without adding to the management burden.

### ***Modular Growth at the Storage Tier***

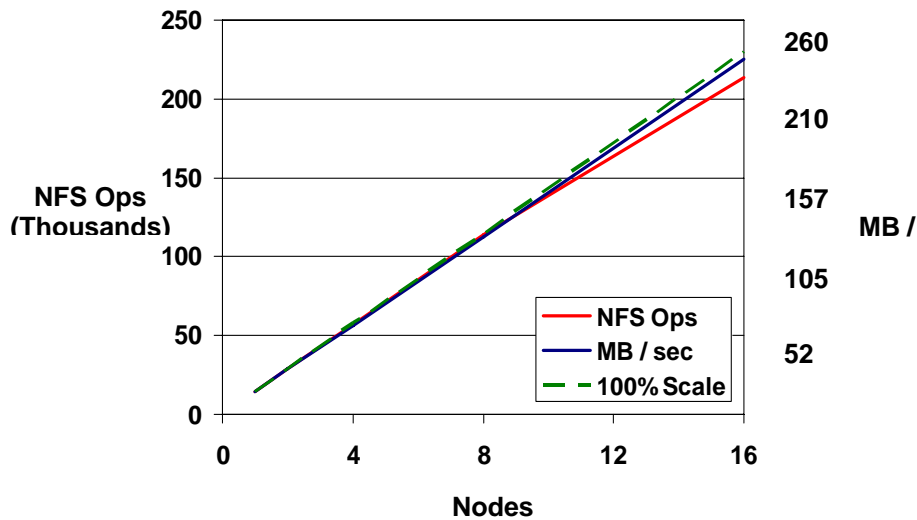
SFS uses CVM and DMP to discover and use new storage resources; this is also an on-line operation. A storage administrator can configure a new array or even add new LUNs from existing ones to the SFS cluster nodes. SFS can then scan the storage automatically, see the new LUNs, and place them under SFS control for use in the file systems.

Additionally at the storage tier all growth of existing file systems can be done on-line with no interruption of service. There is a simple command to add space to an existing file system which automatically uses newly available storage. Existing files system shares can be reduced in size in order to free up or reclaim tier-one storage. This is a unique offering of the SFS solution and is a powerful compliment to the Dynamic Storage Tiering solution mentioned earlier. Finally SFS provides the flexibility for a user to specify which LUNs a given file system should reside on. This can be critical for some performance based applications.

SFS growth allows a customer to adopt a pay as you grow model. This is often highly desirable for service providers or any customer who has implemented a chargeback model.

### Near Linear Scaling

The product provides nearly linear scaling in terms of both NFS operations per second as well as total I/O throughput. Figure 4 below depicts this scaling capability.



**Figure 4 SFS Scalability**

Note: SFS is initially qualified to support up to 8 nodes for the 5.0 product release

### Availability, zero interruption of file services for company critical data

SFS provides an “always-on” file service. The loss of single or even multiple nodes does not interrupt I/O operations on the client tier. I/Os are rerouted in approximately one minute. This is in stark contrast to the traditional NFS active/passive failover paradigm that companies such as NetApp, EMC, and others offer where failover times in the field have been measured in the 5+ minute range.

The convergence of two key trends is driving the capabilities that SFS delivers. The first is the dramatic growth of file based (unstructured) data assets. The second is the critical nature of these assets; they have to be available at all times to the application.

The architecture provides transparent failover for other key services such as NFS lock state, reporting and logging, and backup and restore operations. The installation service itself is also highly available and can seamlessly recover from the initially installed node failing during the install of the remaining nodes in the cluster. Finally, the console service providing access to the centralized menu driven interface is automatically failed over.

## **Example Use Cases**

### **Enable Scale out compute clusters and Heterogeneous Sharing of Data**

The trend to scale out, aka Grid, computing continues. There are significant performance and cost advantages of moving applications off of large Unix Symmetrical Multi-Processing (SMP) or Mainframe environments to a farm of compute servers running a distributed version of the application. One class of these sort of applications are Extract Transform and Load (ETL) products such as Abinitio and Informatica. These applications allow you run the computational processing on heterogeneous physically disjointed servers to allow the removal of all bottlenecks and maximize the benefits of parallelism. For the customer, running these jobs in a distributed manner allows the work to complete faster, enabling more orders or analysis to be done. The hardware costs for these compute farms is often a fraction of the cost of the large SMP environment. Most industry experts foresee the number of commercial applications that can run in a distributed manner increasing with time.

One of the key inhibitors to this scale out computing is the requirement to provide a shared storage infrastructure for the compute nodes. There are few options available today that enable you to share heterogeneously as well as scale up as performance requires. SFS solves both of these issues by providing a highly scalable and shared storage platform at the storage tier and by facilitating NFS heterogeneous sharing on the compute tier. Additionally the management of NFS client locks enables those applications that require using them.

An example architecture might look like the following: A HP-UX server functioning as a large (10TB) data staging and ingest server; a series of 6 Linux blades with the latest high end processors serving as CPU intensive compute servers for transformations of the data assets as well as sorting; and finally there is an AIX reporting server that needs to access the results to produce customer deliverables. All of these servers need access to the same files. A single NAS head or two heads in an active/passive configuration cannot provide enough NFS file serving performance to meet the throughput and response time requirements. Additionally the loss of a NAS head would result in a long running batch job failing and potentially having to be restarted. SFS can provide the performance and availability necessary for the NFS storage tier. It provides more than adequate throughput and seamless failover for this architecture.

### **Enable virtual machine mobility among physical nodes**

Moving virtual machines between different physical machines presents management challenges on block storage. FC masking is fragile and requires the masking of a single LUN to many physical HBAs in different hosts. This masking results in significant security exposures as well as the management burden of changing the masking on the fly. Finally some guest operating systems cannot dynamically discover the new storage device after the masking changes.

Traditional file storage interface (NAS) does enable mobility of virtual machines (files). Each physical host can share the same NFS exported file system and the underlying files. But complexity of management is a key driver in this solution.

Consolidation of storage resources is a common goal in the enterprise with great economic benefits, however, most NFS serving architectures do not scale as you add physical and virtual machines to the shared file system pool. This limitation on scaling prevents the consolidation into fewer more scalable pools that many customers desire with these virtual machine solutions.

## Infrastructure checklist

### Server Hardware Requirements

- The server hardware must be capable of running SUSE Linux Enterprise Server 9 SP3 for AMD64 and Intel EM64T.
  - Dual or Quad core processors at 3.0 GHz or above, are recommended for performance
  - Itanium is not supported
- 8 GB Minimum error-correcting code (ECC) random-access memory (RAM).
  - 16 GB recommended
- Internal Drives: Minimum single disk, dual drives recommended
- 4 Gigabit Ethernet NIC Ports
  - Dual embedded Gigabit Ethernet
  - 2 Additional Gigabit Ethernet
- 2 Fiber Channel HBAs
- Internal DVD Drive
- PXE Boot capable BIOS
- Redundant power supply

### Storage Hardware Requirements

Please Review the following Document from the public support website.

<http://seer.entsupport.symantec.com/docs/283161.htm>

### Summary

Veritas Storage Foundation Scalable File Server from Symantec is a clustered NAS gateway appliance that provides enterprise-wide NFS file services in a highly available and scalable environment. The server software is installed on industry standard hardware to give IT organizations flexibility in choice of hardware vendors. SFS is based on the industry-proven Veritas Storage Foundation Cluster File System.

Additional SFS servers can be added to the cluster without disruption and the file system can be grown dynamically without affecting file services. SFS combines the stability of a well-proven product with enterprise-class scalability, low cost of deployment, and a simplified administration model.

SFS lowers the cost of NFS file services by reducing the proliferation of storage islands, reducing the complexity of storage management and by improving storage utilization.

## About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world. The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at [www.symantec.com](http://www.symantec.com).

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation  
World Headquarters  
20330 Stevens Creek Boulevard  
Cupertino, CA 95014 USA  
+1 (408) 517 8000  
1 (800) 721 3934  
[www.symantec.com](http://www.symantec.com)

Copyright © 2007 Symantec Corporation. All rights reserved. Symantec and the Symantec logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

01/08