White Paper

SGI Network-Attached Storage: Building HPC Functionality Into Today's Appliances

Table of Contents

Executive Summary

Network-attached storage (NAS) has come front and center as IT departments are pressured to consolidate, centralize, and in any other way possible streamline operations and reduce operating costs (OPEX). Duplicated efforts can be avoided and capacity more effectively utilized when large pools of storage are managed centrally to overcome the scalability limitation of NAS islands With today's high-speed networks, storage can be made very accessible to users on the network and users raise no objections to the network-attached model as long as the file system shields topologies and performance remains comparable to local storage.

Silicon Graphics, with a long history in high-performance server and data management architectures, provides NAS solutions based on a trickle-down migration of technology. Originating in the most demanding high-performance computing environments, SGI NAS solutions have been characterized by optimized performance, high-density, and maximum capacity for science, engineering, oil exploration, and government research. The latest architectural enhancements also bring the SGI optimized NAS architecture within reach of a broad range of enterprise environments such as on-line analytical processing and data mining.

This paper overviews the SGI NAS strategy, the platform alternatives, and the effects of the latest innovations on users and networks.

1.0 SGI and NAS: Then and Now

The SGI history in data management focused on meeting the requirements for the most technical users and very large-scale data sets. As a result, the SGI NAS product line evolved to include leading-edge solutions focused on performance and capacity. Today, the traditional markets for the company require a much more diverse range of solutions and more closely mirror the demands of many enterprise environments. The driving forces for data management, from both scientific/CAD and commercial customers, include demands for:

- Ease of use – Users want solutions that are easy and fast to deploy, and can be supported by any IT professional (no requirements for NAS specialists).
- Lower OPEX – Having the right tools to efficiently handle data management is essential, and users also want solutions that are cost-effective in terms of facilities requirements (floor space, power, and cooling).
- Choice of capacity or performance solutions – While some projects require the best possible performance, others place higher priorities on capacity and cost. Enterprises and project teams want solutions that give them different price-performance points.
- Unified management – Choices in the platforms must be transparent to the user and IT. Unified software and system management are required to shield users from platform complexities and differences.
- Standards compliance – Global operations require that NAS solutions comply with region-specific regulations. In particular, RoHS compliance is required throughout Europe.

2.0 Multi-Workflow Storage Requirements

In the past, a few key R&D leaders in one or two fields would drive requirements. Today, requirements come from all scientific and engineering disciplines as well as business functions spanning the accounting department to the factory floor, product development, marketing, and traditional research and development. A single NAS platform cannot meet the scale and workflow requirements for every technical and business application. Rather than force-fit one NAS storage solution into all environments, Silicon Graphics has developed an NAS product line that applies best-of-breed technologies suited to the two diverging categories of computing and the related storage requirements:

- Performance workflows: These workflows and applications are data-centric and require extremely fast access to data sets. Often characterized by highly non-deterministic resource demands and high levels of user interactivity, performance workflows are typical in very rapid development cycles within research and development organizations, and in enterprise applications such as streaming media to thousands of users or high-end transaction processing.
- Capacity workflows: Typical within production environments, these workflows are often driven by operational costs rather than time to insight. Capacity workflows are typically highly automated with little if any user interactivity. Adherence to standardization is important for operational efficiency.

3.0 Storage Appliances

To provide storage solutions for the full range of workflows, the SGI NAS storage family includes two products. The high-performance SGI InfiniteStorage NAS 4550 and the high-capacity SGI InfiniteStorage NAS 4050 can be further tailored to different application environments (see Table 1).

While diverging in terms of the technologies employed to address specific requirements, the product line is unified with:

- Plug-and-play deployment: These are true storage appliances, preconfigured at the factory to speed installation and simplify overall administration. Complete set up and installation of any SGI NAS appliance can be done in less than 20 minutes.
- Single, tunable file system: With the ability to support a single file-system across all the available capacity, the SGI NAS solutions simplify the data management for very large data sets. Scalability and throughput are also maximized with this architecture.
- Intuitive management software: All of the SGI NAS appliances integrate the newest NAS manager 3.2 to provide an exceptional user experience including graphical representation of system status, menu -driven configuration, and comprehensive optimization tools.
- Integrated data service tools: Network aggregation, storage tuning, iSCSI target, snapshot and replication software module are all integrated into NAS manager 3.2 for ease of management.
- Advanced system and network diagnostics: Troubleshooting can be done much more quickly with the integrated suite of diagnostics.

Table 1. SGI NAS Appliances for Multi-Workflow Applications and Environments

SGI NAS Appliances:	Configurations:	
NAS 4550 for HPC IS-4500 disk subsystem High-availability controllers Fiber Channel Drives 16 drives per enclosure Max: 14 enclosures per system	**Streaming:** High-performance technology for media production, petrochemical research and exploration, science, engineering, and data mining 2 GB controller cache 10K RPM 300 GB Fiber Channel drives for large capacity applications 4+ 1 RAID 5 group for optimal capacity and performance Base system: 32 x 300 GB drives Add-on enclosures with 300 GB or 146 GB Fiber Channel drives or 500GB Serial ATA (SATA) drives	**Ultra:** Throughput optimized for data base transaction, log file, and streaming loads, and multiuser workloads (500+ users) 16 GB controller cache 15K RPM 146 GB Fiber Channel drives increases spindles for high IOPS workloads 4+1 RAID5 group for optimal capacity and performance Base system: 32 x 146 GB drives Add-on enclosures with 146 GB or 73 GB FC drives or 500GB SATA drives.
NAS 4050 for High-Capacity 1 to 4 redundant controller pairs 2 GB controller cache SATA drives: - 500 GB 7,500 RPM drives - High density 3u enclosure with16 drives	**SATA Server:** Affordable SATA technology for secondary, low-cost storage for office file serving, imaging, online backups, secondary storage, and more Single subsystem Base system: 16 drives (mounted in the controller) Up to 96 add-on drives (total of 112 drives in 6 enclosures and controller)	**Maximum Capacity:** Suitable for corporate archiving, e-discovery, data/server consolidation initiatives, online research data sets, and multimedia storage 2 to 4 subsystems Base system: 2 controllers and 160 drives (8 enclosures) Up to 284 add-on drives (448 drives total in four controllers and 24 disk enclosures)

3.1 Capability: SGI 4550

For high-capability, performance-intensive applications and workflows, the SGI NAS family includes the SGI InfiniteStorage NAS 4550. Fiber Channel enables optimal throughput for HPC workgroups performing scientific research, or commercial teams in aerospace, energy, and other performance-intensive fields. This platform is uniquely configured to deliver industry-leading performance:

- 67 TB raw disk capacity with Fiber Channel drive
- 48 GB server memory, plus 2 or 16 GB disk cache, to deliver the performance required for high-performance applications and systems that support many users
- 12 Gigabit Ethernet ports, standard
- Four 4 GB Fiber Channel ports

3.2 Capacity: SGI 4050

The SGI InfiniteStorage NAS 4050 provides the highest capacity, with the ability to store up to 224 TB of usable, online data. This platform employs affordable serial ATA (SATA) drives and supports a single, scalable file system, resulting in the lowest cost per TB for data management. This storage platform suits research teams with extensive data sets, as well as enterprises that want to start small and be able to expand as the company or organization grows.

3.3 High Availability

All SGI NAS storage appliances have been designed for high availability. Features that ensure non-disrupted data access include:

- Mirrored root disk on the NAS server
- Dual array controllers with failover capabilities
- Mirrored write cache on the array controller
- RAID 1 and 5 protection for disk drives
- N+1 power supply and fans in each controller and enclosure

4.0 Software

All SGI NAS solutions are factory integrated and tested, with the software environment based on SGI and third-party innovations.

SGI NAS Manager software shortens the installation process with:

- Setup wizard via HTTP connection for fast initial system setup
- Quick network address and file system choices
- Built-in management and diagnostic functionality
- Easily accessible over-rides and details

The SGI NAS Manager platform also enhances NFS and SGI XFS file system performance and improves scalability within NAS environments. Additional enhancements have been integrated to complement the new appliances. These include the abilities to:

- Export iSCSI targets (using iSCSI target from SLES10), with full integration in the NAS Manager graphical user interface (GUI)
- Add storage and grow existing file systems in full tray increments
- Support multiple raid heads
- Monitor performance in real time using NAS Manager (via updated PCP and nasavg PMDA)
- Value-add historical trending function to help system administrators take proactive actions
- Use XVM for storage virtualization as well as volume snapshot for Linux, with full GUI integration
- Comply with RoHS and export requirements
- Support new Engenio 4Gb FC storage
- Support Samba/CIFS (with server performance stats) and other Samba features
- Perform more functions using the generic client-side GUI (field checking, progress indicators, etc.)

Bakbone NetVault Replicator software has been fully integrated with SGI NAS Manager, and can support replications for backing up SGI NAS solutions as well as other storage platforms. This utility allows:

- Remote copy of data for disaster protection, optimized for fast backups (only changed data is copied)
- Real-time, byte granularity for optimized control
- Flexible configurations (one-to-one and many-to-one for centralization)
- No distance limitations

XVM Snapshot software is also included for:

- Point-in-time copies of data
- Facilities backup processing (snap first and backup the snapshot)

5.0 Next Steps

The newest SGI NAS solutions can be configured as monolithic servers. In the next phase of the architecture evolution, the SGI solutions will be extended with additional high-availability features, and support for clusters.

Conclusion

The latest SGI NAS solutions demonstrate Silicon Graphic's continued leadership for HPC storage solutions, and also provide affordable alternatives to enterprises that are unsatisfied with the performance, capacity, or manageability of their existing storage solutions. By utilizing a combination of best-of-breed technologies, the SGI NAS solutions allow users to choose the best platform for each environment. The economies of commonality across the SGI platforms and the simplicity of the new fully integrated appliances provide lowered OPEX for both existing and new users. Compared to the competitive offerings, the SGI NAS solutions excel in all categories: performance, capacity, cost per TB, price-performance, and ease of use. The range of pre-configured platforms and expansion options allow project teams and businesses to start small today, and easily scale to maximum capacity over the years of deployment.

For more information about the newest SGI NAS appliances, please visit *www.sgi.com/storage*.

Corporate Office
1200 Crittenden Lane
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3971 [10.24.2006]

J15199