

Scale I/O Bandwidth on Linux Clusters, Simply and Reliably, with HP StorageWorks Scalable File Share, Based on Lustre™ Technology



Executive Overview	2
Linux Clusters Thrive in Prime Time	3
Lower Cost, Increase Power, and Simplify	3
Selecting a Scalable Filesystem Solution	4
NFS	4
Lustre Open Source Technology and HP Engineering	4
Benefits of HP SFS – An Overview	7
Scalable Bandwidth	7
Scalable Capacity	7
Scalable Resilience	7
Scalable Simplicity	8
Scalable Value	8
HP SFS Case Study	9
PNNL	9
Architecture of the HP SFS Product	10
StorageWorks Grid and Smart Cell Strategy	10
Major Components of HP SFS	10
Data Smart Cells	11
Metadata Smart Cells	11
Lustre Clients	11
LDAP Server	12
Message-Passing Interconnect	12
Options Available in HP SFS smart cells	12
HP Additional Value	13
Available Server Platforms for HP SFS	14
Expanding or Changing HP SFS after It Is Installed and Running	14
Accessing HP SFS Stored Data from Servers That Are Not Lustre-Enabled	14
Support Services Offered for HP SFS	14
The Future of HP SFS	15
Conclusions	15
For More Information:	16



Executive Overview

Linux clusters are scalable, computational engines that can deliver immense computing power—trillions of calculations per second—to meet the most demanding compute-intensive research projects. However, many Linux clusters use slow, shared I/O techniques, such as Network File System (NFS), the current de facto standard for sharing files. The resulting slow I/O can limit the speed and throughput of the Linux cluster. Additionally, programmers must often implement numerous time consuming and difficult techniques to use hundreds of disjointed, distributed file systems. As a result, applications run slowly and users waste valuable time and effort on filesystem housekeeping.

This paper examines the benefits of implementing HP StorageWorks Scalable File Share (HP SFS), based on Lustre™ technology. Lustre is an open, high-performance, highly scalable file share engineered to provide extremely fast, scalable, and reliable I/O for Linux clusters. Through rigorous product development and testing methods, HP augments the Lustre software with unique usability features on carefully selected and fully resilient StorageWorks hardware. HP is the first and only tier one vendor to offer a fully supported and integrated, hardware and software, Lustre-based filesystem. As a result, customers receive the benefits of leading edge, open standard technology combined with the simplicity, quality, and worldwide support of HP-engineered products.

HP StorageWorks Scalable File Share is recommended for Linux clusters running applications that require fast I/O across dozens to thousands of distributed compute nodes. HP SFS eliminates the need to pre-stage files, a time intensive and costly practice commonly used today to solve the I/O bottleneck. Its multi-GB/s speed delivers data quickly to applications as they request it, reducing the time to solution.

In this paper, we discuss the major challenges for Linux scalability and how to effectively combat the distributed file system I/O bottleneck. A brief background and critique of the current de facto standard, NFS, is explored, including the advantages of using Lustre-based, HP SFS for specific applications or in certain environments. In support of our positive assertions of the products' unique capabilities, we offer proven results of the Lustre implementation. Finally, the architecture of the HP SFS is detailed, including HP value-added features to the Lustre base code, product specifications, options, and future capabilities of the product.

Linux Clusters Thrive in Prime Time

Compute clusters are quickly replacing the legacy shared-memory parallel systems that have been the traditional technology used in complicated economic and financial modeling, drug research and design, and design performance simulations. Multiple servers in the cluster run portions of the application in parallel—transparently to the user—which can dramatically increase the performance of a single job. In many cases, compute cluster technology has been proven to provide a lower total cost of ownership and reduce time to market by performing complex simulations faster than shared-memory supercomputers.

Once found primarily in university or scientific research settings, the Linux compute cluster is becoming an accepted production technology for the enterprise based on its openness, flexibility, low cost, and reliability. HP is committed to develop and supply a wide range of Linux-based products and services targeted specifically at the clustering environment. HP has also forged important alliances in both the open source and commercial application communities and leverages them to benefit customers with expert systems design support and training.

Lower Cost, Increase Power, and Simplify

Cluster computing lowers cost by distributing workloads across numerous, less expensive, yet scalable platforms. Because clusters do not have a single point of failure, they are highly available for numerous mission-critical applications. In addition, because Linux clusters are based on open—rather than proprietary—standards, customers are not locked into one vendor's solution. As technology advances, new nodes can be attached and older technology repurposed, thus protecting hardware and software investments.

Yet there are some obstacles hindering the broad adoption of Linux clusters. As nodes are added to large clusters, the communications and administration loads affect performance and price. In the past, available management and support infrastructure for large, Linux-based clusters has not been robust enough for production sites. In addition, multiple source hardware and software components present maintenance and support issues.

HP, the leader in HPTC clustering technology, is removing these obstacles by simplifying and supporting the deployment of Linux clusters. The HPTC Linux cluster portfolio from HP simplifies cluster complexity, offering single, scalable management and application environments for diverse workload and user groups. HP lowers the total cost of ownership, delivering turnkey, supported cluster systems, and providing an industrial-strength environment for the most demanding applications. This cluster portfolio enables rapid adoption of new and increasingly cost-effective technologies, including the three leading industry standard 64-bit processors (Xeon™, Opteron™ and Itanium™ 2). Users can add functionality and capacity as needed, scaling from dozens to thousands of nodes, with limitless bandwidth and scalable mega-pixel capability for limitless visualization scalability. HP has strong relationships with best-in-class application partners, middleware providers, and members of the Linux community, ensuring superior support, maintenance and optimum performance of the Linux kernel.

Selecting a Scalable Filesystem Solution

Choosing the most effective, scaleable filesystem in a cluster environment depends on many factors, such as operating system of choice, applications being run, time, and budget. Below, we discuss and compare the history of the current de facto standard, NFS, and the new, Lustre-based HP SFS solution.

Criteria for selecting HP SFS

The main criteria for selecting HP SFS include the following:

- Installation of a current or future Linux cluster
- I/O bandwidth or capacity that is larger than easily supplied by an NFS server
- I/O bandwidth from hundreds of MB/s up to tens of GB/s
- Management of 10s, 100s, and 1000s of terabytes of storage in a single filesystem
- The simplicity of dozens to thousands of Linux clients simultaneously accessing a shared filesystem and shared files in parallel

Other characteristics not always found in other distributed filesystems, include:

- Coherent I/O, with logical buffer management (as if one were using shared-memory buffers on an SMP).
- A full POSIX-compliant filesystem interface, a feature lacking in NFS.
- Efficient intent-based record locking.

NFS

NFS (Network File System) is the current industry standard for NAS (Network Attached Storage) sharable storage on UNIX and Linux servers. Its scalability and lack of coherency can be limitations for some high bandwidth, I/O intensive applications, causing an imposing I/O bottleneck. However, if scalability and coherency are not issues, then NFS is a viable solution.

UNIX and Windows-based clusters tend to use existing NAS standards, such as NFS and CIFS (Common Internet File system from Microsoft) because Lustre has not yet been ported to support these operating systems. Lustre clients can and probably will be implemented on non-Linux platforms, but as of this writing, Lustre is available only on Linux.

Additionally, organizations are not comfortable with new technology or who do not consider themselves early adopters, may want to continue using NFS while they examine the correct time to move up to Lustre. This new Lustre protocol is running at dozens of sites and is performing reliably at large computing sites such as Pacific Northwest National Lab (PNNL) and Lawrence Livermore National Lab (LLNL). Lustre-based solutions are ready for organizations that need leading edge technology to solve I/O bottlenecks in their Linux clusters.

The reliability and performance of Lustre is already well advanced and will become more so as it is deployed at hundreds and eventually thousands of sites. Lustre's rapid adoption rate is expected to lead to a broad support in the Linux community followed by Lustre clients on a variety of operating systems.

Lustre Open Source Technology and HP Engineering

Lustre is an open source standard protocol and reference code initially written by Clustered File Systems, Inc. (CFS). Lustre was funded by the DoE through the Hendrix project with ASCI Path-Forward funds to design the system, develop the code, test it at scale, and bring it to market as

open source GPL code. HP co-funded the Hendrix project and works jointly on this contract with CFS¹ and Intel.

The DoE required an open source sharable filesystem with scalability to match the largest Linux clusters in the world, such as those installed at their national laboratories. Lustre is architected to match the continuing and rapidly increasing scalability of clusters, orders of magnitude larger than the initial and extremely large cluster installations in 2003 and 2004. Thus, it has been initially deployed on 1,000 node clusters, but is designed to scale to serve 10,000 and even 100,000 node clusters.

HP SFS advances Lustre technology to the next level by simplifying the use and administration of Linux clusters to deliver faster processing and a higher return on investment. HP's unique and valuable engineering transforms the well-crafted, open source, Lustre distribution into the HP SFS product that is easy to install, manage, and use. HP SFS is highly reliable, highly scalable, and easily deployed, delivering ten to a hundred times more bandwidth than standard NFS servers.

HP SFS is a self-contained fileserver built from multiple, industry standard HP ProLiant servers and StorageWorks disk arrays. The HP SFS server runs a combination of Lustre and HP-specific value added software. HP organizes these parallel servers, storage arrays and software into a highly reliable product with the simplicity of a single, scalable file system. HP SFS is typically paired with Linux compute clusters, such as high performance technical computing (HPTC) clusters of HP ProLiant or HP Integrity servers. HP SFS support is built into all HP XC clusters and is an integrated option for all cluster configurations in the HP HPTC cluster portfolio.

HP SFS demonstrates HP's commitment to use industry standard and open technologies to ensure maximum long-term customer value, stability, and agility. HP SFS combines accessible, open source technology with a well-engineered product. The filesystem is fully MPI-I/O and POSIX compliant, which includes file locking, coherency, and other standard and easy-to-use behaviors that are traditionally delivered on SMP-based filesystems.

When deploying new technology, organizations look for a partner with experience and leadership to help ensure that success will come quickly, easily, and painlessly. HP leads the industry in investment and deployment of Lustre technology and is by far the market-share leader in open SAN technology². HP is the first and only major vendor to offer a fully supported and carefully engineered Lustre-based file server product, stemming from its joint research and development project with the DoE and CFS. The DoE selected HP to provide program management and development services to support the Lustre project. With decades of cluster filesystem experience and leadership in Lustre technology, HP provides a solid foundation for the Lustre technology. Figure 1 below summarizes features of NFS compared with HP SFS.

¹ Cluster File Systems, Inc. (CFS) is the principle developer of the Lustre architecture and code base. HP and CFS cooperate to support Lustre and advance the adoption of its open technology..

² Source: IDC and the Gartner Group: HP has three times the market share of any competitor.

Figure 1. A comparison of the features of NFS and HP SFS.


	NFS	HP SFS
Network Attached Storage	Yes	Yes
Open standard	Yes	Yes
Scalable bandwidth	multiple 100 MB/s	to 10s of GB/s
Scalable bandwidth	limited	10 to 100 times faster
Scalable filesystem capacity	Usually 16 TB or less	100s or 1000s of TB
Scalable metadata	No	Planned for v2
Smart cell agility	No	Yes
Add b/w as needed	No	Yes
Coherent I/O	No	Yes
POSIX API	No	Yes
POSIX record locking	No	Yes
Interconnect neutral	No	Yes
Supports Lustre protocol	No	Yes
Ubiquitous standard	Yes	Not Yet
Low latency I/O	No	Yes
Low overhead I/O	No	Yes

Benefits of HP SFS – An Overview

Scalable Bandwidth

With HP SFS, bandwidth is scaled by distributing the files in parallel across a scalable StorageWorks grid, a standards-based architecture developed by HP that enables storage services to be delivered across a massively scalable, centrally managed system. The StorageWorks grid architecture divides storage, indexing, search and retrieval tasks across a distinct set of computing nodes or storage “smart cells” that cooperate to form a single shared filesystem. Each HP SFS smart cell is an intelligent storage server, running the Lustre protocol, which works in parallel with other smart cells on a shared StorageWorks grid. To scale to the desired bandwidth level, users simply add smart cells. The smart cell is built with the agility of an object-based storage framework, which is architected for increasingly sophisticated features that improve responsiveness, security, reliability, and resiliency.

Figure 2. An overview of HP SFS



Scalable Capacity

HP StorageWorks Scalable File Share also delivers potentially unlimited scalable storage capacity (from terabytes to petabytes of storage) by simply adding smart cells as needed to existing StorageWorks grids. Because each smart cell contains additional storage capacity, this “pay as you grow” agility allows an organization to become an Adaptive Enterprise by expanding their capacity on existing filesystems as needed.

Scalable Resilience

Lustre technology in HP SFS is designed to scale while maintaining resiliency. As servers are added to a typical technical cluster environment, failures become more likely due to the increasing number of physical components. Lustre’s support for resilient, redundant hardware provides protection from inevitable hardware failures through transparent fail-over and recovery.

Although resiliency is elegantly architected in the Lustre code base, a simple port of the open code is not sufficient to maintain true resiliency in a highly complex, scaled environment. Without rigorous attention to the details of matching the hardware to the software, potential reliability might not be delivered. HP solves this challenge by carefully designing and selecting resilient HP hardware with no single points of failure, and then adding key HP-specific, fail-over hardware

and software components. Additionally, HP thoroughly tests these configurations to ensure that reliability is delivered as the servers are scaled.

Scalable Simplicity

HP SFS can span dozens to thousands of Linux clients, dramatically simplifying the ability to run clustered applications. HP SFS removes the I/O bottleneck, saving users hours of programming time. With HP SFS, users avoid the complexity of running applications on many individual, distributed file systems.

With HP SFS, applications see a single, high-bandwidth filesystem image regardless of the number of Linux clients. HP SFS is performance-tuned, offering ease of use and a lower cost of ownership. The result is a centralized, easy to use storage system that is also simple to manage and expand. In addition, because the interface for these Lustre filesystems is fully POSIX compliant with all of the standard interfaces in place, programs can run without modification.

HP SFS also simplifies installation, administration, and support. HP adds significant features on top of the open source code base that dramatically reduce the work required to deploy and operate a distributed-scalable file share. From a single console, using simple and automated HP-specific tools, administrators can easily install and update software, create and update filesystems, detect errors and monitor performance, thereby streamlining the system operation.

Scalable Value

The StorageWorks grid and Lustre architectures enable the use of industry-standard components and open technology that lowers the cost of the file share while delivering unprecedented scalability. This increased scalability lowers total cost of ownership (TCO) by delivering easy to use, centralized management. Filesystems are easy to access and can be expanded as needed with additional smart cells.

HP SFS Case Study

PNNL

The U.S. Department of Energy's (DoE) Pacific Northwest National Laboratory (PNNL) has installed one of the largest Linux clusters in the world to address biological, chemical, and environmental challenges. For example, the massive Linux cluster at PNNL is being used to help model environmental cleanup scenarios. PNNL wanted to create a supercomputer that would allow scientists to model grand challenges that wouldn't be feasible to test in real life. Using the cluster supercomputer, scientists are able to more rapidly model a wide variety of substances, including enzymes, in their quest for the most appropriate cleanup material.

"This enormous cluster is centralized, easy to use and manage, and simple to expand with a single shared Lustre filesystem. Its resilient and cost-effective scalability translates into faster time to solution and better science."—Scott Studham, Associate Director for Advanced Computing at PNNL.

HP and PNNL partnered on the design, installation, integration and support of the cluster, resulting in one of the "Top 10" fastest systems in the world. The HP Linux supercluster, with more than 1,800 Itanium 2 processors, is rated at more than 11 TFLOPS. PNNL has run Lustre for more than a year and currently sustains over 3.2 GB/s of bandwidth running production loads on a 53 terabyte Lustre-based file share. Individual Linux clients are able to write data to the parallel Lustre servers at more than 650 MB/s.

"HP's Lustre implementation in our HP Integrity-Linux supercomputer allows us to achieve faster, more accurate analysis so we can better understand how numerous scenarios impact our environment. We attain answers faster from complex, I/O hungry applications, because Lustre scales the high-bandwidth I/O needed to match the large data files produced and consumed by our scalable simulations. HP has worked with PNNL to ensure Lustre is highly reliable and stable. We have never lost any data entrusted to this scalable-shareable filesystem." Scott Studham, Associate Director for Advanced Computing

The capabilities of the Lustre technology combined with HP Integrity servers allow scientists at PNNL to achieve the scalable bandwidth, scalable capacity and scalable resiliency needed for highly sustained performance that is also I/O intensive.

Architecture of the HP SFS Product

StorageWorks Grid and Smart Cell Strategy

Over the years, HP has built a leading portfolio of network storage solutions that have revolutionized IT environments worldwide and has provided the foundation for the storage utility. Now HP is taking the next step in delivering that utility by defining the need for real time information services and a plan to deliver them on a StorageWorks grid. The StorageWorks grid describes the physical infrastructure and software upon which services are delivered and the storage utility will be built. The building blocks of the StorageWorks grid include smart cell technology. These smart cells are standard, modular components that are designed to be dynamically changeable to deliver unique, real-time, information services.

In 2004, HP released two products that support the StorageWorks grid: The StorageWorks Reference Information Storage System (RISS) and the StorageWorks Scalable File Share. RISS not only addresses the business need of securely archiving application data, but also turns data into information by providing fast retrieval of archived content. Although the functionality of RISS is different from HP SFS, both products are built on a distinct set of smart cells.

HP Labs has prototyped other storage hardware and software based on smart cells. This HP research is expected to deliver technology and platforms shared across diverse smart cell-based products that will increase performance and lower costs while augmenting resiliency.


Major Components of HP SFS

Lustre is arguably the most scalable filesystem in existence today. It uses “object based storage” technology³ at the core of its scalable architecture. This architecture and implementation details are described in papers from CFS at <http://www.clusterfs.com>. Thus, only a summary of this architecture need be included in this HP SFS paper.

Several major components on a Lustre server include: the data server smart cells (Object Storage Servers), the metadata server (MDS) smart cells, the compute clients, an LDAP server, and one or more message-passing interconnects. These components work together to provide a single, coherent, scalable filesystem.

³ <http://www.snia.org/>

Figure 3. The major components of HP SFS



Data Smart Cells⁴

The data smart cells contain the main data storage. Files are striped across the data smart cells with each cell serving part of the file simultaneously for scalable-parallel bandwidth. The corresponding stripe width depends on the file size and other criteria. Files can be striped across one or more cells. Small files are typically stored on a single cell, whereas large files tend to be striped across many cells. The administrator and/or user control the width of the file stripe with settable parameters.

In resilient HP SFS configurations, which is the norm for HP SFS, data smart cells are deployed in active-active pairs, with their storage cross-connected. The result is no single point of hardware failure. Thus, any hardware component in a data smart cell or its storage can fail and the client application will maintain a transparent and uninterrupted path to its storage. As the number of components in the fileserver increases, the likelihood of hardware failure also increases. Therefore, resiliency is often an important consideration.

Metadata Smart Cells⁵

The metadata smart cells contain the information about the files, such as the file names, the directory names and structure, the logical location of the data, file status, security, mount points, and symbolic links. A client application contacts the MDS directly when creating or opening files, when searching directories, or for other metadata queries and transactions. For an actual read or write transaction, the client uses this metadata when interacting with the data smart cells.

The MDS servers are deployed as an active-passive pair. Two MDS smart cells are cross-connected to a shared disk array. Should the active MDS fail, its passive counterpart will transparently assume its responsibilities. HP SFS always employs the MDS smart cells as resilient pairs, even if the data smart cells are not resilient.

Lustre Clients

The Lustre client code on the compute servers provides a standard filesystem interface to Lustre. Each client installs a path to the Lustre filesystem and presents a POSIX compliant filesystem

⁴ Data smart cells in Lustre parlance are called "Object Storage Servers" (OSSes) running "Object Storage Target" (OST) code

⁵ Metadata smart cells are called Metadata Servers (MDSes) in Lustre parlance.

interface to the applications. Thus, the applications perform normal open, close, read, write, record locking and other POSIX I/O commands. Lustre maintains coherency among the buffers in these distributed clients so the I/O behaves logically as if one were using shared-memory buffers on an SMP. This simple and standard API eases the portability of parallel applications. MPI-I/O also works across this interface but is not necessarily needed, given each client has direct, coherent, and high bandwidth access to the shared filesystem.

LDAP Server

The network data resource and administrative services are delivered using an LDAP (Light-Weight Directory Access Protocol) server. LDAP is an open and commonly used standard for managing configuration and security information. The LDAP server is also used in the resilient fail-over mechanisms to determine the redundant routes among the clients and the servers.

Message-Passing Interconnect

Lustre is architected to be portable to any message-passing interconnect. It uses an intermediate protocol layer called Portals, derived from work at Sandia National Labs. Portals provide low-latency OS bypass and zero-copy interfaces on hardware that supports such features. It also runs on standard TCP/IP protocols over standard networks such as Ethernet. The latency and bandwidth tends to be slower than that of more efficient OS-bypass hardware, but may be sufficient for many Lustre applications.


The interconnect for the HP SFS server tends to be driven by the interconnect used in a given compute cluster. An existing LAN or interconnect can be used to connect the HP SFS smart cells to each other and to the compute clients. For example, if Quadrics were used in the compute cluster, it would also be used to connect to the HP SFS smart cells.

Options Available in HP SFS smart cells

HP SFS supports various hardware options ranging from low to high resiliency in its various hardware options. The initial offering in 2004 includes resilient configurations based on StorageWorks EVA disk arrays and a non-resilient option based on StorageWorks MSA disk arrays. Later HP SFS releases are planned to include continuously improving storage and server hardware to provide optimum choices for the right level of resiliency at the right price.

The data and metadata smart cell configurations are similar. They differ in that an active-passive MDS pair is cross-connected to one resilient set of storage. Only one MDS is active at a time, therefore only one dedicated disk array is needed. Because of the active-active configuration, the data smart cells need one or more dedicated disk arrays per cell. These disks are cross-cabled for fail-over.

Figure 4. Multiple resiliency techniques are employed in data and metadata cells.



The disk arrays are also configured with no single point of failure. This is usually accomplished with RAID-5 controllers that provide a logically redundant parity disk, but could be a combination of mirroring, RAID-5, and striping at several levels. This flexibility allows for integration of new and potentially different smart cell hardware over time.

As server and storage technology advances and HP smart cell designs are standardized for a wider set of functions, the smart cells will be implemented on even higher volume, higher-performance, and less expensive hardware designs.

HP Additional Value

HP contributes additional value over the existing open source code in the following areas:

Automated installation and management: HP SFS provides tools that automate the installation, discovery, and configuration of the Lustre technology. These tools are a significant enhancement, ensuring ease of administration by allowing installation and management from a single point.

File system creation: HP SFS provides a simple "configure file system" command which presents user with set of legal options and guides user through the set up process. All set up is managed from a single management point with the data automatically loaded into a database that is viewed from the single management point.

Filesystem management: HP SFS includes simple management commands that show the status of the filesystem on each server, from a single management point.

High availability: HP SFS provides carefully designed, selected, tested, and verified configurations ready to deploy.

Hardware/software compatibility, qualification and testing: The HP SFS hardware, software, storage, and interconnects are compatible, qualified and tested as a complete product.

Performance characterization and optimization: The standardized HP SFS configurations allow HP to optimize performance, study anomalies, and better tune the I/O.

Hardware and Software Support: HP backs the HP SFS product with worldwide support packages and HP QuickStart services.

The integrated HP SFS product includes the following categories of components:

- Data smart cell hardware (minimum 2)
 - ProLiant servers with attached StorageWorks disk arrays.
 - The data smart cells are replicated as needed to scale the bandwidth and capacity.
 - Metadata smart cell hardware (minimum 2)
 - ProLiant servers with attached StorageWorks disk arrays.
 - A pair of metadata smart cells are supplied: one active and one backup.
 - Disk storage systems.
 - HP StorageWorks EVA disk arrays for maximum resiliency.
 - HP StorageWorks MSA disk arrays for optimized price/performance.
 - One or more dedicated disk arrays per data smart cell.
 - One dedicated disk array per pair of metadata smart cells.
 - Racks, cables, and other hardware to house and connect the storage to the data and metadata smart cells.
 - The HP SFS software distribution.
 - Includes a base of open source Lustre code.
 - Includes HP-specific code for improved ease of installation, maintenance, use, and performance.
 - Includes HP-specific software for the HP-specific hardware.
 - A thorough set of HP SFS documentation for easy administration.
 - Warranty and separately priced installation, maintenance, and support services.
-

Available Server Platforms for HP SFS

HP SFS is tested and warranted to work on most HP ProLiant and HP Integrity Linux servers. HP SFS provides the Linux kernels for these selected platforms, including the needed Lustre client modifications. The Linux servers in the HP XC cluster products come pre-configured with these Lustre modifications, ready to integrate with HP SFS.

The Lustre client code is open source and is supported on many vendors' server platforms running Linux. HP can support these multi-vendor platforms through HP Consulting and Integration services. For product quality control and support purposes, the initial support for HP SFS clients is on approved HP hardware, with other arrangements by specific exception.

Expanding or Changing HP SFS after It Is Installed and Running

To add I/O bandwidth (GB/s) or capacity (terabytes) to an existing HP SFS server, one adds data smart cells. Each smart cell adds bandwidth and capacity to the system.

An existing filesystem can be reconfigured to include the new smart cells. Existing files in the filesystem remain striped across the original set of smart cells. When new files are created, they can be striped across both the older and newer smart cells.

Accessing HP SFS Stored Data from Servers That Are Not Lustre-Enabled

The HP SFS storage is exported through two protocols simultaneously: Lustre and NFS. The Lustre service delivers significantly higher bandwidth and is inherently scalable. However, not all clients in a data center are likely to be Lustre-enabled and so NFS service is provided as an alternate access method.

Support Services Offered for HP SFS

HP SFS comes complete with warranty and worldwide support, including a broad range of service agreements up to 24x7 on-site coverage. HP Care Pack Services are available in addition

to the standard product warranty of HP SFS. A choice of support levels ensures all business requirements are met -- from basic to mission-critical. Upgraded service levels extend and expand standard product warranty with easy-to-purchase and easy-to-use support packages that help organizations make the most of every hardware and software investment.

The Future of HP SFS

HP's continuing enhancements to HP SFS fall into several categories:

- HP improvements to Lustre (contributions to the open source Lustre base)
- New releases of the open source Lustre base, with will be incorporated into HP SFS
- HP enhancements to HP SFS-specific code and tools above the Lustre base
- HP SFS support on ever improving smart cell hardware
- Additional options for lower-cost smart cells, delivered with additional resiliency and price options
- Multiple active metadata servers (in a major release of Lustre slated for 2005).
- Additional tuning and scalability for NFS service
- Security enhancements
- Ongoing tuning, performance, and scalability enhancements.

Conclusions

HP SFS solves the distributed file system I/O challenge while simplifying the use and administration of Linux clusters to deliver faster processing and a higher return on investment. HP SFS is based on Lustre technology, a high-performance, highly scalable file share engineered to provide extremely fast, scalable, and reliable I/O for Linux clusters. HP implemented numerous value-added capabilities to the open source Lustre distribution to improve ease of use, deployment, and support and to increase functionality and performance.

For More Information:

www.hp.com/go/hptc

HP High-performance technical computing Web site

All about HP scalable-technical solutions, including Linux clusters and HP SFS

www.hp.com/storage

HP Storage Web site

All about HP StorageWorks solutions

www.lustre.org

The main source for Lustre news, overviews, and technical information

Lustre is a trademark of CFS.

www.clusterfs.com

Home page of Cluster File Systems, Inc., the principle developer of the Lustre architecture and code base.

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Itanium is a trademark or registered trademark of Intel Corporation in the U.S. and other countries and is used under license.

5982-9347EN, 10/2004

