



Improving Scalability, Manageability, and Performance for Network-Attached Storage (NAS) Solutions



1 Introduction

No one disputes the theoretical advantages of network-attached storage (NAS) compared to direct-attached storage. The inefficiencies involved with multiple, separate data storage platforms can be greatly improved with an alternative NAS solution. Scalability, availability, and performance can also be enhanced with the currently available NAS platforms. However, many NAS solutions introduce complexities and hidden costs. The expertise required to effectively deploy and manage some NAS products can require adding IT specialists to the team, and can introduce administration and monitoring tasks that raise the total cost of ownership. For these reasons, NAS has been deemed too expensive by some organizations that could otherwise benefit from the inherent advantages such as simplified data sharing and increased data access performance.

With many years of experience as an innovator of storage solutions, SGI® recently introduced a new generation of NAS solutions that overcome these complexities and cost barriers. The SGI InfiniteStorage NEXIS NAS Family and SGI InfiniteStorage Appliance Manager software make scalable, high-performance NAS solutions available to a broad range of businesses and organizations.

This paper introduces the SGI InfiniteStorage NEXIS NAS Family and SGI InfiniteStorage Appliance Manager, and overviews the SGI innovations that optimize the performance and cost efficiencies available within an NAS environment.

2.0 The SGI InfiniteStorage NEXIS NAS Family and SGI InfiniteStorage Appliance Manager

2.1 Current Family Offering

The SGI InfiniteStorage NEXIS NAS Family (see Table 1) provides scaleable investment protection, reduced manageability cost, and high performance for vertical market and business applications.

2.2 The NEXIS Difference

Two primary areas of innovation set the SGI NAS platforms apart from the competition:

- **Optimized file system.** The SGI XFS file system enables efficient scalability of storage capacity, file directory size, and performance, and has been tailored for NAS and other performance-critical environments.
- **Ease of management.** The new SGI InfiniteStorage Appliance Manager provides a single, unified storage management interface featuring automated storage asset discovery and configuration guidance for fast and easy initial deployment and future storage capacity expansion. Granular resolution guidance and storage utilization tools help reduce manageability costs and maximize storage efficiencies.

The SGI innovations give users the ability to non-disruptively scale capacity and directories with minimal performance impact. High file directory scalability reduces management complexity by enabling a large single directory for archiving or for expanding multiple departmental directories while improving performance. This scalability, combined with the simplified management enabled by the SGI InfiniteStorage Appliance Manager, add up to lower total cost of ownership for the overall NAS solution.

3.0 XFS: A Scalable File System for Fast-Growing Storage

At the heart of each SGI InfiniteStorage NEXIS solution is the XFS® file system. XFS was originally developed by SGI to provide a 64-bit journaled file system that can handle large files quickly and reliably. Fully embraced by the open source community today, XFS leverages SGI's intimate knowledge of mixed work load serving environments. As a next-generation product, XFS is a new fully integrated product, rather than a modification of an existing product. Thus, important file system features are seamlessly integrated into XFS instead of having been awkwardly added to an older, existing file system.

High storage capacity and file count scalability provide the foundation for seamless file directory growth, reducing the need for multiple directory configurations and associated management complexities. The fast and efficient XFS file system provides highly scalable performance for a wide range of applications ranging from small workgroups to demanding enterprise and vertical market requirements.

3.1 Feature Summary

XFS provides a full 64-bit file system capable of scaling easily to handle extremely large files and file systems that can grow to 1 terabyte (up to 9 million terabytes in successor releases). XFS's major features include:

- Full 64-bit file capabilities to support large files, file systems, and numbers of files, as well as sparse files (files with holes)
- A rapid and reliable recovery of file system structure using journalling technology:
 - Integrated, full-function volume manager called XLV
 - Extremely high I/O performance that scales well on multiprocessing systems
 - Guaranteed rate I/O for multimedia and data acquisition uses
 - Backup of active file systems (while still in use), significantly reducing administrative overhead and disruption to users
 - XFS is compatible with existing applications, EFS file systems, and NFS

3.2 A New File System Design

Unlike alternative file systems, XFS is not simply a rewrite or port



InfiniteStorage NEXIS Model	Description	Target Applications	Key Features
SGI InfiniteStorage NEXIS 7000 H/A	SGI's best NAS performance and availability for mixed application workloads	Database transactions, data mining, media-rich apps, and online social networks	Industry-leading NFS performance. Fast and scalable file directories. Performance optimization and management tools. Fast and automated data path failover.
SGI InfiniteStorage NEXIS 7000	High performance without redundant data path failover for dedicated workloads	Media production, oil and gas, science and engineering modeling, business intelligence/OLAP	Industry-leading NFS performance. Fast and scalable file directories. Performance optimization and management tools.
SGI Infinite Storage NEXIS 5000	High storage capacity, moderate performance for large-file sequential workloads and online archives	Disk backups, corporate archiving, e-discovery, online research libraries, audio/video & graphic object repositories, compliance archives	Large root directory storage capacity matches available storage. Storage utilization management tools. Non-disruptive storage capacity expansion.
SGI InfiniteStorage NEXIS 2000 (SAS)	Entry-level performance and storage capacity	Media, education, and scientific workgroup storage consolidation and file sharing	Price/performance. Fast & scalable file directories. Performance optimization management tools.
SGI InfiniteStorage NEXIS 2000 (SATA)	Entry-level storage capacity	Small online archives, Disk-to-tape backups, moderate performance for large-file sequential workloads	Large root directory capacity matches available storage. Storage utilization tools. Non-disruptive storage capacity expansion.
SGI InfiniteStorage NEXIS 500	Integrated NAS server and storage in a single 2U chassis	Small workgroup file sharing, dedicated scratch file storage, and workstation storage expansion	NAS server and storage in a single 2U chassis save space. Easy to deploy, configure and manage.
SGI InfiniteStorage NEXIS 1000 Gateway	NAS connectivity to existing SGI CXFS & XFS compliant storage	Similar applications to the NEXIS 2000	Investment protection by utilizing existing SGI storage

Table 1. The SGI InfiniteStorage NEXIS NAS Family

of existing technology such as the UNIX System V or BSD file system. XFS was "built from scratch" and SGI was therefore able to integrate key features such as journaling for high reliability and redesign important areas such as the allocation algorithms for increased performance with large file systems. As a result, XFS is the first file system built for today's large, demanding, data-intensive environments.

Most other file systems on the market evolved in an age of smaller file systems, lower processing power, and limited storage capacity. As a result, many file systems are based on architectures that emphasize conserving storage space at the expense of performance. The revolutionary growth of CPU and storage technology has fueled a dramatic increase in the size and complexity of data, outstripping the capabilities of many

file systems. SGI designed XFS to meet this increasing need to manage large complex file systems with high performance and reliability. Thus, completely different design targets were established for XFS:

- Scalable features and performance from small to truly huge data (petabytes)
- Huge numbers of files (millions)
- Exceptional performance
- Designed with log/database (journal) technology as a fundamental part not just an extension to an existing file system
- Mission-critical reliability

The cutting-edge XFS technology opens up new opportunities in the marketplace by enabling users to manage large amounts of

data with unsurpassed speed. This is particularly advantageous in NAS environments. Additionally, new programmatic features and interfaces such as guaranteed rate I/O provide unique advantages to users. Finally, the growing necessity for business-critical reliability is driving corporations toward robust journalled file systems such as XFS.

3.3 File System Efficiencies

XFS uses sophisticated file system management techniques such as extents and B-Tree indices to efficiently support:

- **Very large files (64-bit size).** There is little or no performance penalty to access blocks in different areas of a large file. XFS creates large sizeable extents to locate file data close together for faster reading and writing of data.
- **Very small files.** Most symbolic links and directory files are small files. XFS allows these files to be stored in inodes for increased performance. XFS also uses delayed writes to wait to gather the entire small file in the buffer cache before writing to disk. This reduces the number of writes to disk and minimizes the number of extents.
- **Sparse files.** Sparse files are files that contain arbitrary “holes,” areas of the file that are untouched and that read back as zeroes. XFS supports holes in a manner that avoids wasted space.
- **Mapped files.** Memory mapping of files is supported by XFS. It allows a program to attach a file such that the file appears to be part of the program, and lets XFS manage the disk I/O operations.
- **Large directories.** Large directories are indexed in directory files to expedite searches, insertions, and deletions. Operations on directories containing millions of files are almost as fast as on directories containing only hundreds of files.

3.4 High Performance

XFS is a very high-performance file system with support for contiguous data (reduced I/Os) and both direct (non-buffered) and asynchronous I/O file system. SGI has also made significant enhancements to the integrated NFS functionality within XFS. These advancements reduce latency, and tailor the otherwise “chatty” NFS protocol for performance critical environments such as NAS. As a result, these highly shared environments benefit from increased file system throughput.

XFS dump/restore provides high performance by:

- Using multi-threaded processing that streams the drives so that they never “starve” for data and evenly distributes the dump across drives
- Employing very large record sizes (typically 2MB) to reduce I/O operations

4.0 SGI InfiniteStorage Appliance Manager: Minimizing Complexity in NAS Environments

The new SGI InfiniteStorage Appliance Manager software greatly simplifies implementation and optimization of NAS solutions.

The intuitive, browser-based GUI facilitates step-by-step set up, and helps automate the tuning of storage systems for the best possible IOPS or throughput performance.

The single, unified storage management interface of the SGI InfiniteStorage Appliance Manager helps to automate storage asset discovery and provides configuration guidance for fast and easy initial deployment and future storage capacity expansion. Granular resolution guidance and storage utilization tools help reduce manageability costs and maximize storage efficiencies.

4.1 Simplifying Administration and Monitoring

With the SGI InfiniteStorage Appliance Manager, users can:

- Perform initial system configuration using the Setup Wizard.
- Configure the system components.
- Perform general system administration tasks.
- Monitor current and historical views of a fileserver’s state and performance (including disk space, CPU usage, and network throughput).
- View connected clients and determine how each of these contribute to the current workload.
- Detect and investigate problems.



Figure 1. SGI InfiniteStorage Appliance Manager interface

The “summary” menu option displays a graphic summary of system utilization, including file system and node status, number of alerts, CPU usage, disk space, disk throughput, network throughput, current clients, and uptime.

A “monitoring” menu option gives users the ability to quickly review several categories (see Figure 2):

- **Alerts** displays messages from the system logs. These provide informative messages, notifications of unusual events, and error conditions.
- **Resources** shows finite capacities of the file servers. Current and historical views are provided for the state and the performance of a storage server. This includes CPU usage, disk and network throughput, disk space and quotas, and many other metrics. It also allows viewing connected clients and determining how each of these contribute to the current workload.
- **Services** brings up a list of services provided by the fileserver. By clicking on a service (such as NFS and CIFS,), users can view its status. This screen also displays the current version of installed software.
- **Clients** displays various I/O criteria for the fileserver’s clients.

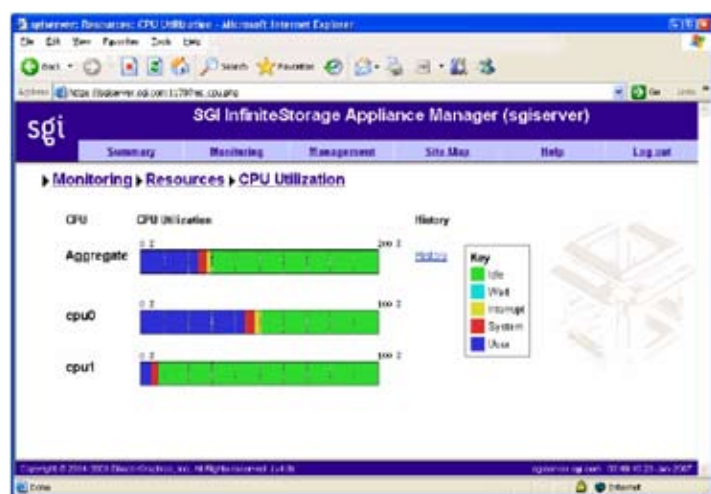


Figure 2. SGI InfiniteStorage Appliance Manager interface for performance monitoring

The “management” menu option makes it fast and easy to perform a variety of administrative tasks including:

- **Save/Restore Configuration** for backing up and restoring the system configuration files created with SGI InfiniteStorage Appliance Manager
- **Configure resources**, such as network interfaces
- **Configure services**
- **Global configuration** and various general system administration task

4.2 Securing NAS Without Impeding Information Sharing

SGI InfiniteStorage Appliance Manager also simplifies the process of configuring and managing security within an NAS. The user and group configuration menus walk administrators through the process of configuring a name service client, defining local users and groups, and setting user and group quotas (or resource utilization limits).

The “Name Service Client” screen makes it easy to specify a name service (or directory service) for the system. A name service is the application that manages the information associated with the network users. For example, it maps user names with user IDs and group names with group IDs. It allows for centralized administration of these management tasks.

Administrators can specify the use of a site-wide network information service (NIS), lightweight directory access protocol (LDAP), Active Directory services, or local files (in the absence of a site-wide protocol; names and IDs are kept locally on server). NIS is a network lookup service that provides a centralized database of information about the network to systems participating in the service. The NIS database is fully replicated on selected systems and can be queried by participating systems on an as-needed basis.

5.0 Summary

The combination of SGI file system innovations and new SGI InfiniteStorage Appliance Manager browser-based management and monitoring solution enable the SGI InfiniteStorage NEXIS NAS Family to provide scalable investment protection, reduced manageability cost, and high performance for vertical market and business applications. By overcoming the complexities and inhibiting cost of ownership of traditional NAS platforms, SGI technology brings the latest storage solutions to a broader range of businesses and markets.

For more information, visit sgi.com/storage