

Sun StorEdge™ 5310 NAS Appliance

Just the Facts

Table of Contents

Sun StorEdge 5310 NAS

Positioning.....	5
Introduction.....	5
Product Family Placement.....	6
Product Availability.....	7
Product Shots of the StorEdge 5310 NAS Appliance.....	8
Key Features.....	9
Key Messages.....	11
Value.....	11
Innovation.....	12
Choice.....	12
Target Markets	13
Sun Software Support.....	15
Third Party Application Support.....	16
Relevant Standards Supported	16
Features and Benefits.....	17
Enabling Technology.....	23
Technology Overview.....	23
Flash Disk Module for Operating System.....	23
Storage	23
2U Chassis.....	24
Motherboard Board Information	24
StorEdge 5300 RAID Expansion Unit & 5300 Expansion Unit.....	25
Software/Operating System.....	26
64-bit Journaling Filesystem.....	26
File Sharing for UNIX (NFS) and Windows (CIFS).....	27
Active Directory Services.....	29
Sun StorEdge File Replicator.....	30
Sun StorEdge File Cluster.....	33
Sun StorEdge File Checkpoint.....	37
High Availability Bonds: (NIC Failover / Port Aggregation):.....	41
Operation.....	42
Recovery.....	42
System Architecture.....	43
Overview.....	43
Hardware Components.....	43
StorEdge 5300 Expansion Units.....	44
Sun StorEdge 5300 Expansion Unit.....	48
Reliability, Availability and Serviceability (RAS).....	50
Reliability.....	50
Availability.....	50
Serviceability.....	50
Tape Backup and Restore.....	52
Server Management.....	54
IPMI.....	54
System Monitoring	54
SMTP.....	56
SNMP	57
Telnet	58
LCD Keypad.....	59
Syslog/Remote Syslog.....	59
Sun StorEdge 5310 NAS Specifications.....	60
Processors.....	60
Memory	60
Standard/Integrated Interfaces.....	60

Mass Storage.....	60
Power Supply.....	60
Environment.....	60
Regulations.....	61
Dimensions and Weight.....	61
Software/Operating System.....	61
StorEdge 5300 RAID Expansion Unit & 5300 Expansion Unit F Specifications.....	63
Requirements and Configuration.....	64
Cabling Diagrams.....	64
Client Operating System Support.....	67
Software Support Matrix.....	67
Peripherals Support Matrix.....	67
Expansion.....	68
Mixing Fibre Channel and S-ATA.....	68
Calculating Usable Capacity.....	68
Ordering Information.....	69
Warranty, Service and Support.....	75
Warranty.....	75
Sun Spectrum Instant Upgrades.....	75
Enterprise Installation Service.....	75
Knowledge Services.....	75
Professional Services.....	76
Glossary.....	77
Materials Abstract.....	82
Competitive Information.....	83
Q & A.....	85

Revision: January 2005

- Corrected SCSI interface description on page 10
- Updated product photo's
- StorEdge File Cluster content
- Tape backup now shows direct attach and 3-way support

February 2005

- Various grammatical and minor technical corrections

March 2005

- Added S-ATA content
- Updated Q&A

April 2005

- Clarified SE5300 RAID EU support for SE5310C in ordering section

May 2005

- Added 3-way and direct attach backup support with NDMP
- Added usable capacity guide for FC & S-ATA
- S-ATA supported with StorEdge 5310C NAS Appliance

June 2005:

- Added support for Legato NetWorker/EBS v7.2
- File system limits
- RAiD group limits
- Replaced stand-alone diskless RAID EU with diskless RAID EU/SATA DEU combo in ordering section

Sun StorEdge 5310 NAS

Positioning



Figure 1. Sun StorEdge 5310 NAS Appliance and StorEdge 5300 RAID EU

Introduction

The Sun StorEdge 5310 NAS is Sun's latest offering in NAS (Network Attached Storage). It is a mid-range storage solution ideally suited for departmental, workgroups, distributed enterprises, and small- to medium-sized companies with a need to consolidate storage, improve availability, share information, and reduce management costs. Purpose-built for simple operation, the StorEdge 5310 NAS combines high performance with low maintenance requirements to create a versatile, NAS solution for mixed NFS/CIFS environments.

The SE5310 NAS is based on a dedicated, storage-optimized OS and Xeon hardware platform. It typically requires minutes to install and less than three minutes to boot, and delivers high availability with minimal maintenance.

Featured is a 64-bit journaling file system, checkpoints, RAID 5, battery-backed cache, hot-swappable drives and power supplies, global spares, and support for a UPS. The file system aids in ensuring data integrity across unforeseen events such as power outages. Sun StorEdge File Checkpoint creates "checkpoints" (also known as snapshots) that provide static images of the file system, enabling rapid recovery of mistakenly deleted files.

The SE5310's space-saving design conserves precious IT real estate, requiring only 2U of rack space. With the addition of up to 4 Sun StorEdge 5300 RAID Expansion Units and 28 StorEdge 5300 Expansion Units (up to seven StorEdge 5300 EU's may be attached to one StorEdge 5300 RAID EU), the total

Sun StorEdge 5310 NAS

Just the Facts



capacity can be expanded to over 65TB (raw) utilizing 146GB Fibre Channel drives. 400GB S-ATA drives may be used for a total raw capacity of 179TB's.

All Sun StorEdge 5310 NAS models include:

- Operating System with a user license for an unlimited number of users
- CIFS and NFS file protocol support
- Sun StorEdge File Checkpoint (snapshot)
- Dual redundant power supplies
- Embedded dual 10/100/1000 Gigabit Ethernet NIC ports
- Configuration with 2-port optical (MMF) NIC support
- 1 or 2 Two-port 2Gb/s PCI Fibre Channel Host Bus Adapters

Product Family Placement

The Sun StorEdge 5310 NAS offers quick installation, simplicity of configuration and ease of management that scales in capacity from 2TB to over 65TB (FC), and from 5.6TB to 179TB (S-ATA). and is the latest member of the Sun NAS family. The StorEdge 5310 NAS provides customers with high-performance file serving and file storage.

The following table illustrates where NAS might best fit and environments better suited to other solutions.

<i>Criteria</i>	<i>DAS</i>	<i>SAN</i>	<i>NAS</i>
Cost	Inexpensive, but difficult to manage. Fundamental cause for the IT "Islands of Data" problem.	Very Expensive and Labor Intensive and difficult to manage	Inexpensive and Easy to Manage
Connectivity	Minimal, only a single connection per Disk	Good, multiple connections per Storage Device	Excellent, can be in the hundreds or thousands
Performance	Moderate to Fast, depending on Technology	Extremely Fast	Fast and in many cases, higher performance than DAS.
Price/Performance	Marginal due to high administrative overhead and lack of manageability	Moderate due to expensive hardware and high labor costs	Excellent due to inexpensive hardware, easy to manage and minimal labor costs
Storage Management	Marginal, accomplished "server by server" contributing to the "Islands of Data" IT problem.	Storage Management tools are available but some lack maturity. Data availability problems may plague LUN management functionality such as expanding LUNs "on the fly" as contrasted with NAS Logical Volume Management.	Built-In to Filer Platforms. Volume Management, DTQ/User Quotas & Snapshotting provide abstract layer storage management not available in SAN or DAS environments.
Storage Sharing	None.	Storage is consolidated but not shared. But, fundamentally, same as DAS. Centrally managed "Islands of Data" Platform.	Hardware is shared and Data is shared, fundamentally an optimum functional solution.
Data Sharing	None, this is what causes "Islands of Data" problem.	None, this is what causes "Islands of Data" problem.	All data sharable. Resolves the "Islands of Data" dilemma.
IT Management	Difficult, causes IT Management problems such as extended outages and undue risk to production data	Moderately Difficult, multiple connectivity options as well LUN management tooling enables Best Practices but can be cumbersome due to proximity to Server Support software systems and data.	Excellent, NAS Logical Volume Manager, Mirroring and Snapshotting separates data from Server Infrastructure providing built-in protection to production data during IT Management functions.

<i>Criteria</i>	<i>DAS</i>	<i>SAN</i>	<i>NAS</i>
Backup/Restore	Difficult due to loss of Data Availability when backing up mission critical highly volatile filesystem(s). This is exacerbated by the "Islands of Data" architecture	Better, data is separated from production support systems, proximity to B/R tooling is enabled by FC SAN connectivity.	Excellent due to filesystem features that enables B/R Best Practices. Features include Snapshottng as well as ability to participate in SAN Connectivity to B/R tooling. Snapshots enable uninterrupted access to Production Data at all times.
OLTP (Low Transaction Rate)	On board DAS, could work well, but IT management issues marginalize the implementation.	Excellent. Server and SAN I/O Throughput exceeds the I/O and Network requirement.	Excellent. Server and Network I/O throughput exceed the requirement for performance. Snapshottng enables application processes that enhance processing of data.
OLTP (Medium Transaction Rate)	Marginal. Other disadvantages as well as peak workloads may exceed ability to preserve performance	Excellent. Server and SAN I/O Throughput exceeds the I/O and Network requirement.	Very good. Server workload and Network I/O will meet this requirement. Care should be taken that txn levels don't exceed the network's ability to transport data to/from the Filer's NAS Filesystem.
OLTP (High Transaction Rate)	Marginal, depending on disk based technology. FC Implementation would work well, but IT Management issues marginalize data availability. High Txn environments exceed typical SCSI performance and are usually not recommended.	Excellent. Server and SAN I/O Throughput meets or exceeds the I/O and Network requirement.	Marginal. Network throughput and overhead impacts overall application performance.
CAD/CAM	Good, single user(s) mode only. Can expect users to move/copy drawings/designs to local workstations due to inadequate server performance.	Good, but wastes the storage in single user mode. Lack of data sharing capability will inhibit CAD/CAM design collaboration processes.	Excellent. Shared Storage enables multiple users to reference the data across platforms and OS architectures due to File Level Locking. Home Directory support functionality enables multiple users/projects on the NAS Filesystem. Snapshottng and other data protection mechanisms enable change control of the CAD/CAM design(s).

Product Availability

• Sun Product Introduction (Presto)	December 7, 2004
• WEBDESK Orderability	January 11, 2005
• Public Announcement	November 15, 2004
• Revenue Release (RR)	December 21, 2004
• General Availability (GA)	January 14, 2005

Product Shots of the StorEdge 5310 NAS Appliance



Figure 2. Front View



Figure 3. Rear View



Figure 4. StorEdge 5310 NAS Appliance and StorEdge 5300 RAID EU

Key Features

The initial release of the Sun StorEdge 5310 NAS includes the following features:

Hardware

Feature	Specifications
Motherboard	Intel Xeon class Server motherboard with 533MHz Front Side Bus and dual PCI-X 64bit/100 Bus
Processor	Intel Xeon 3.06GHz with 512KB Level 2 cache
Maximum number of CPUs	One
Default Memory	4GB, DDR RAM
Expansion Slots	6 PCI-X
Network Ports	Dual embedded 10/100/1000 Gigabit Ethernet ports
Boot Disk	Reliable 256MB Flash media
External Hardware RAID	External Fibre Channel hardware RAID-5, with redundant controllers and mirrored, battery-backed cache
RAID Level Support	RAID 5
LCD Keypad Interface	LCD interface on the front of the unit for easy setup and visual status notification
Bezel	Supplied with standard configuration
Cooling	Redundant fans
Power Supply	Dual 500W redundant hot-swappable power supplies
Rack Height	2U (3.5").
Dimensions	H 3.5" x W 19" x D 25.51"
Expansion Module	<ul style="list-style-type: none"> 1 – 2 (4 via “Get to Yes” with a 60 day lead-time) StorEdge 5300 RAID Expansion Units, dual RAID controllers, 14 (2.04TB) 146GB FC drives. StorEdge 5300 RAID Expansion Units are diskless for S-ATA configurations. Up to 28 (7 per StorEdge 5300 RAID EU) StorEdge 5300 Expansion Unit F, providing an additional raw capacity of 2.04TB with 14/146GB FC drives per tray Up to 32 (8 per diskless StorEdge 5300 RAID EU) StorEdge Expansion Unit S, providing additional raw capacity of 5.6TB with 14/400GB S-ATA drives per tray.
Warranty	2 years, 1 st year on site, 2 nd year 15 day return to depot

Software

Feature	Specifications
Operating System	StorEdge NAS OS, a highly efficient storage-optimized Operating System
File System	64-bit Journaling File system, with Dynamic Volume Expansion
File Access Protocols	CIFS/SMB, NetBIOS, NFS v2 and v3, FTP
Point-in-time Copies (snapshots)	Checkpoints and Fast Checkpoint Restore support
Quotas Support	User, Group, and Directory level quotas support
Directory and Name Services	AD (LDAP, Kerberos v5), NT 4.0 Multiple Master Domains (MMD), DNS, WINS, NIS, NIS+, Local files
Supported Clients	If a client implementation follows the NFS version 2 or 3 protocol or the CIFS specifications, it is supported with the StorEdge 5x10 NAS Appliance.
Dynamic Shares Support	Autohome shares feature
Remote Management	SNMP, Remote Syslog
System Administration	Web GUI (HTTP/Java based), Telnet, Console Command Line Interface (CLI)
Mail Services	SMTP
System Monitoring	<ul style="list-style-type: none"> • Disk subsystem via SCSI Enclosure Services (SES) • Fans, power supplies, temperature, voltages via IPMI protocol • UPS • Network monitor
Tape Backup	<ul style="list-style-type: none"> • NDMP v2 and v3 • Backup with full Unicode • “Remote Configuration” backup support with Veritas NBU
RAID Manager	Integrated through Web-based GUI
Optional Software	<ul style="list-style-type: none"> • StorEdge File Replicator Software for IP-based remote replication • StorEdge Compliance Archiving Software

Key Messages

With the introduction of the Sun StorEdge 5310 NAS, Sun continues to offer customers value, innovation and choice.

Value

- **Simplified Management:** The StorEdge 5310 NAS provides an intuitive installation wizard for easy setup and configuration, and advanced software features that simplify network storage management. This simplified management enables organizations to consolidate dedicated storage from multiple servers and to scale as capacity requirements grow. The browser-based user interface with remote web administration enables system management and monitoring from anywhere in the world.
- **Investment protection:** The StorEdge 5310 NAS incorporates state-of-the-art components that allow a seamless integration into any existing network infrastructure. It enables organizations to consolidate storage, reducing unused pools of expensive storage, and to scale as capacity requirements grow. Combined with simplified management, maximized system availability, and no additional user license fees required, the StorEdge 5310 NAS provides an effective investment protection for current and future growth requirements.
- **World-class Services :** Sun's world-class service organization will provide the same high level of service and support for your StorEdge 5310 NAS Appliance as it provides for a multi-million dollar Sun storage solution implementation. Internal automated diagnostics and remote notification capabilities, combined with a design that implement redundant components, help simplify the serviceability of your StorEdge NAS Appliance.
- **Sun's "One-Stop Shop":** Buy your StorEdge 5310 NAS, software, service, consulting and training from Sun's "One-Stop Shop." Sun makes it easy – we understand secure enterprise network storage computing. Sun has what you need and Sun is here to help you.

Innovation

- **Advanced Technology:** The Sun StorEdge 5310 NAS combines software and hardware technology in an easy-to-manage, high performance storage appliance. The Sun StorEdge NAS operating system is optimized for file serving and file storage and uses a 64-bit journaling file system specifically designed to ensure data integrity. Sun StorEdge File Checkpoint and File Replicator are important applications that provides an additional level of data protection. The hardware design is based on the server-class Intel Xeon platform with dual PCI-X buses that provide excellent performance.
- **Compact Design:** The system architecture of the StorEdge 5310 NAS offers two built-in 10/100/1000 Gigabit Ethernet ports, a single external Ultra320SCSI multimode (SE/LVD) port, two serial ports, and 6 PCI-X slots for expandability, all within a compact 2U form factor design.
- **Robustness:** The journalized filesystem helps to ensure data integrity across unforeseen events such as power outages. Checkpoints provide static images of the filesystem enabling rapid recovery of mistakenly deleted files or objects, and facilitating backup of the system. Combined with hardware RAID, battery-backed RAID cache, hot swappable drives, global spares, redundant hot-swappable power supplies and optional UPS backup, the StorEdge 5310 NAS provides exceptional robustness.
- **Scalability:** The StorEdge 5310 NAS has 4GB of memory and ships standard with six PCI-X slots. A standard Fibre Channel configuration can support the connection of up to 2 StorEdge 5300 RAID EU's and as many as 14 StorEdge 5300 EU F's for a total raw capacity of over 32.7TB. Standard S-ATA configurations can support the connection of up to 2 diskless StorEdge 5300 RAID EU's and as many as 16 StorEdge 5300 EU S's for a total raw capacity of over 89TB. Configurations incorporating 4 StorEdge 5300 RAID EU's must utilize the "Get to Yes" program and have a 60 day lead-time.

Choice

- **Multiple Architectural Choices:** Sun offers choices in capacity and network connectivity to help tailor the StorEdge 5310 NAS to the customers environment.

Target Markets

The StorEdge 5310 NAS is ideally suited for departmental, workgroups, distributed enterprises, and small- to medium-sized companies with UNIX network environments and running any of the following applications:

- E-mail Servers: StorEdge 5310 NAS Appliances provide cross-platform file sharing support for e-mail applications such as SendMail and Lotus Notes.
- Storage Consolidation: The cross-platform file sharing support built into StorEdge 5310 NAS Appliances aids in reducing the need for organizations to dedicate large storage resources to any single server.
- Engineering, CAD: StorEdge 5310 NAS Appliances provide cross-platform file sharing support for engineering applications such as CAD and computer software R&D.
- Imaging and Graphics: Huge graphic files demand maximum performance that can only be achieved with a filer that is optimized for storage services. StorEdge 5310 NAS Appliances with StorEdge File Cluster can be best utilized for this purpose.

Some applications impose extremely high and sustained transaction-per-second loads, require very high QoS (low latency and low latency fluctuation) and have very high availability requirements necessitating fault-tolerance and/or failover capabilities. Examples of such applications are credit card transaction processing and airline ticketing. Other applications, such as Microsoft Exchange 2000, are designed to only run on DAS (Direct Attached Storage) or SAN (Storage Area Networks) devices¹. The StorEdge 5310 NAS Appliance would not be well-suited to these environments. Other Sun products such as the StorEdge 6300 series and the StorEdge 9900 series are specifically designed to meet the requirements of these applications.

The following chart lists applications with particular vertical markets. The checkmark does not necessarily indicate that the application is specific for that industry, but that a company within that industry could have a need for the listed application. For example, Home Directory/File Sharing is not a Telco application, but a company in the Telco industry could use the StorEdge 5310 NAS to provide Home Directory/File Sharing capabilities to its employees.

Application/Industries	Telco	Education	Mfg.	Gov't	Financial	Health Care	Retail	Business Services
Imaging (Document, Medical, GIS)	✓	✓	✓	✓	✓	✓	✓	✓
Home Directory/File Sharing	✓	✓	✓	✓	✓	✓	✓	✓
CAD/CAM			✓					
VoIP/Video	✓							
Web Server Consolidation Storage	✓	✓	✓	✓	✓	✓	✓	✓
DW/DM	✓	✓	✓	✓	✓	✓	✓	✓
Email (SendMail, Notes)	*	*	*	*	*	*	*	*
Microsoft Exchange	X	X	X	X	X	X	X	X
OLTP	X	X	X	X	X	X	X	X

Legend:

✓ = Recommended for StorEdge 5310 NAS.

X = Not recommended with StorEdge 5310. Other Sun products should be used.

* = Non notable installations on this application

¹ In early 2004 Microsoft announced NAS device support for Microsoft Exchange if NAS devices are built on Windows Storage Server 2003 and pass the Microsoft qualification process. This announcement by Microsoft does not state that Exchange will not work on NAS devices that are not based on WSS03 (most other vendors use a proprietary NAS implementation). It just says that Microsoft will not provide support for it

Sun StorEdge NAS OS provides advanced software features that facilitate the deployment of StorEdge 5310 in the recommended applications. The table below lists the main software features that apply for the specific applications, along with the benefits that these features provide in those application environments.

Features/Application	Imaging	Home Dir/ File Sharing	CAD/ CAM	VoIP/ Video	Web Server Consolidation Storage	DW/DM
Checkpoints	✓	✓	✓	✓	✓	✓
StorEdge File Replicator	✓			✓	✓	✓
StorEdge File Cluster	✓	✓	✓	✓	✓	✓
StorEdgr Compliance Archiving Software	✓	✓				
Heterogeneous Access	✓	✓	✓			
Dynamic Volume Expansion	✓	✓	✓	✓	✓	✓
Quotas Support		✓	✓			
Autohome Shares		✓	✓			

Sun StorEdge File Checkpoint Software: Enables the creation of point-in-time images of a StorEdge NAS filesystem taken at the volume level. Creation speeds are almost instantaneous, without disrupting normal operation, and consume minimal storage space. While the active file volume can be modified with read/write operations, a virtual volume produced at the time of checkpoint creation remains available in a static, read-only state.

- Provide a facility for self recovery of lost and/or accidentally deleted objects
- Provide for business continuance without requiring massive expenditures for completely redundant or mirrored systems
- Increase backup window, easing the burden on administrators to complete backup operations between production cycles, and obviates the need to undergo length-time and operator-intensive tasks such as restores from tape
- Facilitates migration to / from new applications

StorEdge File Cluster: A feature of the StorEdge 5310C NAS that enables a dual active-active server architecture for the StorEdge 5310C NAS Appliance. StorEdge File Cluster helps to enable nearly zero down time by eliminating any single point of failure when combined with other high-availability features. The architecture relies on a dedicated heartbeat interface between the server modules which triggers a transparent online failover of all the network interfaces in the event of a server module failure. In case of a failure the system failover takes place within a few seconds, making this an ideal feature for business and other mission critical applications.

- Active-active architecture may be used to enhance performance and/or availability
- Failover typically performed in less than 5 seconds

StorEdge File Replicator Software: An optional StorEdge 5310 NAS application that provides IP-based remote replication of data. Individual StorEdge NAS systems volumes can be mirrored from any system to any other system (even StorEdge NAS systems from different families).

- Provides for business continuance, even in the face of catastrophic loss of the primary (source) system
- Fast reversion from source to target, providing maximum data availability and protection

- Network (Ethernet) based mirroring means lower cost, and provides the ability to exploit technology familiar to users and without requiring expensive gateways

StorEdge Compliance Archiving Software: An optional StorEdge 5310 NAS application, the Sun StorEdge Compliance Archiving System couples the Sun StorEdge 5310 NAS Appliance with Sun StorEdge Compliance Archiving Software to provide compliance-enabling features for authenticity, integrity, ready access, and security. The compliance archiving software was designed from the ground up in consultation with information management compliance and enterprise content management industry experts to help address the most stringent requirements for electronic storage media retention and protection. For more in depth information, please consult the StorEdge Compliance Archiving Software “Just The Facts”.

Heterogeneous Access : Enables file-level access over an IP-based Ethernet network for NFS, CIFS, and FTP file access protocols.

- Enables robust and safe file sharing in heterogeneous UNIX, Windows, and Linux environments

Dynamic Volume Expansion: Provides the capability to add storage to a filesystem without downtime, maintaining data availability.

- Downtime due to capacity limitations eliminated
- Can be effected during normal business operations

Quotas Support: Allows for the restriction of disk space or number of files written to volumes in the StorEdge 5310 NAS. This limit can be determined for a user or group (user or group quota) or for a directory (directory tree quota).

- Enables administrators to quickly allocate space to departments or individuals
- Enables administrators to partition large, individual volumes from a capacity standpoint, providing for volume-like control of space allocations but without requiring them to administer multiple actual volumes
- Simple backup policies (e.g., volume-based full, differential, incremental backups) can be retained across the entire volume, even if capacity is being managed locally

Autohome Shares: Automatically creates temporary SMB/CIFS shares when a user logs on to the system and removes the shares when the user logs off.

- Facilitates IT Administrator’s management tasks by eliminating the need to create and manage shares manually for potentially thousands of users

Sun Software Support

- VERITAS NetBackup v5.x
- Enterprise Backup Software v7.2
- StarOffice v7

Sun StorEdge 5310 NAS



Just the Facts

Third Party Application Support

- Oracle (Database)
- Adobe Acrobat Suite
- Microsoft Office

Relevant Standards Supported

The following standards are supported on the Sun StorEdge 5310 NAS:

- Java™ technology
- Intelligent Platform Management Interface (IPMI)
- IEEE 802.3 - Ethernet (10Mb/s)
- IEEE 802.3U - Ethernet (100Mb/s)
- IEEE 802.3ad – Link aggregation/Port Aggregation Protocol
- IEEE 802.3ab - Ethernet (1000Mb/s, Twisted Pair) Relevant standards supported
- IEEE 802.3z – Ethernet (1000Mb/s, Fiber) with multimode fiber NIC
- SNMP versions 1 and 2
- NDMP versions 2 and 3

Features and Benefits

Features	Benefits
RELIABILITY	
Flash-resident OS	<ul style="list-style-type: none"> OS resides on flash disk, which is more reliable than a traditional magnetic disk
External Hardware RAID	<ul style="list-style-type: none"> External Fibre Channel hardware RAID-5, with redundant controllers and mirrored, battery-backed cache
AVAILABILITY	
Redundant Server Power Supplies	<ul style="list-style-type: none"> Reduces system outages in the event of StorEdge 5310 NAS power supply failure
Redundant Cooling Fans	<ul style="list-style-type: none"> Reduces system outages in the event of cooling fan failure
Redundant Disk Subsystem Power Supplies (each Expansion Module)	<ul style="list-style-type: none"> Reduces system outages in the event of StorEdge 5300 Expansion Unit power supply failure
NIC Failover	<ul style="list-style-type: none"> Provides for a redundant path to the same IP address on the server through a different switch and using a cooperative (with the host) Failover topology
Link-down Failover	<ul style="list-style-type: none"> Enables the system to deal intelligently with an intermittent link problem but sustain a simple port swap or switch by an administrator User-configurable for the number of outages within a specified period
Port Aggregation	<ul style="list-style-type: none"> Enables the use of the same IP address across multiple NIC ports from a single switch, increasing bandwidth and providing some degree of fault tolerance
StorEdge File Cluster (StorEdge 5310C)	<ul style="list-style-type: none"> Transparent failover enables greater uptime
SERVICEABILITY	
<p>Multiple components are hot-swappable:</p> <ul style="list-style-type: none"> Server head power supplies Disk subsystem power supplies Disk subsystem I/O modules Disk Drives 	<ul style="list-style-type: none"> Ensures that the maximum level of data availability is provided by eliminating downtime when failed components need to be replaced or when proactive maintenance is performed
Dual-boot OS	<ul style="list-style-type: none"> Enables the user to quickly and easily backtrack to a prior version of the OS should problems manifest during or subsequent to an OS upgrade operation
Tool-less Access	<ul style="list-style-type: none"> Facilitates maintenance by eliminating the need for tools when servicing the system (most components)
Indicator LEDs	<ul style="list-style-type: none"> Enables operations personnel to easily identify problematic components and prevents mistakes in servicing
DIAGNOSTICS	

Email Diagnostics	<ul style="list-style-type: none"> Enables Sun's support representatives or customers to quickly get a clear and comprehensive report describing system configuration and health
Memory Tracking Facilities	<ul style="list-style-type: none"> mbuf tracking for network issues malloc tracking
Onboard Packet Capture	<ul style="list-style-type: none"> Enables Sun's support representatives to create packet capture sequences describing particular protocol errors, thereby empowering the Tech Support team to diagnose customer situations more quickly
Automated Diagnostic Gathering	<ul style="list-style-type: none"> Gathers and saves diagnostics in the event of system problem
MONITORING	
Disk Subsystem Monitoring via SCSI Enclosure Services (SES)	<ul style="list-style-type: none"> Standards-based disk subsystem monitoring for enclosure, controller, power supply, fan and disk health
UPS Monitoring	<ul style="list-style-type: none"> RAID controller cache switched to write-thru mode if living off of battery (UPS) power, helping to ensure data integrity. Graceful shutdown effected if battery charge drops below acceptable range System does not revert to write-back mode until on AC power and battery sufficiently charged
IPMI Support	<ul style="list-style-type: none"> Enables the system to monitor its own server components (power supplies, fans – including the dedicated processor fan and temperature) and inform the user should any problems be discovered, thereby preventing downtime due to catastrophic failure
NOTIFICATION	
LCD Posting	<ul style="list-style-type: none"> Provides a local reference enabling an administrator to get a picture of overall system health when viewing the unit in its immediate presence
Syslog (Local)	<ul style="list-style-type: none"> Enables the user to exploit a standard and familiar facility for providing log-based information
Remote Syslog	<ul style="list-style-type: none"> Enables the administrator to log StorEdge NAS OS events on a remote system (usually in concert/context with events from other systems) using a standard and familiar facility
SNMP Support	<ul style="list-style-type: none"> Provides for simple integration of StorEdge NAS into an environment where systems are monitored centrally using familiar, industry-standard tools (Sun NetManager, Computer Associates Unicenter, IBM Tivoli, HP OpenView)
SMTP Email	<ul style="list-style-type: none"> Provides a simple, effective and ubiquitous mechanism for critical system event notification
FILE SYSTEM	
Dynamic Volume Expansion	<ul style="list-style-type: none"> Volume expansion can be effected during normal business operations Downtime due to capacity limitations eliminated

64-bit Filesystem	<ul style="list-style-type: none"> • File system scales to 16TB volume sizes – larger volume size support eases administrator’s management tasks by reducing the number of objects that they must manage • Up to 512 file systems • 2TB+ LUN (RAID group) capacity • Up to 256 RAID groups per NAS head for single-head config's; 128 for cluster
Journaling Filesystem (JFS)	<ul style="list-style-type: none"> • Journaling – a two-phase commit operation – helps to ensure filesystem integrity across unforeseen events like power outages. In the two phase commit process, the first phase of journaling happens on NVRAM on the RAID controller. The second phase commits the data to disk thus ensuring data integrity. (See ‘On-disk Journaling’) • Enhances data availability by obviating the need for lengthy filesystem check operations subsequent to power outages or other unplanned events • Aids with data integrity by journaling both metadata and user data (for synchronous writes), ensuring that the client and StorEdge NAS Appliance have a consistent perspective of the on-disk data
On-disk Journaling	<ul style="list-style-type: none"> • The StorEdge NAS products journal to disk – as opposed to journaling to an in-head NVRAM card. The lack of data on the head unit precludes any possibility of a data integrity issue due to head unit failure or malfunction.
Directory Tree Quotas	<ul style="list-style-type: none"> • Enables administrators to allocate space quickly to departments or individuals • Enables administrators to partition large, individual volumes from a capacity standpoint, providing for volume-like control of space allocations but without requiring them to administer multiple actual volumes • Simple backup policies (e.g., volume-based full, differential, incremental backups) can be retained across the entire volume, even if capacity is being managed locally
SUN STOREDGE FILE CHECKPOINT	

Implementation	<ul style="list-style-type: none"> • Provides for business continuance without requiring massive expenditures for completely redundant or mirrored systems • Point-in-time image of filesystem increases backup window, easing the burden on administrators to complete backup operations between production cycles • Creation speeds are sub-second, almost instantaneous, providing administrators a very useful, powerful tool which obviates the need to undergo lengthy and operator-intensive tasks such as restores from tape • Sophisticated technology retains only those filesystem blocks that actually change, e.g., a change to a 16k portion of a 3GB file would create a checkpoint of size 16k, enabling administrators to exploit the capability without drastically increasing space requirements • Facilitates migration to / from new applications • Checkpoints use a “copy-on-write” technology which is the technique of duplicating disk blocks only as they are modified. • Checkpoints provide a facility for self recovery of lost and/or accidentally deleted objects
StorEdge File Checkpoint Scheduling	<ul style="list-style-type: none"> • Provides for automatic creation of checkpoints based upon a user-defined schedule • Provides for user-defined retention criteria and automatic deletion of checkpoints at specified intervals
Local Directory Checkpoint Access (LDCA)	<ul style="list-style-type: none"> • Enables users to self-recover (i.e., without administrator assistance) prior versions of key files in the event they are accidentally deleted • Obviates administrator involvement, restore from tape drives etc. in the recovery process
Fast StorEdge File Checkpoint Restore	<ul style="list-style-type: none"> • Provides for a metadata-effected restore of a checkpointed file • Facilitates recovery of accidentally-deleted data – even by a user • Does not necessitate recovery of an entire volume (as does some competitor implementations)
DATA PROTECTION FACILITIES	
Backup	<ul style="list-style-type: none"> • Remote/over the LAN utilizing NDMP with VERITAS NetBackup Server and NDMP option

<p>StorEdge File Replicator: Block-level Mirroring (option)</p>	<ul style="list-style-type: none"> • Provides for business continuance, even in the face of catastrophic loss of the primary (source) system • All writes communicated in time order from source to target, guaranteeing a consistent view of the state of user data at the remote site, any point in time. (Note that asynchronous mirroring does not guarantee that the state of the local and remote sites is identical or changes simultaneous, only that each site is consistent and updates in the same sequential order.) • When promoted to source (or master), the remote system always provides access to every in-order transaction it received from the original source (master), thereby providing the optimal architecture for data preservation in an asynchronous mirroring model • Preserves write ordering, ensuring that files remain consistent on the mirror target, and ensuring that up-to-the-second updates are available should mirror promotion occur • Individual StorEdge NAS Appliance volumes can be mirrored from any system to any other Sun StorEdge NAS family system, providing the user with flexible and cost-effective options for assuring data availability • Fast reversion from source to target, providing maximum data availability and protection • Network (Ethernet) based mirroring means lower cost, and provides the ability to exploit technology familiar to users and without requiring expensive gateways.
OTHER	
<p>IP Aliasing</p>	<ul style="list-style-type: none"> • Enabler for active/active NIC's • End user can set up a single port with multiple IP addresses, facilitating network management and topology layout
<p>Multi-byte Character Support</p>	<ul style="list-style-type: none"> • Unicode supported, facilitating deployment of StorEdge NAS throughout the world
<p>DHCP</p>	<ul style="list-style-type: none"> • Facilitates installation in environments where DHCP is active, enabling the system to acquire an IP address (which is posted on the LCD) dynamically, which the user can then access to setup the system
<p>Installation Wizard</p>	<ul style="list-style-type: none"> • Facilitates installation by automatically detecting first-time access and leading the user through an intuitive, 10-15 minute set of dialogs, after which the StorEdge NAS system will be largely ready to use
<p>LCD Keypad</p>	<ul style="list-style-type: none"> • Simplifies installation by communicating success or failure, and acquired IP address if successful • Provides the ability to effect a 'soft' shutdown from the front panel • Provides the ability to set an IP address should DHCP fail • Posting area for communication of system problems (e.g., disk failures, fan failures, power supply failures etc.)
<p>Multiple Administration Interfaces</p> <ul style="list-style-type: none"> • Web GUI • Telnet 	<ul style="list-style-type: none"> • Provides the user with a choice of a user interface, enabling them to select and use the interface with which they feel the most comfortable

SSH Support	<ul style="list-style-type: none"> Provides for secure, remote menu access to StorEdge NAS administrative facilities
Native Administration Tool Support	<ul style="list-style-type: none"> Facilitates administration thorough the provision of support for tools with which the administrator is already familiar, e.g., Windows Explorer can be used to set directory/file permissions
Comprehensive Quota Support <ul style="list-style-type: none"> User Quotas Group Quotas Default Quotas 	<ul style="list-style-type: none"> Enables the administrator to simply and easily prevent rogue applications or users from consuming too much disk space. Default quotas provide for a simple mechanism to set a global limit for all users and/or groups without setting each user/group individually
Cross-Protocol Support	<ul style="list-style-type: none"> Simplifies the administrator's storage management tasks by providing for access in heterogeneous environments, enabling administrators to create and manage storage without artificially partitioning it between Unix and Windows clients Facilitates file sharing between groups and users in the enterprise
User/Group Mapping	<ul style="list-style-type: none"> Facilitates administration in heterogeneous environments by enabling the administrator to map Unix users/groups to Windows users/groups, either through automatic means or manually
Cross-Protocol File Locking	<ul style="list-style-type: none"> Enables robust and safe file sharing in heterogeneous environments. Cross-protocol file locking is limited to blocking NFS writes to a file when a SMB/CIFS client has the file open. The converse is not implemented, in part because the often stateless nature of NFS access makes it difficult to determine when a file is being modified via NFS.
Cross-Protocol Security Mapping	<ul style="list-style-type: none"> Guarantees security by providing for consistent security policies for objects created through one filesystem protocol (e.g., Unix) and then accessed through another (e.g., Windows)
Permission Synthesis	<ul style="list-style-type: none"> Facilitates securing of data by synthesizing permissions during access of objects via non-creating protocols – for example, Windows permissions are synthesized if an object created under Unix is accessed through Windows
NTP Support	<ul style="list-style-type: none"> Enables the administrator to utilize industry-standard means of synchronizing StorEdge NAS time with network time (the clocks of all other servers on the network)
FTP	<ul style="list-style-type: none"> Provides a very fast, familiar and industry-standard tool for uploading and downloading files to StorEdge NAS
Autohome Shares	<ul style="list-style-type: none"> Temporary SMB/CIFS shares that are automatically created when a user logs on to the system and removed when the user logs off IT Administrator's do not have to create and manage shares manually for potentially thousands of users

Enabling Technology

Technology Overview

The Sun StorEdge 5310 NAS utilizes the following technologies:

- Flash Disk Module for Operating System
- FC and S-ATA storage
- 2U chassis dual redundant power supplies and fans
- Intel Xeon Server-class Motherboard and CPU
- Software/Operating System

Flash Disk Module for Operating System

The Sun NAS Operating System resides on a 256MB solid-state Flash Disk Module that connects directly to the IDE port. The main features the Flash Disk Module are:

- Solid-State (no moving parts)
- High Shock and Vibration Limits
- Rugged and compact
- 512 Byte Sector and ECC Defect Management Compatible to IDE Hard Disk Drives

Storage

The StorEdge 5310 NAS supports up to 2 StorEdge 5300 RAID EU's with up to 7 StorEdge 5300 EU F's connected to each StorEdge 5300 RAID EU for a total of 16 “trays”. Each “tray” holds fourteen 146GB FC drives for approximately 2.044TB raw capacity in each tray.

Useable capacity calculated as follows:

Each 14 drive StorEdge RAID 5300 EU has 2 preconfigured RAID groups of 5+1 with 2 hot spares (4 drives for overhead). The 146GB FC drive has a useable capacity of 133GB.

14 drives – 4 drives (overhead) = 10 drives * 133GB = 1.330TB useable per StorEdge RAID EU tray

Each 14 drive StorEdge 5300 EU F has 2 preconfigured RAID groups of 6+1 (2 drives for overhead). The 146GB FC drive has a useable capacity of 133GB.

14 drives – 2 drives (overhead) = 12 drives * 133GB = 1.596TB useable per StorEdge EU F tray/25.04TB useable per fully configured Fibre Channel system.

For S-ATA configurations, the above guidelines apply with the exception that the StorEdge 5300 RAID EU's are **diskless** and each diskless StorEdge RAID EU supports up to 8 StorEdge 5300 EU S's. Each S-ATA tray holds fourteen 400GB S-ATA drives for approximately 5.6TB raw capacity in each tray.

Each 14 drive StorEdge 5300 EU S has 2 preconfigured RAID groups of 6+1 and 5+1 with 1 hot spare (3 drives for overhead). The 400GB S-ATA drive has a useable capacity of 364GB after formatting.

14 drives – 3 drives (overhead) = 11 drives * 364GB = 4.004TB useable per StorEdge EU S tray/64.064TB useable per fully configured S-ATA system.

Configurations greater than 32.7TB up to 65.4TB (raw) for FC and over 89TB up to 179TB (raw) for S-ATA will be supported via the “Get to Yes” program with extended lead-times. These configurations require 4 StorEdge 5300 RAID EU's attached to a single StorEdge 5310 or 5310C NAS Appliance.

2U Chassis

The 2U chassis is a standard 19” rack-mountable with dual redundant and hot-swappable 500 watt power supplies and redundant fans. The chassis’ front contains multiple LEDs indicators. Refer to the Server Architecture section for a detailed diagram of all the components on the front of the chassis.

Motherboard Board Information

The Sun StorEdge 5310 NAS is based on the Intel® Xeon-class Server Board SE7501WV2. The main features of the motherboard are listed below.

- Support for one Intel® Xeon™ processor with 512KB L2 Advanced Transfer Cache, Hyper-Threading Technology, and a 533MHz system bus.
- Support for 4 GB of registered ECC DDR200/266 SDRAM
- Two integrated Intel® PRO/1000 MT Server Gigabit Network Connections (Intel® 82546EB controller)

StorEdge 5300 RAID Expansion Unit & 5300 Expansion Unit



Figure 5. Front panel of StorEdge 5300 RAID Expansion Unit




Figure 6. Rear of StorEdge 5300 RAID Expansion Unit

The Sun StorEdge 5300 RAID Expansion Unit (EU) and 5300 EU F and EU S provide modular, rackmountable and scalable back end storage to the StorEdge 5310 NAS Appliance. The minimum configuration consists of the StorEdge 5310 NAS Appliance and a single StorEdge 5300 RAID EU with dual-active FC RAID controllers integrated and either 14 or 0 dual ported fibre channel drives in a 3U enclosure. Dynamically, seven StorEdge 5300 EU F trays can be added to the StorEdge 5300 RAID EU to create a maximum 1x8 configuration with 112 drives (eight trays). Eight StorEdge 5300 EU S trays may be added to a **diskless** StorEdge 5300 RAID EU for a maximum 1x9 configuration with 112 drives.

Software/Operating System

The Sun NAS Operating System has been designed for the sole purpose of high-performing file serving and file storage. It utilizes the following technologies:

- 64-bit Journaling Filesystem
- File Sharing for NFS and CIFS clients
 - Unified Lock Management
- Active Directory Services
- Sun StorEdge File Checkpoint
- Sun StorEdge File Replicator (optional)
- Sun StorEdge File Cluster (optional)
- High Availability Bonds (NIC Failover/Port Aggregation)

64-bit Journaling Filesystem:

The StorEdge NAS OS filesystem is a 64-bit journaling filesystem that guarantees filesystem integrity across unforeseen events, such as power outages, enhancing data availability by obviating the need for lengthy filesystem check operations subsequent to power outages or other unplanned events. The OS supports up to 512 file systems with file systems scaling up to 16TB in volume size, simplifying an IT administrator's management tasks since larger volume sizes reduces the number of objects they must manage.

Journaling ensures filesystem consistency and fast recovery in the unlikely event of a system crash. The OS maintains a log or journal of the block level activity that has taken place on disk as directed in the two phase commit process.

The two phase commit process is similar to a typical data base environment. When a client executes a transaction, the first phase of the commit process, a write, either synchronous or asynchronous, as determined by the transaction, initiates an entry in the StorEdge NAS memory; the meta-data and content (data-data). In the event of a synchronous write, both the meta-data and the content are subsequently written to disk and the transaction is journaled. Conversely, in a asynchronous write, the content and meta-data are written to memory with only the meta-data being journaled allowing for greater performance and lower latency as experienced by the client and/or transaction. However, due to the architecture of asynchronous writes in both NFS and SMB the client and StorEdge systems will have a consistent perspective of data on disk.

The second phase, the commit, as dictated by the client, "flushes" the content left in the StorEdge memory, as with a asynchronous write, to disk and subsequently a notification is sent to the client from the StorEdge NAS Appliance of the content being written to disk. At this time, content is also flushed from the client's cache thus "completing" the transaction.

In the unlikely event of an abnormal shutdown of the StorEdge NAS Appliance, any lost data can be recreated because any updates to the metadata in directories and bitmaps are written to the journal. This alleviates the need for fsck's.

With this journaling methodology filers can recover from an abnormal shutdown in a few minutes as opposed to requiring a costly, time-consuming consistency check. The journaling filesystem not only returns the data to the pre-crash configuration but also recovers unsaved data and stores it in the location it would have been stored had there not been an unexpected interruption.

Drawbacks encountered due to costly NVRAM topologies are non-existent by journaling to industry standard hard drive technologies. Since the data is on disk, there is no dependency on batteries to maintain journal information, for it is on fault-tolerant RAID drives.

In the event that the filesystem needs to be expanded, an available LUN can be assimilated by an existing volume via a simple operation. This feature provides for the dynamic expandability through the addition of unallocated extension segments while the StorEdge NAS Appliance is online and accepting data I/O. As such, this feature allows an administrator the flexibility to allocate storage on an as needed basis. Because of the dynamic scalability of a filesystem (volume), what was once a time consuming task requiring system downtime can be accomplished without the traditional negatives associated with such a task. Available storage made simple.

File Sharing for UNIX (NFS) and Windows (CIFS)

The Sun StorEdge 5310 NAS delivers secure file services for CIFS and NFS clients. The integration of these two security models presents a challenge due to the many differences between these protocols. Windows file servers use access control lists, or ACL's, to specify access rights for defined lists of users and groups. NFS servers use traditional UNIX permissions - user/group/world - that specify a limited set of permissions. The Sun StorEdge 5310 NAS provides an integrated security model in which Windows ACL's and UNIX permissions are simultaneously maintained for filesystem objects, and which preserves the semantics and features of the native security models.

With filesystem support for both CIFS and NFS security, the Sun StorEdge NAS OS provides an integrated model for file sharing between Windows and UNIX systems. This model includes support for NFS permissions, Windows ACL's, Windows Domain integration, and cross-platform services like credential mapping. The Sun StorEdge 5310 NAS also supports network name resolution via Network Information Services (NIS, NIS+), the Berkeley Domain Name System (DNS) and Microsoft Windows Internet Naming Service (WINS). The following are notable aspects of the design:

- The key to the Sun StorEdge 5310 NAS support for cross-platform file sharing is credential mapping. Credential mapping enables users with both NFS and CIFS user accounts to seamlessly exercise their ownership and access rights from either client platform via an association (or mapping) between the two accounts. This association can be created automatically, or it can be established manually in environments where the Windows and UNIX account names do not match. Credential mapping facilitates file sharing between the platforms by mapping each NFS user or group to a CIFS user or group and vice versa. This allows users to assign access rights to a user or group from either platform using that platform's native toolset.
- Files created via NFS are classified as UNIX files and files created via CIFS are classified as Windows files. Windows files typically have security descriptors associated with them while UNIX files do not. Editing or copying a UNIX file (a UNIX-created file) using a CIFS client will add a security descriptor to it, and cause it to be reclassified as a Windows file. Removing the security descriptor from a Windows file (possible only when configured by the administrator to allow this) will cause the file to be reclassified as a UNIX file. For native file service requests (NFS to UNIX files and CIFS to Windows files) the security model honors the native CIFS or NFS access conventions. For non-native requests, a heuristic model is employed to synthesize permissions of the native (or creating) protocol. The guiding principle for Sun StorEdge 5310 NAS cross-platform file access is that a user will not be granted greater access via a non-native protocol than would have been granted via the native protocol.
- As with Windows systems, each Sun StorEdge 5310 NAS will establish membership in a Windows Domain. When a user connects to a Sun StorEdge 5310 NAS system, the user is authenticated by a Domain Controller (DC) using pass-through security. Domain membership and pass-through authentication allow the Sun StorEdge 5310 NAS to take advantage of

Windows availability and scalability features, such as Primary (PDC) and Backup (BDC) Domain Controllers and inter-Domain trust relationships. An access token, obtained from the Domain Controller, is used to identify the user and contains the user's ID and the ID of each group of which the user is a member. The access token is used to represent the user when an attempt is made to access secured objects, such as files and directories.

- Windows files and directories are secured via security descriptors (SD). Each SD contains the ID of the object owner, the ID of the owning group and an access control list (ACL). The ACL is a list of access control entries (ACE), each of which grants or denies specific access rights to a specific user or group.
- The CIFS security model implemented on the Sun StorEdge 5310 NAS is the standard Windows security model. Each request to access an object contains a set of desired access rights. These desired access rights are checked against the access control information defined in the object's security descriptor to determine whether access should be granted or denied.
- The Sun StorEdge 5310 NAS provides support for several Windows built-in local groups. These groups are the standard Windows solution for assignment of system wide privileges for common administration tasks. A well-known example is the backup/restore privilege granted to the Backup Operators group. These groups can contain users from either the local domain or trusted domains. The Domain Admins group of the domain which the Sun StorEdge 5310 NAS is a member of is automatically given membership in the local group Administrators. This behavior is expected from Windows Domain member servers, and therefore improves compatibility.
- The Sun StorEdge 5310 NAS automatically creates administrative CIFS shares for each volume. This provides access to all files and directories for Domain Administrators without the need to create shares at the root level. Again, this behavior is expected from Windows Domain member servers and improves compatibility with Windows and Windows-aware applications.
- The Sun StorEdge 5310 NAS provides file access from Windows systems via CIFS/SMB with support for either share level security or Windows Domain security. When used in a domain environment, Sun StorEdge 5310 NAS users are authenticated with a Domain Controller using pass-through authentication. In the workgroup security mode, access is authenticated locally on the Sun StorEdge 5310 NAS using a share password. ACL's are not used in workgroup mode, resulting in what we would designate as a "UNIX file". Workgroup security is weak in comparison to Windows Domain Security, and is generally not recommended.
- File Locking. The Sun StorEdge 5310 NAS employs an integrated file locking mechanism that supports the following file locking schemes:
 - Windows (CIFS) whole-file locks
 - Windows (CIFS) byte-range locks
 - Windows (CIFS) Opportunistic locks (optional file locks acquired at file open time, "breakable" by CIFS file access)
 - UNIX (NFS) byte-range locks (advisory locks, via lockd and statd)

In addition to supporting single-protocol file locking, the Sun StorEdge 5310 NAS also protects user data with the support of an integrated cross-platform file locking mechanism. This functionality prevents NFS clients from writing to files that are open via the CIFS protocol.

File locking caveats:

- Level II Opportunistic Locks (oplocks) are not supported.
- Cross-protocol file locking is limited to blocking NFS writes to a file when a SMB/CIFS client has the file open. The converse is not implemented, in part because the often stateless nature of NFS access makes it difficult to determine when a file is being modified via NFS.

Sun StorEdge 5310 NAS



Just the Facts

Active Directory Services

Active Directory is the Windows directory service that provides access to domain information such as users, groups and shared resources. Active Directory clients access this information on the network using the Lightweight Directory Access Protocol (LDAP). The AD relies on the Internet Domain Name System (DNS) to provide name resolution services. DNS is the industry standard service used throughout the Internet to resolve hostnames to Internet Protocol (IP) addresses. The DNS provided with AD supports the ability for clients to dynamically update their entries in the DNS database; this is known as dynamic DNS. Sun StorEdge NAS OS supports LDAP version 3, which is defined in RFC2251. Further information is also available on the Microsoft web site.

To support Windows environments, Sun StorEdge NAS OS provides the ability to publish shares in AD and update DNS records using dynamic DNS. AD clients can then access Sun StorEdge NAS OS shares by looking in the Active Directory. Shares appear in AD as folder objects that point directly to the appropriate shares on Sun StorEdge NAS OS. Share can be placed in any container in which a share folder object can be created. One of the properties of a share folder object is its UNC name. The Universal Naming Convention (UNC) name specifies both the server on which the share exists (i.e. server.sun.com) and the name of the share, i.e. \\sun.com\sharename. AD provides the ability to control access to share objects within the directory service, in addition to any access controls placed on the Sun StorEdge NAS OS directory being shared.

Sun StorEdge NAS OS supports Kerberos version 5 authentication for secure AD and dynamic DNS updates. Kerberos uses secret-key cryptography to provide a network authentication protocol for client/server applications. Each secure AD or dynamic DNS update is performed within the context of an AD user. Sun StorEdge NAS OS communicates with a Key Distribution Center (KDC), which normally resides on a domain controller, to authenticate the user prior to performing an update²

The AD domain is equivalent to concatenating the NT domain name and the DNS domain name. For example, if the NT domain is SALES and the DNS domain name is sun.com, the AD domain would be sales.sun.com.

The username and password provide the user credentials used to authenticate AD updates. The user must be a valid Window user, defined in AD, that has the administrative rights required to perform secure AD updates, i.e. a Domain Administrator.

The User Container field specifies the container in which the username resides; for example, the administrator user resides in the users folder. The field format is in LDAP distinguished name notation, without the domain. Thus, for the administrator user, the container would be specified as: cn=users

If the user definition resides in an organizational unit folder (for example, user_folder) then the container would be specified as: ou=user_folder

If the user resides in a sub-container folder of a parent folder, then the container would be specified as: ou=sub_folder,ou=parent_folder. Note the reverse order definition.

Sun StorEdge NAS OS provides support for dynamic DNS, which can be used to automatically update the IP address and hostname of a Sun StorEdge NAS OS server. In an AD environment, it is important that the DNS information is correct since AD clients will use DNS to locate the Sun StorEdge NAS OS when they access shared resources.

² If the administrator user is used to update AD or DNS and the update fails then the DES encryption keys required for Kerberos may have to be regenerated by changing the administrator user password. This is only required for the administrator user and the previous password can be reused.

Sun StorEdge NAS OS updates DNS when it boots up or when DNS settings are saved. Dynamic DNS must also be enabled in DNS. Sun StorEdge NAS OS can update a DNS using the secure or non-secure mechanisms. It will first attempt to update DNS using the non-secure method. If this fails, it will try the secure method. In order for Sun StorEdge NAS OS to update a DNS zone securely, the zone has to be an AD integrated zone, meaning the zone is stored in AD. Sun StorEdge NAS OS will use Kerberos to authenticate an AD user before DNS can be updated securely.

Sun StorEdge File Replicator

Leveraging the Sun NAS OS's advanced journaling filesystem, the StorEdge File Replicator application can be used to create a copy of a volume on a remote StorEdge NAS Appliance. This is performed using the Safe Asynchronous Mirroring (SAM) architecture which ensures file system consistency at both source and target. The following are important aspects of the StorEdge File Replicator architecture:

Source/Target Interconnection

A StorEdge File Replicator implementation consists of a target and source file volume, along with the interconnecting TCP/IP network that is responsible for propagating the data between them. The target StorEdge NAS Appliance can be located virtually any distance from the source. The target StorEdge NAS Appliance can be on the same campus as the source, or on the other side of the world, as long as the interconnecting network has sufficient bandwidth to carry the data across. There are three options:

- **One-to-One Mirroring:** This is the most straightforward and popular approach to mirroring. There are one Master system and one Target system in this configuration and is described above.
- **Many-to-One Mirroring:** In Many-to-One mirroring, several Master systems are mirrored to a single Mirror target. Many-to-One mirroring is frequently used by customers for Disaster Recovery purposes, to collect exact duplicates of multiple systems at multiple remote locations at a single, central site, where recovery scenarios can be centrally managed.
- **Bi-directional Mirroring:** Bi-directional Mirroring refers to the ability for systems at sister locations to mirror to each other. For instance, a system in Los Angeles may be configured to mirror its volumes to a sister system in Houston, which in turn and simultaneously mirrors its volumes to the Los Angeles system. In the event either site experiences a problem, the data is readily available at the sister site.

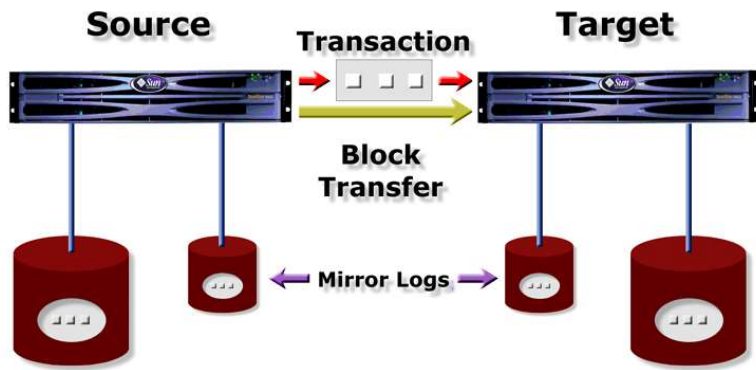
Block-level Mirroring

With StorEdge File Replicator, data replication is performed at the block level. All checkpoints, symbolic links, and other data structures existing on the target will be identical to those of the source volume. In fact, once the source and target have been synchronized, the target will be an exact, block-level image of the active disk. StorEdge File Replicator's block-level implementation dramatically reduces the required bandwidth on the interconnecting network by transmitting only those data blocks that have changed. For example, if only a single 4 KB block of a one hundred megabyte (MB) file is updated, only that 4 KB block is transferred, not the entire 100 MB file.

Safe Asynchronous Mirroring (SAM)

In order to ensure data protection without compromising source StorEdge NAS Appliance performance or introducing checkpoint latency, StorEdge File Replicator was architected with Safe Asynchronous Mirroring technology, a unique asynchronous replicating scheme that utilizes database technology

to ensure file system integrity. On both ends of the mirror, there are dedicated mirror logs that log changes to both source and target file systems. Updates to a target StorEdge NAS Appliance are comprised of a series of write transactions, each containing a number of data block updates. The figure on the right shows the mirror logs and the transactions propagating between StorEdge File Replicator source and target. Each transaction is sent from source to target in an ordered fashion and entered into the target mirror log. Each of these write transactions contain a number of key file system updates such as inode delete and update.



Data Integrity

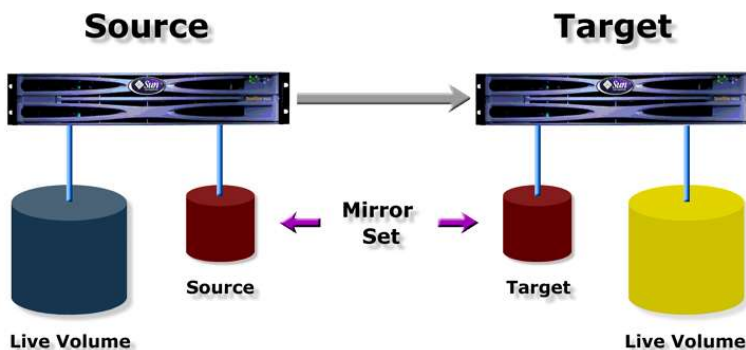
Extensive effort has gone into ensuring that the target's data does not lag that of the source, thereby maintaining a high state of synchronization. Due to the asynchronous architecture of StorEdge File Replicator, the target may lag the source volume by an amount of time depending on the source StorEdge NAS Appliance's load and the speed of the network connection. StorEdge File Replicator mirrors on a block level to ensure high performance, but does not commit data to the target volume on transactional boundaries. This guarantees the integrity of the file system on the target and also preserves write ordering.

Error Handling

The SAM architecture also aides in handling broken network connections between the source and target. If the intraconnecting link is temporarily broken, the mirror log on the source continues to record file system block changes, serving as a buffer for unsent data. When the link is reestablished, the logged data can then be propagated to the target StorEdge NAS Appliance. This ensures that temporary link-down scenarios can be gracefully handled and that the target data is brought in sync with the source in a timely fashion.

Deployment of Target StorEdge NAS Appliance

StorEdge File Replicator replicates data at the volume level. Both of the mirrored StorEdge filers need not be identical and can have other normal, active volumes. This provides the IT manager the flexibility to decide which volume to mirror, and to utilize both filers for other file serving purposes.



Management

The source StorEdge NAS Appliance acts as the focal point for StorEdge File Replicator management. Configurations are performed on the source StorEdge NAS Appliance, which in turn affects the configurations changes on the target StorEdge NAS Appliance. Sun NAS OS's Java-based graphical user interface (GUI) provides an intuitive and easy-to-use browser-based interface for StorEdge File Replicator management. Management can be performed on any browser-equipped workstation that can access the target StorEdge NAS StorEdge NAS Appliance's network. This fits with StorEdge NAS' appliance approach, simplifying data management by eliminating the need to install additional software on management workstations.

StorEdge File Replicator Applications

StorEdge File Replicator can be used to help address many data management challenges facing IT professionals today:

Disaster Recovery: Without a reliance on slow tape media, StorEdge File Replicator eliminates the need for lengthy tape restores. StorEdge File Replicator enhances recovery time in case of a complete loss of data, as businesses can now access mission-critical data from an online backup on a mirror StorEdge NAS Appliance. With StorEdge File Replicator properly configured, the target, or mirror, is assured to be an accurate, near realtime representation of the offline source StorEdge NAS Appliance. The StorEdge target file volume can be brought online quickly to ensure uninterrupted operations.

Backup: A StorEdge File Replicator target volume may be dedicated for backing up source volumes. Without affecting production operations, replicated data can be backed up on the target. StorEdge File Replicator enhances operations by moving backup I/O to the remote volume. This shadow processing capability reduces CPU load on the production StorEdge NAS Appliance, streamlining operations.

Sun StorEdge File Cluster

Traditional server-attached storage architectures are equipped with numerous features that help maintain data availability through most hardware failures, but few of them provide true redundancy throughout the entire data I/O channel. Hardware features such as hot-swappable hard drives and power supplies only provide a part of the story. In order to ensure that data is available at all times, a different breed of storage product is necessary to provide redundancy of all hardware components. The Sun StorEdge 5310C NAS Appliance helps to ensure that a path to critical data is always available.

Overview

The clustered failover architecture of the Sun StorEdge 5310C NAS can be succinctly explained as redundant from top to bottom. As shown in Figure 7, the system consists of a pair of active-active server modules that share access to pairs of RAID controllers and multiple LAN subnets. The RAID controllers are connected to each server module via fibre-channel connections. A dedicated heartbeat cable connects the two server modules and monitors the corresponding partner's health status.



Figure 7. A StorEdge File Cluster configuration

In normal operation, each server module (or head unit) operates independently, with responsibility for a subset of the disk volumes. Simple disk, fan or power supply failures are handled gracefully using internal redundant components in the same manner as with any Sun StorEdge NAS Appliance. However, if a StorEdge 5310C server module suffers a hardware failure that renders a data path unavailable, its healthy peer server will automatically assume ownership of the RAID resources formerly managed by the failed head. The active server module will also assume the identity of the failed server on the network, in addition to maintaining its own.

With dual active-active server modules, dual active-active RAID controllers, and RAID volume distribution through several drive arrays, a Sun StorEdge 5310 NAS Cluster allows for a failure in nearly any single major hardware component without rendering data unavailable.

The foundation of the ultra high-availability features of the Sun StorEdge 5310 NAS Cluster is its ability to respond to hardware failures by initiating a server failover. All facets of the failed server, including RAID volume ownership and network interface addressing are transferred in this failover operation. Later we shall see how the redundant servers and RAID controllers explained below work together to provide a fault-tolerant failover solution.

RAID Architecture

In addition to the inherent resiliency of Redundant Array of Inexpensive Disks (RAID) technology on the Sun StorEdge 5310 NAS Cluster, additional failover functionality is offered with dual active/active RAID controllers, a standard, integral part of this storage solution. By providing an added layer of protection, it extends availability above and beyond standard server failover. It also delivers an important layer of abstraction between the server module and the disk drives; providing a redundant, reliable platform on which to calculate processing-intensive RAID parity information. This allows the server module to focus its precious CPU cycles on the delivery of data on the network.

The intelligent RAID controller delivers field-proven fault-tolerance that protects against most types of failures. If any component on the data path (fibre channel cable, RAID controller, or server fibre channel host adapter) causes the SE5310 server to lose communication with a controller, the RAID controller automatically reroutes all I/O operations to the other available controller. Further down the hardware path to the drive level, the fibre channel architecture ensures that redundant paths are available to all drives.

Server Architecture

Acting as the data pumps that make high-availability NAS possible, the StorEdge 5310 NAS server modules provide the interface between the disks and the network. Two identical, highly optimized server modules running a proprietary microkernel operating system comprise the brains of the StorEdge 5310 NAS Cluster and provide compatibility with both Windows and UNIX clients. With kernel-level support for both CIFS and NFS, the StorEdge 5310 NAS provides an integrated model for sharing data between Windows and UNIX systems, including WINS, ACL, domain and Active Directory Services (AD) support, and support for network name and user/group resolution via Network Information Services (NIS), NIS+, and the Berkeley Domain Name System (DNS). Beyond the operating system imbedded throughout all of the products in the Sun StorEdge NAS product line, the StorEdge 5310 NAS Cluster includes additional intelligence that provides active-active failover functionality. This added functionality includes special failover support for the network and storage interfaces.

Network Interfaces

With redundant network interfaces, the StorEdge 5310 NAS ensures that an active port is available at all times. In case of a failure anywhere on the data path managed by one of the server modules, a predefined network interface of the server module not affected by the failure will assume the other's network identity. In this failover situation, using IP aliasing³, the partner's NIC assumes the MAC and IP addresses of the NIC in the failed server module. This fast transition is almost transparent to the network clients, simply detected as a few seconds of no traffic. The clustered failover architecture supports the following NIC interfaces: single Ethernet/Fast Ethernet, dual Ethernet/Fast Ethernet, and Gigabit Ethernet interfaces. As a general rule, the interface types in both filers should match.

Storage Management

Armed with dual active-active RAID controllers, the Sun StorEdge 5310 NAS offers the utmost in RAID

3 IP Aliasing IP is simply a mechanism that enables a single physical port to assume responsibility for more than one IP address.

redundancy. This high-availability storage architecture would be worthless, however, if the server modules were not able to manage the transition of the RAID logical units (LUNs) from one controller to the other. Advanced RAID monitoring and control features specifically engineered for the StorEdge 5310 provide the ability to respond to failures that require a RAID controller failover. In the event of such a failure, requests to the failed RAID controller would be rerouted to the active server module.


In the event of a failover caused by failure in any piece of the system other than the RAID, e.g. a network interface failure, the available server module would immediately initiate a RAID controller failover. As a result, the downed server module would transfer the “ownership” of its assigned LUNs to the healthy module.

Heartbeat Interface

A critical piece in the clustered failover architecture is the interconnection between the symmetric server modules. Each server module is equipped with a dedicated network port that constantly monitors the status of the other module using a heartbeat connection. This dedicated connection is a local, point-to-point ethernet connection utilizing a passive ethernet network cable. Communication between the server modules includes overall health status, notification of configuration changes, and synchronization of shutdown or reboots across the modules.

Failover/Failback

The failover operation of the Sun StorEdge 5310C NAS Appliance is designed to provide an alternate data path in case of a failure that renders a single data path unavailable. The failover transition does not impact the operation of the healthy server module. Once a failure is detected, the ownership transfer usually is done under 5 seconds and hence remains mostly transparent from the end user’s perspective. The following diagrams illustrate the failover process. The normal state, where both server modules are functioning in an active-active configuration:

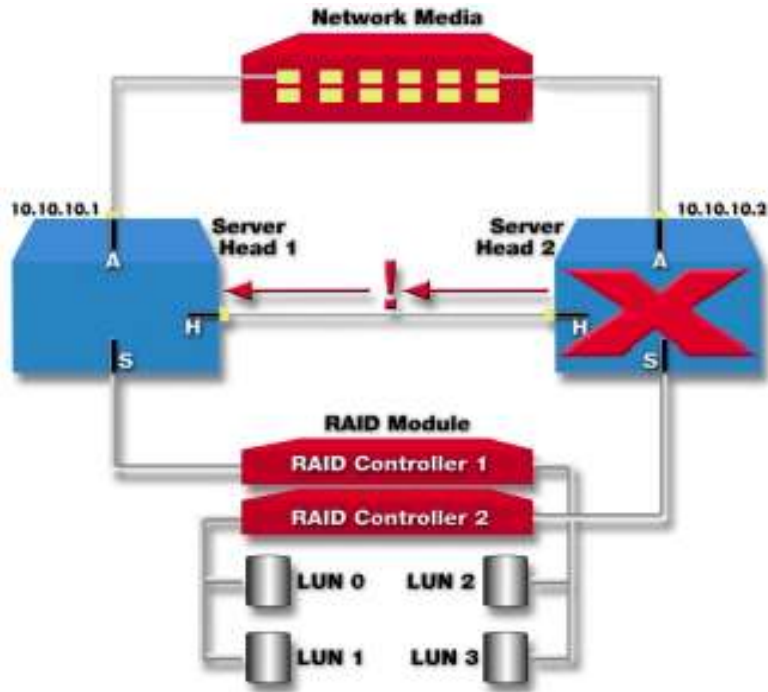


A failure in the second server module has been detected and is communicated via the dedicated
Sun StorEdge 5310 NAS

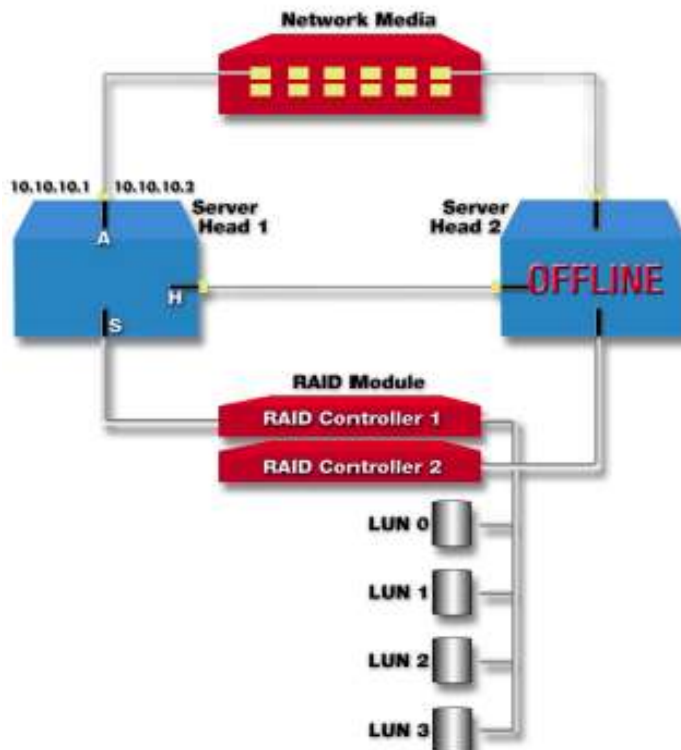
Just the Facts



heartbeat interface. Failover is initiated:



The failover process is completed, with Server 1 handling all requests. Note that the NIC on the Server Head 1 now also assumes the network identity of the failed head, while the connected RAID controller is managing all of the drive arrays:



The recovery process – known as failback – can be initiated when the failed head is repaired and is brought back online. In this operator-initiated recovery process, the failed server module reassumes the ownership of its original storage volumes, while the NICs in the recovered head also return to their original MAC and IP address settings. Like failover, the failback process also is completed in seconds.

Network Visibility

Depending on the timing of the failover or failback and the type of clients served there will be varying impacts on network clients. Clearly, a failover or failback that occurs when the StorEdge 5310 NAS is not serving clients will be transparent when the clients later make a request. The following describes varying scenarios in which a failure happens while the StorEdge 5310 NAS is servicing a client machine with an active connection.

NFS Clients

Given the statelessness of the NFS protocol, a typical NFS client will gracefully handle a StorEdge 5310C NAS failover. The client would simply detect that the connection has been broken and then attempt to reconnect. Once the failover procedure is completed, the connection is reestablished with the healthy server module now acting as the active network interface. Similarly, failback is handled robustly by the stateless reconnection attempts of the NFS protocol.

CIFS Clients

Conversely, CIFS clients do not respond to network disconnects in the same manner as NFS clients. This is due to the fact that CIFS is a stateful protocol, so the failover transition is not transparent to the client machine. Once the connection is broken and detected by a CIFS client, the client must manually re-establish the connection. However, to make the process as smooth as possible in Microsoft Windows network environments, the Sun StorEdge 5310 NAS uses NetBIOS name aliasing.

Quite simply, with this feature the standby NIC assumes the NetBIOS name of the active NIC in the process of failover transition. This reduces the effort involved in reestablishing a lost connection to simply clicking on the mapped directory in the client's window. The effect on the clients when the system recovers (failback) is the same as with failover: NFS clients are transparently recovered and SMB clients might need to re-establish their sessions.

The Sun StorEdge 5310 NAS with StorEdge File Cluster option delivers a higher level of data availability by offering complete fault tolerance through full hardware redundancy and integrated software failover support. Quite simply, this translates to the elimination of the dreaded “single point of failure” that can cause the loss of data availability. By extending the functionality of the base StorEdge NAS operating system, Sun has built a storage architecture on which mission-critical data can be stored with confidence.

Sun StorEdge File Checkpoint

Sun StorEdge File Checkpoint is a point in time image (or “checkpoint”) of a filesystem taken at the volume level that resides on a NAS device. While the active filesystem(s), can be modified with read / write (R/W) operations, the “copy” or logical image can only be read and is not modifiable. It is a virtual and static read-only point-in-time “copy” of the active filesystem. A “copy” or “copies” co-exist with existing, live, filesystem(s). One should not consider a point in time image as a true “back-up” of the active filesystem. Its virtual existence is directly relational to the live physical filesystem's existence.

While the functionality of using Sun StorEdge File Checkpoint is diverse and vast, the process of implementing and managing checkpoints is simple. When a checkpoint is requested by an authorized user or is automatically “taken”, the time required and performance overhead are negligible to both the network infrastructure and the live, active filesystem.

Access to checkpoints is provided via multiple protocols: NFS and CIFS clients. There can be up to sixteen (16) checkpoints maintained at any one time per filesystem (volume) on a NAS device.

Sun StorEdge File Checkpoint Applications

- **Rapid File Recovery:** When files are modified, deleted, checkpoints can be used to quickly access a recent image of the file.
- **Online backup of files:** With its checkpoint feature, Sun NAS OS allows the flexibility to backup large file systems without worrying about data being changed during the backup process. This exclusive feature ensures that such file systems will always be backed up in a consistent state.
- **Database backups:** By checkpointing database files on a StorEdge NAS, it is not necessary to take a database offline during the time it takes to back it up. Instead, the database can be paused momentarily, then a checkpoint taken. The checkpoint will contain an image of the database files at a given point in time, which in turn can be backed up at the user’s convenience.

Sun StorEdge File Checkpoint Benefits

- **Reduced Recovery Time:** point in time images dramatically reduce the time it takes to recover lost or damaged files and filesystems. Recovering files from tape is no longer the only option. Point-in-time images of the active filesystem (volume) are accessed in a read-only state and accessed via CIFS share or NFS mount point assuming proper access rights.
- **User Satisfaction:** In the event of lost, corrupted, or damaged files, system administrators are no longer faced with the complex task and lengthy process of recovering to a “clean” version of the file in question. In a Windows environment, accessed through CIFS share, the file can be “dragged” and “dropped” from the checkpoint folder into the proper volume. From a NFS mount point, the file in question can be copied from the checkpoint directory to the proper location.
- **Reduced Cost (TCO):** Lower TCO is achieved, in the event of recovering lost files, in time savings of the administrators doing the recovery. The time usually spent in searching archived files on tape, the “normal” wear-and-tear of tapes and tape devices, and having to do multiple recoveries in the event the wrong file being recovered allows for time better spent doing more meaningful tasks.

Sun StorEdge File Checkpoint Explained

The Sun StorEdge NAS OS file system has i-nodes, which are data structures within the file system holding specific file related information, including the addresses of data blocks. As shown in Figure 1, a live file volume has i-nodes that describe the addresses of the blocks that compose any given file volume.

This addressing is achieved by using pointers to reference the data blocks that are stored elsewhere on disk.

Sun StorEdge NAS OS creates a checkpoint by maintaining all pointers to disk blocks currently in use at the time of the checkpoint. As illustrated in Figure 2, when the checkpoint is taken, a new inode is created to maintain existing pointers to data blocks. Any changes made to files after a checkpoint has been established will result in updates to the current set of pointers. Figure 3 shows the state of the volume after data blocks are updated. This process is similar to when changes are made to non-checkpointed volumes, which also update pointers to point to the new disk blocks.

Because of this adherence to standard volume update processes, performance degradation with checkpoints is limited. The overhead experienced was a zero change in file system size, and 5% performance degradation in terms of iops on a standard benchmark applied to checkpointed volume. It is important to note that the benchmark imposed a very high proportion of write operations against the file system, therefore degradation is actually quite minimal. Actual effects on performance will vary depending on a number of factors including file system size, number of files, number of delete operations, depth of directories, average file size, and outstanding number of checkpoints.

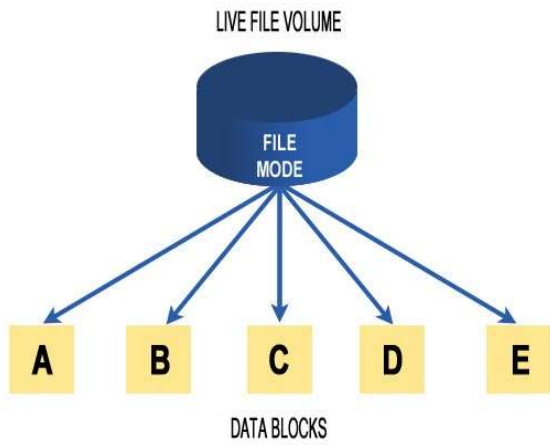


Figure 6 Normal file volume before checkpoint is taken

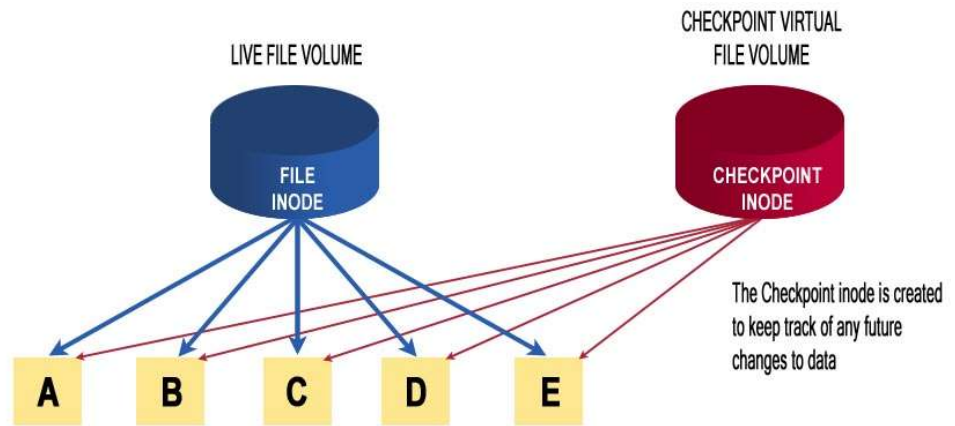


Figure 7 : File volume after checkpoint is taken




Figure 8 New data blocks are created when files are updated

High Availability Bonds: (NIC Failover / Port Aggregation):

The following describes the High Availability (HA) bond features available on the Sun StorEdge NAS OS. In addition, it may also be used to configure a network interface card (NIC) failover.

Network interface bonding is a method of grouping multiple physical NIC's together to form a single logical NIC. StorEdge 5310 NAS systems support bonding for two distinct environments; load sharing and redundancy.

Load sharing bonds are referred to as **port aggregation (PA)** bonds, or Fast EtherChannel. Redundancy bonds are known as high availability (HA) bonds. There can be a maximum of 4 bonds per StorEdge NAS Appliance head.

Port aggregation bonds are used for two different purposes. First, PA bonds can increase the effective network connection bandwidth and secondly, to provide one form of NIC fault tolerance in which the system can tolerate a NIC failure and still provide a path to the network. If a NIC fails, the connection will continue to operate on the remaining operational NIC's.

When establishing a bond, there are considerations to keep in mind. All NIC's in a PA bond must be connected to the same switch and both ends of the network connection must support this form of bonding. In other words, when a Fast EtherChannel NIC is connected to a switch the switch must support Fast EtherChannel technology. Additionally all of the member NIC's must participate in the sending and receiving of packets with each end, the switch and the NIC being responsible for performing load-balancing across the switch's configured ports. The NIC selection algorithm uses the destination IP address as part of the port selection criteria to ensure that all packets to the same host are sent via the same NIC. This will minimize out-of-order packets which can have a detrimental effect on network performance.

High Availability (HA) bonds provide fault tolerance through **NIC failover**. This is an active-passive configuration in which one NIC within the bond is active and handling network traffic while the remaining NIC's are passive backup devices. In the event that the active NIC fails, the system will automatically switch over to one of the backup NIC's in the bond. In this configuration no support is required from the switch and there is no single switch restriction. The member NIC's of an HA bond may be of disparate types and may be connected to different switches or hubs -- avoiding a single point of failure at the switch or hub.

Table 1 provides a comparison of the features of port aggregation (PA) and high availability (HA) bonds.

	<i>HA</i>	<i>PA</i>
Load sharing	No	Yes
Bandwidth scalability	No	Yes
Switch support required (i.e. Fast EtherChannel protocol)	No (any hub/switch)	Yes

Table 1 Bond feature comparison; at GA, one single-port MMF NIC and two on-board 10/100/1000 Ethernet ports provide all supported connectivity.

High Availability (HA) Bond

From an operational point of view, an HA bond works exactly like a single stand-alone NIC. Administration is performed on the master NIC within the bond but with the additional features that each type of bond provides. A bond can have up to 4 IP aliases as in a single stand-alone NIC. User administration tasks, such as changing an IP address, adding a netmask, and deleting aliases, etc., are done through the (master) bond itself and not on the slave NIC's.

Sun StorEdge 5310 NAS



Just the Facts

Operation

When a bond NIC is configured the main NIC becomes the active NIC of the bond. All packets directed to the bond are sent and received through the active NIC. Once a HA bond is up and running, the link status of the active NIC is periodically monitored. If the link status is down then the NIC is considered failed and a backup NIC with a “good” link status is selected and configured to be the active NIC of the bond.

The time it takes for a NIC failover process to complete is minimal. Fail over time is usually less than a few seconds and is dependent on the type of NIC's that are being used in the bond. For link media that supports speed negotiation such as the ‘emc’ and ‘fxp’ NIC, the driver takes a little longer to complete its initialization. Transmission latency introduced to the HA bond software is minimal since no complex decision needs to be made at transmission time. There is no additional latency on packet reception since incoming packets do not go through the bond software.

Recovery

A NIC recovery is performed thus switching the active NIC back to the main NIC. This usually is done to restore network traffic back to the main NIC which might be a higher bandwidth link (1000T) than the backup NIC (100T). The Telnet and WEB GUI are enhanced to support NIC recovery operations.

System Architecture

Overview

The StorEdge 5310 NAS Appliance can be mounted in industry-standard 19-inch equipment racks. A slide-rail/extension kit may be ordered with the NAS Appliance; this slide-rail/extension kit allows you to mount each of these filers into a StorEdge Rack.

Hardware Components

The StorEdge 5310 NAS is 3.50 inches (87 mm) high, 16.93 inches (430 mm) wide and 25.51 inches (648 mm) deep. The airflow direction is from front to back and internal fans are included. I/O ports are located on both the front and rear panels. Informational LEDs are located on the front panel. Access to the two power connections is at the rear of the chassis.

Figure 9. StorEdge 5310 NAS front




Figure 10. Front LED & I/O detail (with bezel removed)





Figure 11. StorEdge 5310 NAS back; 1-port UTP and 2-port MMF NIC's not available together



StorEdge 5300 Expansion Units

Overview

The Sun StorEdge 5300 Expansion Units are a high-performance, enterprise-class, full 2 gigabit per second (Gb/s) Fibre Channel solution that combines outstanding performance with the highest reliability, availability, flexibility and manageability. The Sun StorEdge 5300 Expansion Units consist of two primary components, Controller Tray – 1x1 and Expansion Tray – 0x1. Each of the primary components is fully fault tolerant with automated I/O path failover and redundant, hot-swappable components to ensure the highest availability.


Sun StorEdge 5300 RAID EU

The Sun StorEdge 5300 RAID EU, Controller Tray – 1x1 houses the dual-active, intelligent FC RAID array controllers, provides four 2 Gb/s host connections and supports 14 Fibre Channel drives. Its redundant FC RAID controllers and a host-based failover driver provides automated I/O path protection and ensures continuous data access. In the event of a component failure anywhere in the data path between server and disk drives, I/O is re-routed down the surviving path for uninterrupted access to the data.

Two 2 Gb/s drive channels are configured as redundant FC-AL loops to the drives and any additional Expansion Tray – 0x1. Both channels are active for an aggregate bandwidth of 400 MB/s. The Sun StorEdge 5300 Expansion Unit system supports up to 112 low-profile Fibre Channel drives in the Expansion Trays.

The Sun StorEdge 5300 RAID EU, Controller Tray has two Ethernet ports, one per controller, for remote management from anywhere on the network through Sun StorEdge Configuration Service software.

All primary components of the Sun StorEdge 5300 RAID EU, Controller Tray are hot-swappable, field-replaceable units (FRU's) to help prevent servicing downtime and to help minimize the time to repair a failure. Each FRU, including the FC RAID controller, fans, power supplies and disk drives can be easily accessed and removed or replaced .



Sun StorEdge 5300 RAID EU

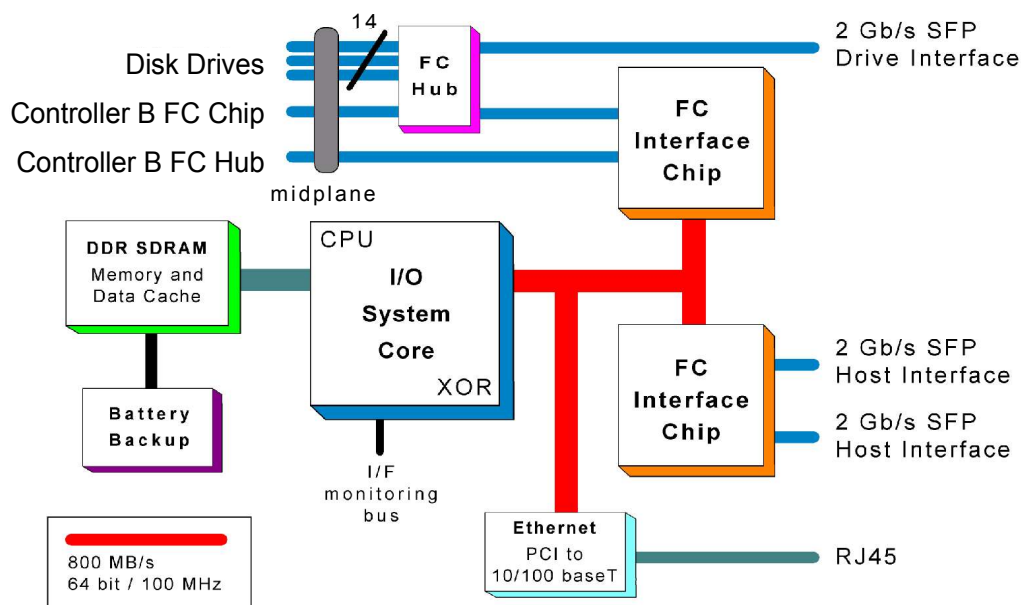
Figure 12. StorEdge RAID EU

Sun StorEdge 5300 RAID EU FC RAID Controller Card

The Sun StorEdge 5300 RAID EU FC RAID controller card benefits from over a decade of design knowledge and firmware development focused entirely on open systems and high-speed, disk-based performance.

The StorEdge 5300 RAID EU controller cost effectively combines high-speed, industry-standard components with robust firmware to deliver enterprise-class functionality and exceptional performance value.

Each StorEdge 5300 RAID EU FC RAID controller card has two 2 Gb/s FC host I/O ports (four per dual-controller storage system), supporting direct host or SAN attachments, and two 2 Gb/s FC drive interfaces. Integrated FC hubs connect the drive-side FC interface chips in each controller – creating dual drives loop that provide both controllers redundant access to all attached disk drives.



Sun StorEdge 5300 RAID EU, FC RAID Controller card

Figure 13. StorEdge 5300 RAID EU controller card

The StorEdge 5300 RAID EU FC RAID controller card's I/O system core provides built-in hardware XOR for high-speed RAID parity calculations, enabling this compute-intensive task of to be handled efficiently and effortlessly. Its 1.6 GB/s bus between the I/O system core and cache ensures lightning fast communication between these critical components.

The StorEdge 5300 RAID EU FC RAID controller card is equally adept at both I/O per second (IOPS) and throughput (MB/s). Its 800 MB/s bus (64 bit, 100 MHz) between the core processor and I/O chips has both the "width" to handle large-block I/O (sustaining up to 400 MB/s on 512k sequential disk reads) and the "speed" to process large amounts of small-block I/O (sustaining 25,000 IOPS on 4k random disk reads).

The Sun StorEdge 5300 RAID EU's RAID and cache algorithms are designed to provide best possible performance without the need for extremely large cache. This is accomplished with a cache management design that has been developed and perfected over the past decade.

The use of large cache in storage has long been driven by the unique requirements of mainframe environments. Open systems servers, however, are seldom able to make good use of large storage cache. Their applications, which typically involve high-levels of random I/O activity, are increasingly unlikely to experience disk subsystem "cache hits." In fact, large-cache "monolithic" systems can substantially increase the cost of the storage without providing any significant performance benefit.

The StorEdge 5300 RAID EU was designed from the outset for open system I/O requirements – recognizing that disk performance is vastly more important than cache performance in these environments. Open systems servers feature their own large, well-managed, on-board caches, and gain more benefit from an efficient storage controller design that maximizes disk performance than from additional cache.

Cache Usage

The Sun StorEdge 5300 RAID EU dynamically utilizes cache for both read and write caching. This allows for more efficient use of available cache by adapting its utilization to match the I/O requests from attached hosts as the needs of the applications change during the computing day. This adaptability is one reason that the StorEdge 5300 RAID EU can sustain higher levels of real-world performance than monolithic and modular systems with larger caches.

Write Policy

The Sun StorEdge 5300 RAID EU offers both write-back and write-through caching options to optimize performance or protection. The write policy determines whether I/O completion is signaled when data is transferred to the cache or when data is written to disk. With write-back caching, the write request is signaled complete as soon as the data is in cache, and actual writing to disk occurs at a later time. Write-through caching signals the completion of a write request when the data is safely stored on disk. Write-back caching creates significantly higher write performance for applications with a high locality of reference, and its inherent risk is eliminated as the StorEdge 5300 RAID EU's battery backup protects the cache at all times. Write performance with a write-through cache is approximately that of a non-cached system, but ensures the greatest data protection.

Write Cache

The Sun StorEdge 5300 RAID EU's cache algorithms are designed to provide the lowest latency I/O access possible when securing the user's data in storage system cache. And the StorEdge 5300 RAID EU ensures that its cache is ready to receive the next write request from the host by efficiently writing cached data to disk in a manner that minimizes disk I/O. By identifying high locality of write data in its cache, the StorEdge 5300 RAID EU is able to group all data in a RAID stripe thus eliminating the need for disk read operations in the computation of parity in RAID 5. This ensures that disk utilization is minimized for a particular usage pattern, allowing maximum I/O throughput per spindle available.

Cache Protection

Cache mirroring protects unwritten data during a controller or cache memory failure. When cache mirroring is enabled, the data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding

write operations. This cache parameter is set on a per volume basis. Data in the controller cache memory is also protected in case of power outages. The StorEdge 5300 RAID EU FC RAID controller cards contain batteries that protect the data in cache by keeping a level of power until the data can be written to the drive media.

Cache Flushing


To prevent data loss or corruption, the controller periodically writes unwritten data residing in its cache to disk (flushes the cache) based on one of two factors:

- A time-based flush writes cache data to disk after a user-defined period of time (in seconds). This value can range from immediate to infinite, with a value of 10 seconds or less recommended for optimal data protection.
- A threshold-based flush occurs when the amount of unwritten data in cache reaches a certain level, called a start percentage. The controller writes cache data to disk until the amount of data in cache drops to a stop percentage level. For example, the controller can start flushing the cache when the cache reaches 80 percent full, and stop flushing the cache when the cache reaches 20 percent full. The start and stop percentage levels are user-defined to maximize data protection and/or performance.

Sun StorEdge 5300 Expansion Unit

The Sun StorEdge 5300 EU, SBOD Expansion Tray is also known as the Common Storage Module, CSM 100. The StorEdge 5300 EU, Expansion Tray is more than “just-a-bunch-of-disks” with features designed to optimize performance, availability and serviceability. Imbedded “loop switch”, redundant, dual-active drive loops access the dual-ported Fibre Channel drives ensuring maximum performance and complete accessibility to all drives in the event of a loop or cable failure.

Like the Sun StorEdge 5300 RAID EU, controller tray, all the primary components in the Expansion Tray are customer replaceable and hot-swappable. This includes the Fibre Channel disk drives, cooling fans, power supplies and FC I/O modules (IOMs). The IOMs deliver online information about component status, such as temperature, fan speed, or faults, to the storage management software and provide the ability to cascade multiple drive modules to a single array module.



StorEdge 5300 EU

The expansion tray supports up to 14 low-profile Fibre Channel disk drives in a 3U enclosure, enabling the Sun StorEdge 5300 EU storage system to maximize “spindle density” and capacity in a minimum amount of space.

Switched BODS (SBODS)

The switched expansion tray has an imbedded “loop switch” which isolates each drive on a private loop with the environmental services module (ESM), enabling direct and detailed FC communication with each individual drive. The switched drive module provides the following key benefits.

Drive isolation enables improved diagnostics for less time in degraded mode

The switched drive module’s IOM-embedded “loop switch” creates what is essentially a private Fibre Channel arbitrated loop between an IOM and a drive. Drive isolation eliminates the risk of a single drive disrupting the loop and potentially causing other drives to fail on the loop (referred to as fault propagation). Drive isolation also enables diagnostics to be performed on an individual drive level, simplifying analysis of drive and loop problems. The longer it takes to identify and diagnose a problem drive, the longer the overall storage system performance is degraded. In loop-based drive enclosures, where multiple drives are operating on a given loop, it can be extremely difficult to identify which drive is faulty. The switched drive module’s drive isolation and direct communication with each individual drive, which allows for more detailed diagnostic information to be gathered, greatly simplifies analysis of drive and loop problems – enabling problem drives to be located and remedied faster.

Fewer nodes creates lower latency and linear scaling for more responsive applications

With loop-based enclosures, each drive is a node on the drive loop. Large loops create a longer and slower I/O path from a controller to a drive and back to the controller. The switched drive module adds only one node (the I/O Module) to each of its redundant drive loops, and its drives are just one point-to-point step away from the now shorter and quicker loop. Fewer nodes and faster communication significantly lowers the latency and provides linearly scalable performance for more responsive applications. Additionally, the switched drive module’s ability to acquire detailed diagnostic information enables more effective performance tuning. A loop-based drive enclosure adds 15 nodes to each redundant drive loop (one IOM and 14 drives). Large loops create a longer and slower I/O path from a controller to a drive and back to the controller. The switched drive module adds only one node (the IOM) to each of its redundant drive loops – a 93 percent reduction. And its drives are just one point-to-point step away from the now shorter and quicker loop. Fewer nodes and faster communication significantly lowers the latency and provides linearly scalable performance for more responsive applications.

Preventive maintenance

The switched drive module puts all drives through an insertion test (port test before insert, or PTBI) before they are added to the Fibre Channel loop. The intent is to identify faulty drives before they can disrupt the loop. PTBI causes a drive being added to the loop to be LIP’d and then looped back to itself. The switched drive module monitors the drive for 5 seconds, and if the loop remains up, the drive is LIP’d again and allowed to join the loop.

Reliability, Availability and Serviceability (RAS)

Reliability

- Flash-resident Operating System enhances reliability
- System monitoring and diagnostic of Sun StorEdge 5310 NAS (power supplies, fans – including the dedicated processor fan and temperature) and StorEdge 5310 Expansion Unit components. Informs the user should any problems be discovered, aiding in the reduction of downtime due to catastrophic failure
- ECC on data cache
- Error checking and correction on disk drives
- Spare cylinders and skip sectors on disk drives

Availability

- Dual hot-swappable redundant load-sharing/load-balancing auto-sensing 110V AC/240V AC power supplies with individual power cords
- Redundant cooling fans
- NIC Failover provides for a redundant path to the same IP address on the server through a different switch and using a cooperative (with the host) failover topology
- NIC Port Aggregation enables the use of the same IP address across multiple NIC ports from a single switch, increasing bandwidth and providing some degree of fault tolerance
- Journaling Filesystem guarantees filesystem integrity across unforeseen events such as power outages enhancing data availability by obviating the need for lengthy filesystem check operations subsequent to power outages or other unplanned events
- Sun StorEdge File Checkpoint “checkpoints” (also known as Snapshots), provide static images of the filesystem enabling rapid recovery of mistakenly deleted files or objects, and facilitating backup to tape
- Dynamic Volume Expansion provides the capability to add storage to a filesystem without system downtime, maintaining data availability
- UPS management port ensures the graceful shutdown of the StorEdge 5310 NAS in the event of an extended power outage.
- StorEdge File Replicator is an optional software application that allows mirroring, through a TCP/IP network, of Sun StorEdge NAS Appliance volumes from one StorEdge NAS Appliance to another StorEdge NAS Appliance anywhere in the world. It provides for business continuance (data availability) in the event of a catastrophic loss of the primary (source) system. When this occurs, the target may be promoted to source (master) and it takes over all the requests to the failed system.

Serviceability

- Tool-less access facilitates maintenance by eliminating the need for tools when servicing the system (most components)
- Hot-swappable dual redundant hot-swappable power supplies easily accessible from the rear of the system
- Dual-boot OS enables the user to quickly and easily backtrack to a prior version of the OS should problems manifest during or subsequent to an OS upgrade operation
- Embedded system diagnostics tools enable Sun support representatives to quickly get a clear and comprehensive report describing the configuration and health of the system. It also enables support

Sun StorEdge 5310 NAS



Just the Facts

representatives to create packet capture sequences describing particular protocol errors, thereby empowering them to more quickly diagnose and service customer situations

- Fault indicator LEDs are present on the front and rear of the system to easily identify problematic components and prevent mistakes in servicing. A fault indicator LED stays on following a fault even if the system has been powered off (but is still connected to the power source)

Tape Backup and Restore

One of the primary goals of enterprise storage management is to back up and restore information in a timely, secure, and cost effective manner over enterprise wide operating systems.

Companies need high-performance backup and the ability to back up data to local media devices. While the data itself may be distributed throughout the enterprise, its cataloging and control must be centralized. The emergence of network-attached storage and dedicated file servers makes storage management more challenging.

Network Data Management Protocol (NDMP) recognizes that these issues must be addressed. NDMP is an opportunity to provide truly enterprise-wide heterogeneous storage management solutions - permitting platforms to be driven at a departmental level and backup at the enterprise level.

The Sun StorEdge 5310 NAS is supported in certain direct attach, 3-way, and “remote” back up configurations utilizing NDMP:




Figure 15. "Remote back up" configuration




Figure 16. 3-way backup configuration

Please see the SE5x10 WWW (SunWIN/Partner Document Library #414203) for details regarding supported configurations, NDMP tiering, etc.

Server Management

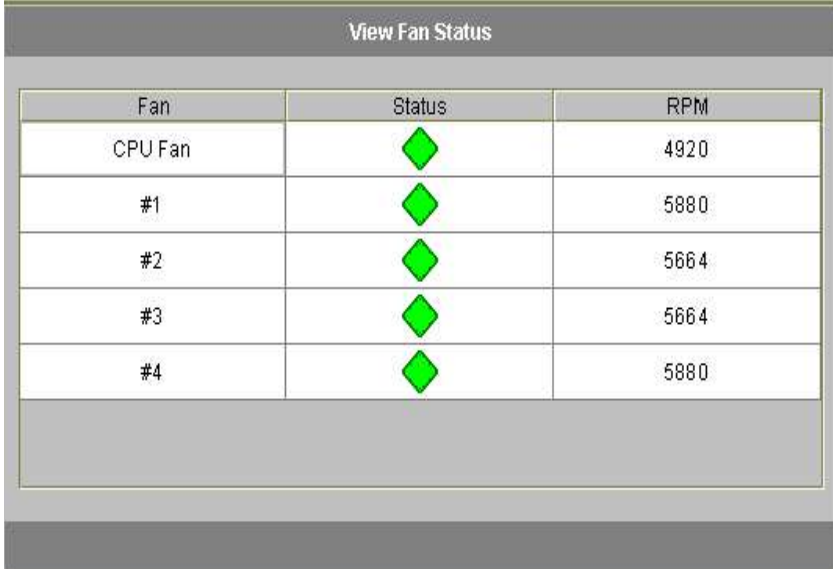
Management of the Sun StorEdge 5310 NAS Appliance is eased by the browser-based user interface with remote web administration for system monitoring from anywhere in the world. The following are protocols and tools that are used to provide system monitoring and management for the StorEdge 5310 NAS Appliance.

IPMI

The OS utilizes Intel's environmental protocol, IPMI (Intelligent Platform Management Interface), to manage and report hardware component statistics. It enables the system to monitor its own server components (power supplies, fans – including the dedicated processor fan and temperature) and inform the user should any problems be discovered, thereby preventing downtime due to catastrophic failure.

System Monitoring

Environmental statistics, such as **fan speeds, chassis temperatures, power, and voltage statuses** are easily viewed through a JAVA based GUI (Graphical User Interface). This offers the NAS administrator a clear, pictorial view of the integrity of the StorEdge NAS Appliance. From this color, graphical depiction of the environmental statuses, appropriate actions can be taken to correct approaching problems or feel confident in how the machine is running and functioning.



The screenshot shows a web-based interface titled "View Fan Status". It contains a table with three columns: "Fan", "Status", and "RPM". The table lists the following data:






Fan	Status	RPM
CPU Fan		4920
#1		5880
#2		5664
#3		5664
#4		5880

Figure 17 View Fan Status

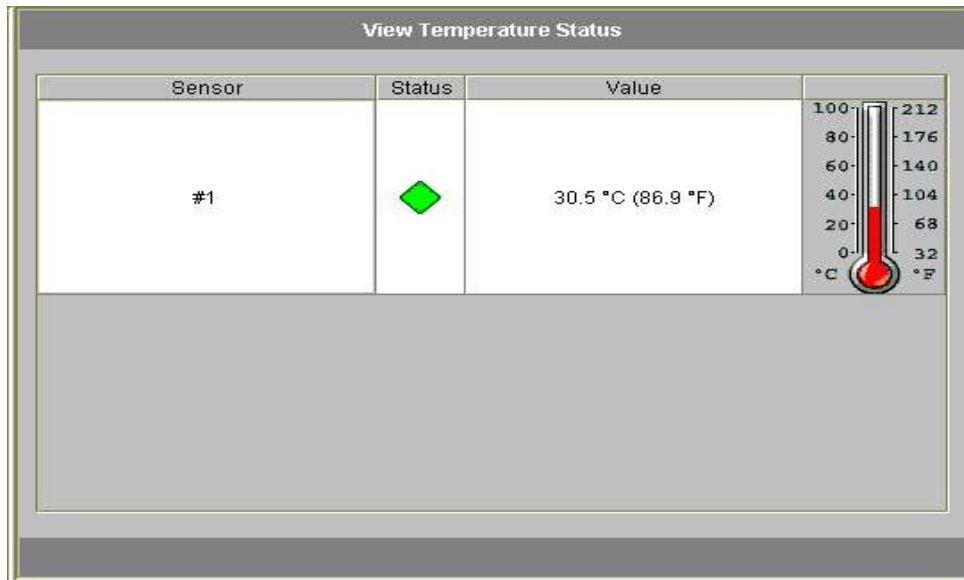


Figure 18. Temperature status panel

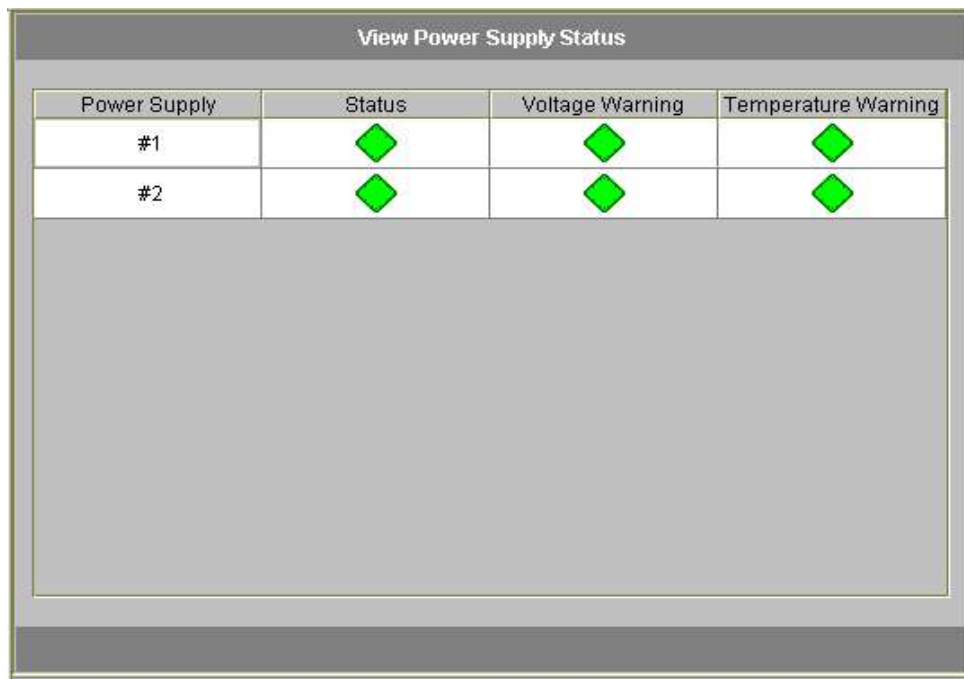


Figure 19. Power Supply status panel







View Voltage Regulator Status		
Voltage Regulator	Status	Current Value
+12V		12
+5V		5.25
+3.3		3.48
Standby		5.01
-5V		-5.07
-12V		-12.24

Figure 20. Voltage Regulator Status panel

SMTP

Using SMTP, an administrator can establish Email notification through a SMTP server. This information can contain either SMTP (Simple Mail Transfer Protocol) server name, as well as designate email notification recipients. In the event of a system error, StorEdge NAS OS will send a detailed email message to recipients via the SMTP server based off of a “Notification Level”. Effectively, the administrator will be notified if/when the StorEdge NAS Appliance encounters situations that have been predefined thus providing quicker and timelier attention to those situations.

This “Notification Level”, which is a system wide setting, is determined when the recipient is added. One of three levels can be chosen:

- **Disk subsystem monitoring via SCSI Enclosure Services (SES):** Standards-based disk subsystem monitoring for enclosure, controller, power supply, fan and disk health.
- **Errors Only:** The server will notify recipients when the system detects errors but not warnings.
- **Errors and Warnings:** The server will notify recipients of all errors and warnings.
- **None:** Effectively disables the email notification option. StorEdge NAS OS will not send any notifications. If “NONE” is chosen as the option, the system should be configured for diagnostic emails at minimum to ensure that proper notification is not missed.

Recipient email addresses are displayed in the **List** box. Once an error has been detected, StorEdge NAS OS logs it in the system log file and sends an email notifications and warnings to the listed recipients as designated.

SNMP

Simple Network Management Protocol is an industry standard to monitor and manage diverse network devices. Supporting SNMP allows the administrator to establish “if/then” rules through SNMP compliant tools allowing faster resolution to adverse scenarios.

One simply enables SNMP by selecting the “Enable SNMP” checkbox in the SNMP administration window. Other attributes are:

- **Server SNMP Community:** SNMP community to which the StorEdge NAS belongs
- **Contact Info:** Name of the person who is responsible for this StorEdge NAS
- **System Location:** The network location which can be logical or physical
- **Destination IP Address:** The TCP/IP address for the server you wish to designate as an SNMP trap destination in the event of system errors
- **Port #:** Port to which StorEdge NAS will send traps
- **Version:** SNMP protocol version (1 or 2)
- **Community:** Community string for the trap destination

Telnet

A Telnet session offers two distinct realms of functionality. The first is an administrative, “green screen” window into the StorEdge NAS OS. From this interface an administrator, security permitting, has the ability to manage all operating functions of the NAS device. As with the GUI, the administrator can alter operational characteristics, configuration parameters, or access controls in a menu driven format.

Secondly, Telnet access allows an administrator, via a secure format, to access a CLI (Command Line Interface). From the CLI, the administrator may issue UNIX style commands furthering operational alterations to the StorEdge NAS OS. Standard MAN and HELP functions detail commands furthering the ease-of-use message.

LCD Keypad

The StorEdge 5310 NAS offers an LCD interface on the front of the unit. The LCD provides a local reference point enabling an administrator to get a "picture" of the overall system health. From the LCD, an administrator has the ability to dynamically set network accesses by setting the GATEWAY, IP address, SUBNET, and NETMASK. The Administrator can also perform system shutdowns and restarts from this panel. Utilizing the "Left/Right" and "Up/Down" arrows, and the "Select" "Exit" buttons, navigating the panel options is made as simple as possible for a NAS administrator.

Syslog/Remote Syslog

Further expanding an administrator's ability to track system activity, a log may be generated and stored locally. Or, the Network System Logger, SYSLOGD, may be enabled. In either instance, storing logs locally or remotely, the administrator can capture a combination of eight (8) different types of messages: emergency, alert, critical, error, warning, notice, info, and/or debug.

If the administrator requires granularity when enabling SYSLOGD, a Facility Value can be identified. Following is a list of Facility Values that the administrator can choose from:

- **Kern:** Messages generated by the kernel
- **User:** Messages generated by random user processes
- **Mail:** The mail system
- **Daemon:** System or network daemons
- **Auth:** Authorization systems such as "LOGIN"
- **Syslog:** Messages generated internally by SYSLOGD
- **LPR:** Line printer spooling system
- **News:** Reserved for USENET network news system
- **Uucp:** Reserved for UUCP system
- **Local0 – Local7:** Reserved for local use

Sun StorEdge 5310 NAS Specifications

Processors

Processor	One 3.06GHz Intel Xeon processor
Cache	512-KB Level 2

Memory

6 DIMM slots, registered DDR-266 ECC SDRAM
System supports 4 GB

Standard/Integrated Interfaces

Network	Two 10/100/1000BaseT Ethernet ports
Serial	One RJ-45 serial port
SCSI	One Ultra320SCSI multimode (SE/LVD)
USB	Two USB 1.1 ports
Expansion bus	Two internal PCI-X 64bit/100MHz
Fibre Channel	1 or 2 dual-port 2Gb/s FC PCI HBA's

Mass Storage

External Disk Fibre Channel	Sun StorEdge 5300 RAID Expansion Unit and Expansion Unit F, each supports 14 x 146GB FC HDDs <ul style="list-style-type: none">• 1 or 2 StorEdge 5300 RAID EU's• Maximum capacity per StorEdge 5300 RAID EU: 2.044TB (raw)/1.330TB (useable)• Up to 14 StorEdge 5300 EU F's (7 per StorEdge 5300 RAID EU)• Maximum capacity per StorEdge 5300 EU F: 2.044TB (raw)/1.596TB (useable)• Maximum capacity: 32.704TB (raw)/25.004TB (useable)
S-ATA	Sun StorEdge 5300 RAID Expansion Unit and Expansion Unit S; EU S supports 14 x 400GB S-ATA HDDs <ul style="list-style-type: none">• 1 or 2 diskless StorEdge 5300 RAID EU's• Maximum capacity per diskless StorEdge 5300 RAID EU: 0TB• Up to 16 StorEdge 5300 EU S's (8 per diskless StorEdge 5300 RAID EU)• Maximum capacity per StorEdge 5300 EU F: 5.6TB (raw)/4.004TB (useable)• Maximum capacity: 89.6TB (raw)/64.064TB (useable)

Power Supply

Equipped with two power supplies. Second is for redundancy	
Maximum DC output	500W

Environment

AC Power	100-240 V, 50/60Hz, 500W
----------	--------------------------

Operating Temperature	10° C to 35° C (50° F to 95° F) (<i>stand-alone server</i>) 10° C to 30° C (50° F to 86° F) (<i>in a rack-environment at sea level</i>) * Subtract 1° C from the maximum temperature for every 500-meter increase in altitude. (Tested to 3000m at 24° C [75.2° F]) 10% to 90% relative humidity, at 27° C max, wet bulb (noncondensing)
Non-operating Temperature	-40° C to 65° C (-40° F to 149° F), up to 93% relative humidity, noncondensing
Altitude (operating)	10° C to 30° C (50° F to 86° F) (<i>stand-alone server for all altitudes</i>) 10° C to 25° C (50° F to 77° F) (<i>in a rack-environment at sea level</i>) * Subtract 1° C from the maximum temperature for every 500-meter increase in altitude. (Tested to 3000m at 19° C [66.2° F])
Altitude (non-operating)	Up to 12000 m

Regulations

Meets or exceeds the following requirements:	
Safety	IEC60950, UL/CSA60950, EN60950
RFI/EMI	FCC Class A, Part 15 47 CFR, EN55022, CISPR 22
Immunity	EN55024
Certifications Safety EMC	cULus Mark, TUV GS Mark, CE Mark CE Mark (93/68/EEC), FCC authorized Class A, VCCI, BSMI, CTICK

Dimensions and Weight

Chassis <ul style="list-style-type: none"> • Height • Width • Depth • Weight 	StorEdge 5310 NAS 3.50 in. (87 mm) 16.93 in. (430 mm) 25.51 in. (648 mm) 46 lb. (20.9 kg)
Enclosure	Supported in 72" StorEdge cabinet and SunRacks

Software/Operating System

Feature	Specifications
Operating System	Storage-optimized Operating System
File System	<ul style="list-style-type: none"> • 64-bit Journaling File system • Up to 16TB per file system; up to 512 file systems • Sun StorEdge File Checkpoint "checkpoints" • Online extensible/expandable capacity • NTFS streams support • Quotas: user, group and directory tree • Unicode

Feature	Specifications
File Access Protocols	CIFS/SMB, NetBIOS, NFS v2 and v3, FTP
Directory and Name Services	AD (LDAP, Kerberos v5), NT 4.0 Multiple Master Domains (MMD), DNS, WINS, NIS, NIS+, Local files
Network Services	<ul style="list-style-type: none"> • DHCP • NTP • RDATE • Network Status Monitor
Access Control	<ul style="list-style-type: none"> • NT4.0 Access Control Lists and Access Tokens • UNIX credentials and permissions • Credential mapping • Directory tree quotas • User and Group Quotas • Sun StorEdge File Checkpoint “checkpoints” • SMB/CIFS Shares • UNIX exports • Approve lists, trusted host lists, netgroups • Cross-platform file locking (limitations) • Cross-platform authentication
Supported clients See the StorEdge 5x10 WWW for most current support information – SunWIN #414203	A client is any computer on the network that requests file services from the StorEdge 5310 NAS. In general, if a client implementation follows the NFS version 2 or 3 protocol or the CIFS specifications, it is supported with the StorEdge 5310 NAS Appliance.
Remote Management	<ul style="list-style-type: none"> • SNMP • Remote Syslog
System Administration	<ul style="list-style-type: none"> • Web (HTTP/Java based) GUI • Telnet • Rlogin • Rsh • SSH • Console Command Line Interface (CLI)
Mail Services	SMTP
System Monitoring	<ul style="list-style-type: none"> • Disk subsystem via SCSI Enclosure Services (SES) • Fans, power supplies, temperature, voltages via IPMI protocol • UPS • Network monitor
Tape Backup	<ul style="list-style-type: none"> • NDMP v2 and v3 • Backup with full Unicode • “Remote backup” and direct attach/3-way configurations supported w/NDMP
Setup and Configuration	Web-based GUI

StorEdge 5300 RAID Expansion Unit & 5300 Expansion Unit F Specifications

Specifications for StorEdge 5300 RAID EU and 5300 EU with 14 hard disk drives

Feature	Specification	
	StorEdge 5300 RAID EU	StorEdge 5300 EU F & S
Physical Planning		
Dimensions	5.21 inches high (13.2 cm) 17.6 inches wide (44.5 cm) 22.1 inches deep (56.1 cm)	5.1 in (12.9 cm) 17.6 in (44.7 cm) 22.1 in (56.1 cm)
Maximum Weight	95 pounds (43 kg)	91 pounds (41 kg)
Environmental (operating)		
Temperature	5° to 40° C (41° to 104° F)	
Relative Humidity	10% to 90% noncondensing	
Altitude	100 ft (30.5 M) below sea level to 9,840 feet (3,000 meters)	
Shock	5.0 g, 11 ms half sinewave	
Vibration	0.20 G, 5 to 500 Hz sinusoidal	
Heat Output	390 Watts (1331 BTU/Hr)	
Environmental (nonoperating)		
Temperature (Storage)	-10° to 50° C (-14° to 120° F)	
Temperature (Transit)	-40° to 60° C (-40° to 140° F)	
Humidity (Storage)	10% - 90%, Max Dew Point is 26° C (79° F), 10% per hour Gradient	
Humidity (Transit)	5% - 95%, Max Dew Point is 26° C (79° F), 10% per hour Gradient	
Altitude	100 ft (30.5 M) below sea level to 40,000 feet (12,000 meters)	
Shock	30 G, 11 msec half sine wave in the side/side & up/down directions 5 G, 11 msec half sine wave in the front/back direction	
Vibration	1.0 G, 5 to 500 Hz sinusoidal	
Power Requirements		
AC Power (RAID EU)	2.9 A Max Operating @ 100 VAC (90 – 136 VAC Range), 50/60 Hz 1.5 A Max Operating @ 240 VAC (198 – 264 VAC Range), 50/60 Hz	
AC Power (EU)	3.2 A Max Operating @ 100 VAC (90 – 136 VAC Range), 50/60 Hz 1.4 A Max Operating @ 240 VAC (198 – 264 VAC Range), 50/60 Hz	
Standards Compliance		
Safety and Emissions	IEC 60950, EN 60950, UL 60950, UL 1950, FCC Part 15 (47CRF15B), CISPR 22 (EN55022 — RF Radiated and Conducted Emissions), IEC 61000-3-2, IEC 61000-3-3	
Immunity	CISPR 24 (EN55024), IEC 61000-4-2, IEC 61000-4-3, IEC 61000-4-4, IEC 61000-4-5, IEC 61000-4-6, IEC 61000-4-8, IEC 61000-4-11	

Requirements and Configuration

Cabling Diagrams

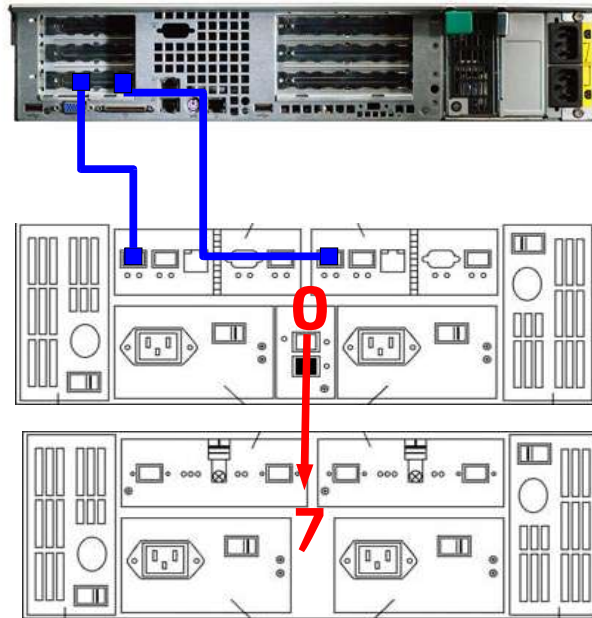


Figure 20: 1 HBA; 1 StorEdge 5300 RAID EU

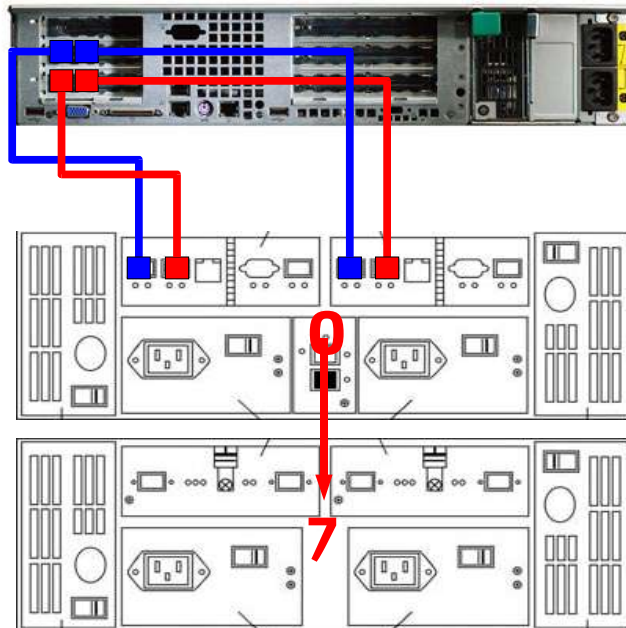


Figure 21: 2 HBA's; 1 StorEdge 5300 RAID EU

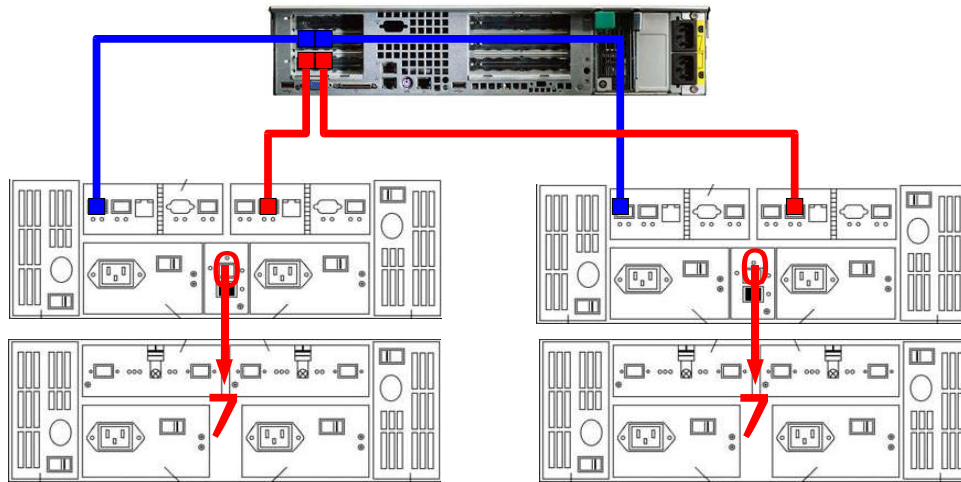


Figure 22: 2 HBA's; 2 StorEdge 5300 RAID EU's




Figure 23. StorEdge File Cluster; 2 heads sharing 1 StorEdge 5300 RAID EU
Heartbeat Connection




Figure 24. StorEdge file Cluster; 2 heads with 2 StorEdge 5300 RAID EU's

The StorEdge 5300 RAID EU's are configured with a “top-down/bottom-up” cabling scheme. This ensures data access to surviving drive modules even in the unlikely event that a full drive module is unavailable.

All StorEdge 5300 RAID EU's optical SFP and I/O Modules copper interconnect connections are designed with port bypass technology. Port bypass technology maintains fibre loop integrity when new Fibre Channel connections are made – or if existing ones are removed – by automatically opening and closing ports as needed. This means full drive modules can be added or removed without interrupting access to data.



Client Operating System Support

A client is any computer on the network that requests file services from the StorEdge 5310 NAS. In general, if a client implementation follows the NFS version 2 or 3 protocol or the CIFS specifications, it is supported with the StorEdge 5310 NAS Appliance.

Software Support Matrix

Category	Vendor	Application	Operating System
Database	Oracle	10g	Solaris 9
Productivity	Sun	StarOffice 7 and later	Solaris 8 and later, Windows NT 4.0 (with SP 6), 2000 (with SP2 or latest), 2003 (with latest SP), XP (with latest SP)
	Microsoft	Office 2000 and later	Windows NT 4.0 (with SP 6), 2000 (with SP2 or latest), 2003 (with latest SP), XP (with latest SP)
	Adobe	Acrobat Suite	Solaris 8 and later, Windows NT 4.0 (with SP 6), 2000 (with SP2 or latest), 2003 (with latest SP), XP (with latest SP)
Backup	VERITAS ¹	NetBackup™ Server 5.x with NDMP option	Solaris 8 and later
	Legato	NetWorker v7.2 w/NDMP	StorEdge NAS oS 4.05 and later

1. Veritas NetBackup 5.0 with NDMP option only supported in a "remote configuration". Support for other backup configurations and backup applications is scheduled for a future announcement. For additional information on Veritas NetBackup with NDMP see the "Veritas NetBackup for NDMP System Admin Guide" on <http://support.veritas.com>

Peripherals Support Matrix

Category	Description
Cables	Used for connecting the StorEdge 5300 RAID EU to the StorEdge 5310 NAS Appliance. All cables are FC. The StorEdge 5310 NAS Appliance ships with 5m LC to LC connector FC cables (2 cables with 1 FC HBA config's; 4 cables with 2 FC HBA config's). If longer cables are required, order the following part numbers: <ul style="list-style-type: none"> X9732A: 2M LC to LC FC Optical Cable X9734A: 15M LC to LC FC Optical Cable X9736A: 25M LC to LC FC Optical Cable X9738A: 50M LC to LC FC Optical Cable
Rack	Sun StorEdge cabinet SG-XARY030A and Sun Racks 900 & 1000
Rackmount Kit	StorEdge 5310 NAS: XTA-5200-2URK-19U
	StorEdge 5300 RAID Expansion Unit: XTACSM1-RK-3RU-19U
	StorEdge 5300 Expansion Unit: XTACSM1-RK-3RU-19U

Category	Description
Power cords	Power cord appropriate for installation requirements to be ordered as an x-option.

Expansion

The StorEdge 5310 NAS supports the following expansion feature:

Fibre Channel

- Up to 2 StorEdge 5300 RAID Expansion Units with up to 7 StorEdge 5300 EU F's on each of the StorEdge 5300 RAID EU's.
- 4 StorEdge 5300 RAID EU's utilizing the "Get To Yes" program with a 60 day lead-time

S-ATA

- Up to 2 diskless StorEdge 5300 RAID Expansion Units with up to 8 StorEdge 5300 EU S's on each of the diskless StorEdge 5300 RAID EU's.
- 4 diskless StorEdge 5300 RAID EU's utilizing the "Get To Yes" program with a 60 day lead-time

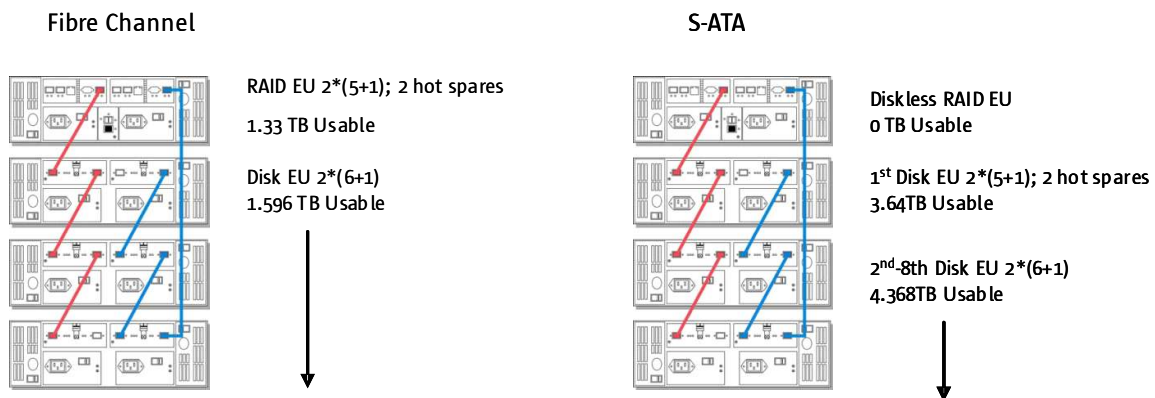
Mixing Fibre Channel and S-ATA

Mixing of Fibre Channel and S-ATA trays on a SE5310 NAS system is supported when the FC trays and S-ATA trays are kept separate. Meaning that the FC trays are comprise one array system and the S-ATA trays comprise another.

Please consult the "What Works With What" (SunWIN/Partner Document Library #414203) for the most current information on support.

Calculating Usable Capacity

2 RAID groups in each EU. Note that the 1st Disk EU for S-ATA has the 2 hot spares as the RAID EU is diskless.



Ordering Information

Ordering Notes

- All StorEdge 5310 NAS and StorEdge 5310 Expansion Units ship with documentation on CD. There is no hardcopy documentation with the exception of the “Quick Start Guide” and the “Safety and Compliance Guide”.
- The StorEdge 5310 NAS and StorEdge 5310 Expansion Units ship with dual power supplies standard. Appropriate power cords are to be ordered.
- The StorEdge 5310 NAS and StorEdge 5300 Expansion Units are rack ready and do not include rack mount kits – to be ordered separately.
 - XTA-5200-2URK-19U for the StorEdge 5310 NAS
 - XTACSM1-RK-3RU-19U for the StorEdge 5300 RAID Expansion Unit and StorEdge 5300 Expansion Unit
- If rack mounting, each unit requires a rack mount kit. Examples:
 - XTA5310R21A1MU2044 requires:
 - 1 XTA-5200-2URK-19U for the StorEdge 5310 NAS head
 - 1 XTACSM1-RK-3RU-19U for the StorEdge 5300 RAID EU
 - XTA5310R32A2MU4088 requires:
 - 1 XTA-5200-2URK-19U for the StorEdge 5310 NAS head
 - 2 XTACSM1-RK-3RU-19U for the 2 StorEdge 5300 RAID EU's
- One StorEdge File Replicator license required for each StorEdge NAS Appliance used for data replication. This means at least 2 StorEdge File Replicator licenses must be ordered to implement data replication. See the Q&A section for more detail.
- StorEdge 5310C: The part numbers contain 2 StorEdge 5310 NAS heads with 2 dual-port FC HBA's and cables. One configuration has 2-port MMF NIC's and the other single-port 10/100 UTP NIC's:
 1. XTA5310R20A4HSCK: Two StorEdge 5310 NAS heads (no storage) each with 2 dual-port FC HBA's; 4x5m FC cables; 1 single-port 10/100 Ethernet NIC
 2. XTA5310R20A4MSCK: Two StorEdge 5310 NAS heads (no storage) each with 2 dual-port FC HBA's; 4x5m FC cables; 1 dual-port 10/100/100 MMF NIC

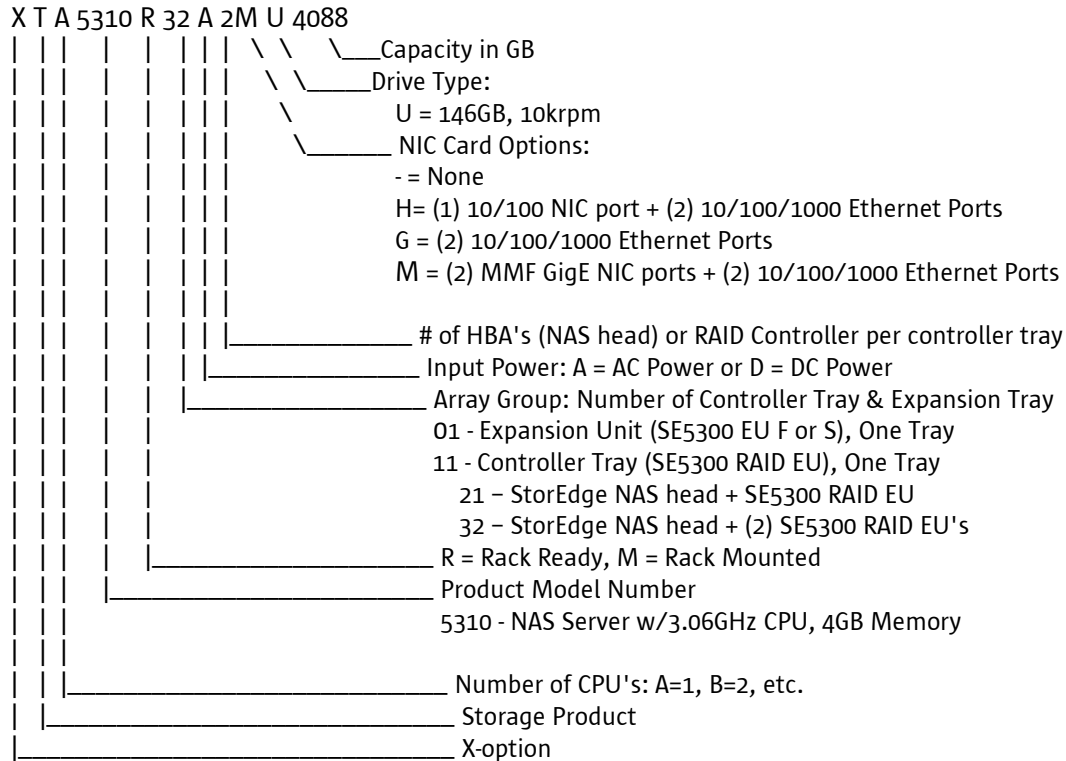
When configuring a StorEdge 5310C NAS, 1 or 2 StorEdge 5300 RAID EU's must be ordered as well. Both Fibre Channel and S-ATA RAID EU's and Disk EU's are supported – normal configuration rules apply. For Fibre Channel, this is the 14 drive StorEdge 5300 RAID EU (XTA5300R11A2-U2044). 7 drive RAID EU's EU F's are not supported with the StorEdge 5310C. When configuring StorEdge 5310C NAS with 2 StorEdge 5300 RAID EU's, it is recommended that both 5300 RAID EU's be configured identically (same number of EU's, disks, etc.).

- All StorEdge 5300 pricing is category “P”.

Part Number Format Description

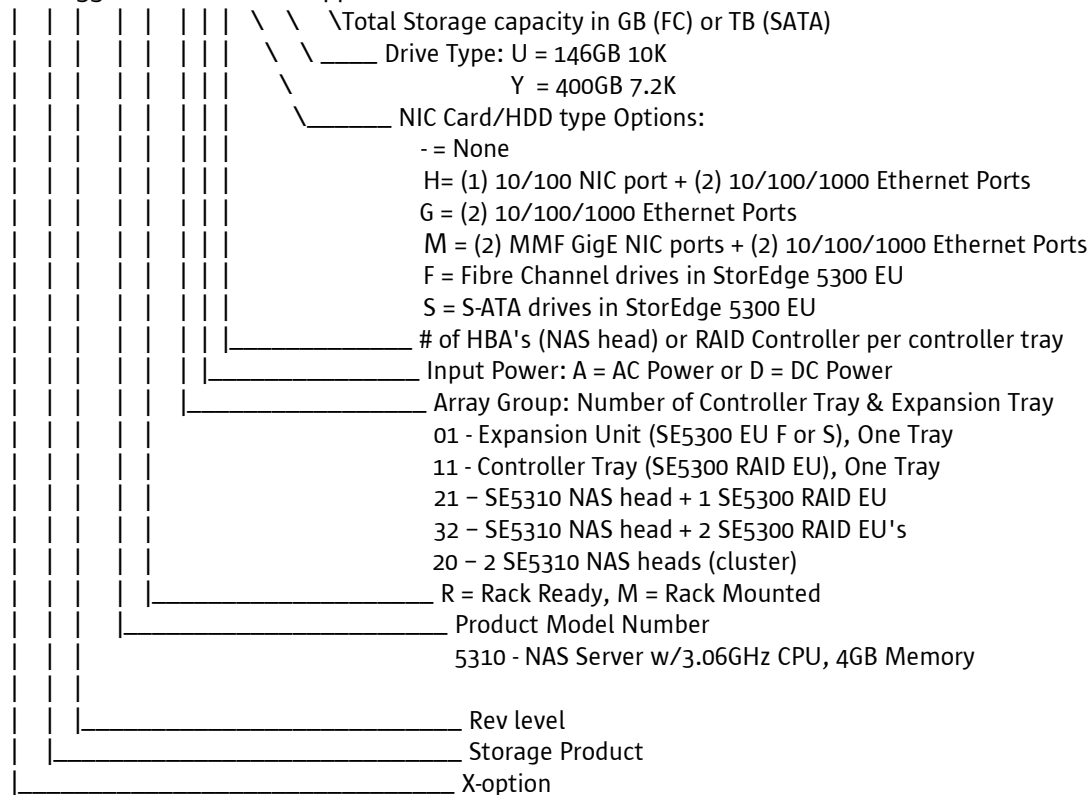
The basic system configuration can be determined from the components of each part number. Here is an example part number, with each component explained.

Example: XTA5310R32A2MU4088



Example: XTA5300R11A2-U2044 StorEdge 5300 RAID Expansion Unit

X T A 5300 R 11 A 2 - U2044



Sun StorEdge 5310 NAS Part Numbers

Part Number	Description
XTA5310R21A1MU2044	Sun StorEdge 5310 NAS Appliance: Rack Ready, 1x3.06GHz CPU, 4GB Memory (4x1GB DIMMs), 2x10/100/1000 Ethernet Ports, 1x2-port 10/100/1000 optical Ethernet NIC, 1x2-port 2Gb/s FC HBA, 2x5m fiber optic cables, 2-Internal AC Power Supplies; 1xStorEdge 5300 RAID EU: Rack-Ready, 2044GB, 14x146GB 10K FC HDDs, 2x1GB-cache memory FC RAID Controllers, 2xhot-swap power supplies, 2xhot-swap cooling fans, 2x copper FC ports, and 4xFC host ports with shortwave SFPs, 2x 6m Ethernet cables; 2x2m copper cables
XTA5310R21A1-U2044	Sun StorEdge 5310 NAS Appliance: Rack Ready, 1x3.06GHz CPU, 4GB Memory (4x1GB DIMMs), 2x10/100/1000 Ethernet Ports, 1x2-port 2Gb/s FC HBA, 2x5m fiber optic cables, 2-Internal AC Power Supplies; 1xStorEdge 5300 RAID Expansion Unit: Rack-Ready, 2044GB, 14x146GB 10K FC HDDs, 2x1GB-cache memory FC RAID Controllers, 2xhot-swap power supplies, 2xhot-swap cooling fans, 2x copper FC ports, and 4xFC host ports with shortwave SFPs, 2x 6m Ethernet cables; 2x2m copper cables

<i>Part Number</i>	<i>Description</i>
XTA5310R21A2MU2044	Sun StorEdge 5310 NAS Appliance: Rack Ready, 1x3.06GHz CPU, 4GB Memory (4x1GB DIMMs), 2x10/100/1000 Ethernet Ports, 1x2-port 10/100/1000 optical Ethernet NIC, 2x2-port 2Gb/s FC HBAs, 4x5m fiber optic cables, 2-Internal AC Power Supplies; 1xStorEdge 5300 RAID EU: Rack-Ready, 2044GB, 14x146GB 10K FC HDDs, 2x 1GB-cache memory FC RAID Controllers, 2xhot-swap power supplies, 2xhot-swap cooling fans, 2x copper FC ports, 4xFC host ports with shortwave SFPs, 2x 6m Ethernet cables; 2x2m copper cables
XTA5310R21A2-U2044	Sun StorEdge 5310 NAS Appliance: Rack Ready, 1x3.06GHz CPU, 4GB Memory (4x1GB DIMMs), 2x10/100/1000 Ethernet Ports, 2x2-port 2Gb/s FC HBAs, 4x5m fiber optic cables, 2-Internal AC Power Supplies; 1xStorEdge 5300 RAID EU: Rack-Ready, 2044GB, 14x146GB 10Krpm FC HDDs, 2x 1GB-cache memory FC RAID Controllers, 2xhot-swap power supplies, 2xhot-swap cooling fans, 2x copper FC ports, 4xFC host ports with shortwave SFPs, 2x 6m Ethernet cables; 2x2m copper cables
XTA5310R32A2MU4088	Sun StorEdge 5310 NAS Appliance: Rack Ready, 1x3.06GHz CPU, 4GB Memory (4x1GB DIMMs), 2x10/100/1000 Ethernet Ports, 1x2-port 10/100/1000 optical Ethernet NIC, 2x2-port 2Gb/s FC HBAs, 4x5m fiber optic cables, 2-Internal AC Power Supplies; 2xStorEdge 5300 RAID EU's: Rack-Ready, 4088GB, 28x146GB 10K FC HDDs, 4x1GB-cache memory FC RAID Controllers, 4xhot-swap power supplies, 4xhot-swap cooling fans, 4x copper FC ports, 8xFC host ports with shortwave SFPs, 4x 6m Ethernet cables; 4x2m copper cables
XTA5310R32A2-U4088	Sun StorEdge 5310 NAS Appliance: Rack Ready, 1x3.06GHz CPU, 4GB Memory (4x1GB DIMMs), 2x10/100/1000 Ethernet Ports, 2x2-port 2Gb/s FC HBAs, 4x5m fiber optic cables, 2-Internal AC Power Supplies; 2xStorEdge 5300 EU's: Rack-Ready, 4088GB, 28x146GB 10K FC HDDs, 4x1GB-cache memory FC RAID Controllers, 4xhot-swap power supplies, 4xhot-swap cooling fans, 4x copper FC ports, 8xFC host ports with shortwave SFPs, 4x 6m Ethernet cables; 4x2m copper cables
XTA5300R01A0FU2044	Sun StorEdge 5300 Expansion Unit F: Rack-Ready, 2044GB, 14x146GB 10Krpm 10Krpm FC HDDs, 2xFC I/O Modules, 2xhot-swap redundant power supplies, 2xhot-swap redundant cooling fans, 2xcopper FC ports (for expansion trays), 2x2m copper cables
XTA5300R11A2-U2044	Sun StorEdge 5300 RAID Expansion Unit: Rack-Ready, 2044GB, 14x146GB 10Krpm FC HDDs, 2x 1GB-cache memory FC RAID Controllers, 2xhot-swap redundant power supplies, 2xhot-swap redundant cooling fans, 2x copper FC ports (for expansion trays), and 4xFC host ports with shortwave SFPs, 2x 6m Ethernet cables
XTA5310R21A1-Y56	Sun StorEdge 5310 NAS:Rack Ready, 3.06GHz CPU,4GB Memory, 2*10/100/1000 Ethernet Ports, 2-port 2Gb/s FC HBA, 2*5m fiber optic cables, 2*AC power supplies;SE5300 RAID EU: 0GB, 2*FC RAID Controllers, 2*hot-swap power supplies, 2*hot-swap cooling fans, 2*copper FC ports,4*FC ports with shortwave SFPs, 2* 6m Ethernet cables; 4*2m copper cables;SE5300 EU S:5.6TB, 14*400GB 7.2K SATA Drives, 2*SATA I/O Modules,2 redundant power supplies,2 cooling fans,2 copper FC ports

<i>Part Number</i>	<i>Description</i>
XTA5310R21A1MY56	Sun StorEdge 5310 NAS:Rack Ready, 3.06GHz CPU,4GB Memory, 2*10/100/1000 Ethernet Ports,2-port 10/100/1000 MMF NIC,2-port 2Gb/s FC HBA,2*5m fiber optic cables, 2*AC power supplies; SE5300 RAID EU: 0GB, 2*FC RAID Controllers, 2*hot-swap power supplies, 2*hot-swap,cooling fans,2*copper FC ports,4*FC ports w/ shortwave SFPs, 2*6m Ethernet cables; 4*2m copper cables; SE5300 EU S:5.6TB, 14*400GB 7.2K SATA Drives,2*SATA I/O Mods,2 redundant power supplies,2 cooling fans,2 copper FC ports
XTA5310R21A2-Y56	Sun StorEdge 5310 NAS:Rack Ready, 3.06GHz CPU,4GB Memory, 2*10/100/1000 Ethernet Ports, 2*2-port 2Gb/s FC HBA, 2*5m fiber optic cables,2*AC power supplies;SE5300 RAID EU:0GB, 2*FC RAID Controllers, 2*hot-swap power supplies, 2*hot-swap cooling fans, 2*copper FC ports,4*FC ports with shortwave SFPs, 2* 6m Ethernet cables; 4*2m copper cables;SE5300 EU S:5.6TB, 14*400GB 7.2K SATA Drives, 2*SATA I/O Modules,2 redundant power supplies,2 cooling fans,2 copper FC ports
XTA5310R21A2MY56	Sun StorEdge 5310 NAS:Rack Ready, 3.06GHz CPU,4GB Memory, 2*10/100/1000 Ethernet Ports,2-port 10/100/1000 MMF NIC, 2*2-port 2Gb/s FC HBA,4*5m fiber optic cables, 2*AC power supplies; SE5300 RAID EU: 0GB, 2*FC RAID Controllers, 2*hot-swap power supplies, 2*hot-swap cooling fans, 2*copper FC ports,4*FC ports w/ shortwave SFPs, 2*6m Ethernet cables; 4*2m copper cables;SE5300 EU S:5.6TB, 14*400GB 7.2K SATA Drives,2*SATA I/O Mod,2 redundant power supplies, 2 cooling fans,2 copper FC ports
XTA5300R12A2SY56	Sun StorEdge 5300 RAID Expansion Unit: Rack Ready, 0GB, diskless, 2x1GB-cache memory FC RAID Controller cards, 2 redundant power supplies, 2 redundant cooling fans, 2 copper FC ports for expansion and 4 host ports with short-wave SFPs, 2x6m Ethernet cables; StorEdge 5300 Expansion Unit S: Rack-Ready, 5.6TB , 14x400GB 7.2Krpm SATA Drives, 2xSATA I/O Modules, 2 redundant power supplies, 2 redundant cooling fans, 2 copper FC ports for expansion, 2x2m copper cables
XTA5300R01A0SY56	Sun StorEdge 5300 Expansion Unit S: Rack-Ready, 5600GB , 14x400GB 7.2Krpm SATA Drives, 2xSATA I/O Modules, 2 redundant power supplies, 2 redundant cooling fans, 2 copper FC ports for expansion, 2x2m copper cables (Standard Configuration)
SEC1B-LCO-9929	StorEdge Compliance (SECompliance) Base License (0 to 32TB raw) for Sun StorEdge Compliance Archiving Software. Software is for use with Sun StorEdge NAS 5310 Appliance with SATA drives
SEC2B-LCO-9929	StorEdge Compliance (SECompliance) Expansion License (add-on 32TB increments) for Sun StorEdge Compliance Archiving Software. Software is for use with Sun StorEdge NAS 5310 Appliance with SATA drives
SFR2A-LCO-E929	Sun StorEdge File Replicator Software (remote replication); License Only; 1 per StorEdge 5300 NAS Appliance
SFR2A-LCO-ED29	Sun StorEdge File Replicator Software (remote replication); License Only; 1 per StorEdge 5300 NAS filer; 90-day eval license
SEC1A-LCO-9929	StorEdge Compliance (SECompliance) Base License (0 to 32TB raw) for Sun StorEdge Compliance Archiving Software. Software is for use with Sun StorEdge NAS 5310 Appliance with Fibre Channel drives
SEC2A-LCO-9929	StorEdge Compliance (SECompliance) Expansion License (over 32TB to 64TB raw) for Sun StorEdge Compliance Archiving Software. Software is for use with Sun StorEdge NAS 5310 Appliance with Fibre Channel drives



<i>Part Number</i>	<i>Description</i>
XTA5310R20A4HSCK	Sun StorEdge 5310C NAS; 2 StorEdge 5310 NAS Rack Ready Appliances: 3.06GHz CPU, 4GB Memory, 2*10/100/1000 Ethernet Ports, 1-port 10/100 Ethernet NIC, 2*2-port 2Gb/s FC HBA, 2 internal AC Power Supplies, 8*5m fiber optic cables, Ethernet cable
XTA5310R20A4MSCK	Sun StorEdge 5310C NAS: 2 StorEdge 5310 NAS Rack Ready Appliances: 3.06GHz CPU, 4GB Memory, 2*10/100/1000 Ethernet Ports, 2-port 10/100/1000 MMF Ethernet NIC, 2*2-port 2Gb/s FC HBA, 2 internal AC Power Supplies, 8*5m fiber optic cables, Ethernet cable
XTA5300-HF2-2G	Host Bus Adapter: PCI Dual-port 2Gb/s FC for StorEdge 5300 NAS Appliance

Warranty, Service and Support

Warranty

The Sun StorEdge 5310 NAS Appliance comes with a two year warranty. Year one provides next business day on-site support and year two provides 15 day Parts Return support.

Sun Spectrum Instant Upgrades

Warranty coverage can be upgraded to standard Silver, Gold, 7x24 Gold or Platinum coverage utilizing Sun Spectrum Instant Upgrades. Sun Spectrum Instant Upgrades allow you to obtain the coverage that best supports your business environment through a variety of service levels including phone and on-site response times, escalation options, event monitoring, account management and skills assessments. Sun Spectrum Instant Upgrades also allow you to maintain consistent support levels in your IT environment with the ability to maintain all your Sun products at the same service level.

Enterprise Installation Service

Installation and configuration of the Sun StorEdge 5310 (scheduled for availability Q1CY05) including installation planning, site audit, acceptance testing and system turnover. The scope of the installation for the SE5310 is expected to get the SE5310 to a basic level of functionality including:

- Planning the Installation
- Installing the SE5310 NAS Server and SE5300 Storage Expansion Unit
- Powering on the System Locally
- Connecting the Grounding Cable
- Connecting the Power Cables
- Establishing a Serial Connection
- Establishing Initial Configuration for one LUN
- Setting and Verifying Default Configuration

Knowledge Services

EA-111 Sun Spectrum Education Account

Help your customers stay ahead of constantly changing technologies by investing in their enterprise's most important asset -- people.

Thousands of students each year are trained by Sun and its authorized centers through Web-based courses and at training sites located in more than 60 countries.

Sun Education's historical business manages a portfolio of IT curriculum focused on Sun technology sold and delivered to Sun external audiences in classroom, CD-ROM, and web formats.

Professional Services

Storage Migration Service - This service can help customers safely transfer data from one storage system to another storage system, without pulling their internal resources from other critical business tasks.

Sun Services can offer multiple levels of consulting services to ensure a smooth data migration from the existing storage to the SE5310. Sun Services offers a wide range of data management and migration services to accommodate most customer environments and circumstances. Contact local Sun Services for further details and quotation for the specific customer environment.

External - <http://www.sun.com/service/storage/>

Internal -

<http://mymarketing.central.sun.com:81/NASApp/myMarketing/Portal/NonLeaf?FunctionalCatId=1656&CategoryId=2916>

Glossary

1U	One rack unit as defined by the Electronic Industries Alliances (EIA). A vertical measurement equal to 1.75 inches.
2U	Two rack units; a vertical measurement equal to 3.5 inches.
AC	Alternating Current.
Access Control	Limits user access to resources on a computer network, most commonly by requiring a user name and password. Usually a single logon is sufficient to authenticate, or verify, a user.
Access Rights	Permissions granted to user accounts to allow access to such system resources as filesystems, applications, and directories. For example, read-only access allows a user to open or list a file without being able to make changes to the file. Users who are granted access rights to a directory usually have the same access rights to all subdirectories.
AD	Short for Active Directory Service. AD is a Windows 2000 namespace that is integrated with the Domain Name System (DNS). AD stores domain information such as users, groups, and shared resources and makes that information available to Active Directory clients.
Alias IP Address	Multiple IP addresses assigned to a single port, in addition to the primary IP address. All IP aliases for a port must be on the same physical network and share the same netmask and broadcast address as the first or primary IP address.
Authentication	The process of validating that the user attempting to logon is truly the owner of the account.
Autohome Shares	Temporary SMB/CIFS shares that are created when a user logs on to the system and removed when the user logs off.
Block or Block Size	Also known as stripe unit size, block size is often used to describe the amount of data sent or received by the host per I/O operation. Also used to describe the size of an atomic read/write operation to/from a disk as well as the size of the data unit being "striped" across disks.
Boot Up	The process of starting a computer. Booting up involves checking all hardware components, initializing system components, and loading the operating system.
Browser	Software used for access to information on the World Wide Web. Microsoft Internet Explorer and Netscape Navigator are examples of browsers. See also Web Browser.
Cache	Cache is the DRAM based staging area used to provide higher performance to applications for reads and writes. During reads, the RAID controller tries to keep the latest and most often accessed data in the cache and also tries to pre-stage the cache with future data during sequential accesses. For writes, cache is used to provide delayed writes to the drives. This delays the parity calculations and drive writes during RAID 5 operations. More optimization and advanced staging algorithms thus provide better performance.
Checkpoint	A point-in-time image of a StorEdge NAS filesystem taken at the volume level. While the active file volume can be modified with read/write operations, a virtual volume produced at the time of checkpoint creation remains available in a static, read-only state.
CIFS	Stands for Common Internet File System. An enhanced version of the SMB file-sharing protocol that allows groups of users to work together and share documents over the Internet in the same way as in local area networks.
Density	Number of units in a given amount of space.
DHCP	Stands for Dynamic Host Configuration Protocol. DHCP provides a mechanism by which a computer can acquire an IP address automatically when it connects to the network. DHCP allows more flexible and efficient use of network resources than static IP addresses.

DNS	Short for Domain Name System. A network service that translates domain names into IP addresses. If you have multiple DNS servers on your network, and one DNS server can't resolve host names, it asks another one, and so on, until the IP address is found. See also Domain Name System.
Domain	A group of computers and devices on a network that are administered as a unit with common rules and procedures.
Domain Name	A name that identifies a domain. See also Domain. The domain name can be the company name, division name, facility name, department name, or other descriptive name.
Domain Name System	The network server that maintains the list of all host names in a domain. StorEdge NAS OS uses the name server to translate domain names to the corresponding IP address. See also DNS.
DTQ	Stands for Directory Tree Quota. A directory tree quota is a quota, or limit, to the space or the number of files that a directory tree (a directory and its subdirectories) can occupy.
Ecache	External cache. Memory cache external to the CPU chip, also referred to as L2 cache.
ECC	Error Correcting Code. A type of memory that corrects errors on the fly.
Error Correction Code	Extra bits added to Words, or Double Words, that correct all single-bit errors, and detect all double-bit errors. A superior technology to parity, which detects, but does not correct, single-bit errors, and cannot detect double-bit errors.
Ethernet	A network communication system developed and standardized by DEC, Intel, and Xerox using baseband transmission, CSMA/CD access, logical bus topology, and coaxial cable. The successor IEEE 802.3 standard provides for integration into the OSI model and extends the physical layer and media with repeaters and implementations that operate on fiber optics, broadband, and unshielded twisted pair.
Ethernet 10/100/1000BaseT	The most widely used LAN access method defined by the IEEE 802.3 standard; uses standard RJ-45 connectors and telephone wire. 100BaseT is also referred to as Fast Ethernet. 1000BaseT is also referred to as Gigabit Ethernet.
Field Replaceable Unit	A component which can be removed and replaced during service in the field.
File Sharing	A feature that allows users of networked computers to make files available to other users.
File Volume	Filesystems created from partitions that have available space. If the file volume does not use up all the available space in a partition, the remaining space is automatically allocated into the next partition. See also Partition.
File Volume Extension	See Segment.
Filer	A Network Attached Storage device focused solely on file service and file storage
FRU	Field Replaceable Unit.
FTP	File Transfer Protocol. A client-server protocol which allows a user on one computer to transfer files to and from another computer over a TCP/IP network. Also the client program the user executes to transfer files. It is defined in STD 9, RFC 959.
General-purpose Server	A server designed to perform any type(s) of function(s). General-purpose servers typically require skilled IT professionals and system administrators to maintain them.
Gigabit Ethernet	An Ethernet standard that enables data transfer rates of up to 1 Gbps running over copper or optical fiber cable.
GUI	Stands for Graphical User Interface. A GUI uses graphical elements to present information to a computer user rather than the traditional text-only command line interface still found in telnet and similar implementations.
HBA	Host bus adapter.
Host ID	The unique identifier assigned to the host computer.
Hot-spare	A drive used by the RAID controller to replace a failed drive. Hot-spares are continuously powered up and spinning, but are not actually part of the array because they contain no data. This allows the array processor to have immediate access to a functioning drive for possible reconstruction of lost data when a disk fails.

Hot-swappable	A hot-swappable component can be installed or removed by simply pulling the component out and putting the new one in. The system automatically recognizes the component change and configures itself as necessary.
HTTP	Stands for Hyper-text Transmission Protocol. A protocol for exchanging HTML pages and forms.
I/O	Input/output. Transferring data between the CPU and any peripherals.
IP Address	A unique 32-bit value that identifies network hosts using TCP/IP. An IP address, or a block of addresses, is assigned upon application to organizations responsible for that function. No two network hosts can be assigned the same IP address. Each address consists of a network number, optional subnetwork number, and host number, written as four numbers separated by periods. Each number can be 0 to 255. See also <i>Address</i> and <i>URL</i> .
ISV	Independent software vendor.
Java	Java is a programming language developed by Sun Microsystems to be portable to any type of computing device. In practice, java allows web browsers to do much more than display information. Java scripts allow much more flexibility and functionality in web access and they run on virtually any type of computer.
Journaling File System	A fault-resilient filesystem in which data integrity is ensured because updates to directories and bitmaps are constantly written to a serial log on disk before the original disk log is updated. In the event of a system failure, a full journaling filesystem ensures that the data on the disk has been restored to its pre-crash configuration.
Kerberos Realm	A kerberos realm is a secured network requiring access through a key. (See also <i>KDC</i> .) Each system or user with a key can access any services or systems that the key opens. The user does not have to enter a user name and password each time he requests a controlled service.
L2 cache	See Ecache.
LAN	Stands for Local Area Network. A communications network that provides high-speed (over 1 Mbps) data transmission and is limited to a specific physical area (up to about six miles). The basic components of a LAN are: adapter boards installed in each computer to provide a cable connector, cabling, server hardware, and network management software.
LCD	Stands for Liquid Crystal Display. On the StorEdge NAS, the LCD is a two line display that shows basic information about system functions and, in conjunction with the control panel, allows you to perform certain system functions, like setting the IP address, directly on the unit, without access through the internet or intranet.
LED	Light emitting diode.
Login	Logging in is a security process designed to prevent access to system settings or other resources by those who should not have access. A login process usually requires a user name and password to verify, or authenticate, a user.
Master Domain Model	One of several types of domain models. In the Master Domain Model, an account domain is trusted by a resource domain.
Mirroring	Mirroring allows you to duplicate any or all of the file volumes of an active server onto a mirror server. In the event that the active server fails, the mirrored file volumes on the mirror server can be made available to network users within minutes.
MTBF	Mean Time Between Failures. The average time a component works without failure.
MTTR	Mean Time To Repair. The average time it takes to repair a component.
NetBIOS	NetBIOS is a protocol used for networking. NetBIOS was designed to support communications between symbolically named stations and the transfer of arbitrary data. NetBIOS manages the use of node names and transport layer connections for higher layer protocols such as SMB.
Network	A series of nodes such as terminals, computer systems, or other peripheral devices connected by a communications channel. See also LAN.
NIC	Stands for Network Interface Card. A NIC is an expansion card that provides access to a network.

NIS	Short for Network Information Service. Along with NFS, NIS provides a distributed database system to centralize (i.e., store one copy, on a single computer) common configuration files, such as the password file (/etc/passwd) and the hosts file (/etc/hosts).
NIS+	Short for Network Information Service Plus (NIS+). NIS+ was designed to replace NIS. NIS+ can provide limited support to NIS clients, but was mainly designed to address problems that NIS cannot address.
NTP	Stands for Network Time Protocol. NTP provides a mechanism for synchronizing the time among a number of computers connected to a network.
Packet	A piece of a message transmitted over a network. Contains the destination address in addition to the data. Once all packets arrive at the destination, they are recompiled into the original message.
Partition	Sections on a LUN. Each partition can either have some space allocated to it, or can be empty. When a LUN is first created, all of the available space is located in the first partition, while the other partitions are empty. Each partition can have only one volume.
Port Bonding	Otherwise known as “channel bonding.” Port bonding allows you to scale network I/O by joining ports. This forms a single network channel of high bandwidth from two or more channels of lower bandwidth.
Protocol	A set of standards or rules that enable computers to connect to one another and exchange data. Using a protocol helps reduce the possibility of errors during data transmission.
Quota	A restriction on disk space or the number of files written to file volumes in the StorEdge NAS. This limit can be determined for a user or group (user or group quota) or for a directory (directory tree quota).
RAID	Redundant array of independent disks. A set of disk drives which appears as a single logical disk drive to the filesystem. Different RAID levels provide different capacity, performance, availability, and cost characteristics.
RAID 0	RAID level 0, or striping without parity or mirroring protection. Data is distributed evenly at the block level among disks for performance. No redundancy is provided, and the loss of a single disk causes the loss of data on all disks. Use this level for high-speed streaming of large file reads (for example, video) of non-critical data which is easily available elsewhere within the organization.
RAID 10	Also known in the industry as RAID 0+1. Implements block interleave data striping and mirroring. RAID 10 is not formally recognized by the RAID Advisory Board (RAB), but, it is an industry standard term. In RAID 10, data is striped across multiple disk drives, and then those drives are mirrored to another set of drives. The performance of RAID 10 is approximately the same as RAID 0 for sequential I/Os. RAID 10 provides an enhanced feature for disk mirroring that stripes data and copies the data across all the drives of the array. The first stripe is the data stripe; the second stripe is the mirror (copy) of the first data stripe, but it is shifted over one drive. Because the data is mirrored, the capacity of the logical drive is 50 percent of the physical capacity of the hard disk drives in the array.
RAID 5	RAID level 5, or striping with distributed parity. Both data and parity are distributed evenly across all the disks in the array at the block level. No single disk can compromise the integrity of the data. RAID 5 balances the optimization of performance, reliability and cost. Use this level for most applications which do not require the special characteristics of the above RAID levels.
RAM	Random Access Memory.
RDATE	RDATE is a time synchronization method that simply asks another computer on the network what the correct time is and resets itself accordingly. RDATE is not particularly accurate, but is adequate for most networks.
Read-ahead	Sequential data read from disk into cache without having actually been requested by the application host, in anticipation that it will be requested by the host. When the request occurs, it can be serviced as a low latency cache hit, thus improving host application performance.

SCSI	Small Computer Systems Interface. Pronounced “scuzzy.” An ANSI standard for controlling peripheral devices by one or more host computers.
Segment	Segments are available space that can be “attached” to a volume when the volume reaches its assigned capacity. This increases the volume’s total capacity. The segment, after being attached, becomes part of the volume and cannot be removed. Otherwise known as volume extensions.
Server	A network host that makes network resources, such as software applications and databases on hard disk or CD-ROM, available to network users. The server provides the centralized, multi-user functionality of the network application, such as data management, information sharing, network administration, or security.
Server Name	Identifies a network server. Server names are used in addition to IP addresses. This allows a server to be advertised on a network with a recognizable name. For example, the first StorEdge NAS server on a network could be identified as cdts0, the second as cdts1, and the third as cdts2 or they could be identified as Fred, Barney, and Wilma.
Shutdown	The multi-user operating system resident on the StorEdge NAS server must be shut down in an orderly sequence prior to turning the power off. The shutdown sequence closes files and terminates running programs to prevent loss or corruption of data.
SMB	Stands for Server Message Block. A Microsoft-compatible network protocol for exchanging files. SMB is typically used by Windows for Workgroups, OS/2 Warp Connect, and DEC Pathworks. See also CIFS.
SMTP	SMTP (Simple Mail Transfer Protocol) is a TCP/IP protocol used in sending and receiving e-mail.
SNMP	Stands for Simple Network Management Protocol. SNMP is primarily used for network monitoring and notification of network errors and other events. In the StorEdge NAS, SNMP also provides notification services through e-mail messages.
TCP/IP	Stands for Transmission Control Protocol/Internet Protocol. A commonly used networking protocol that allows interconnection of different network operating systems.
Telnet	A terminal emulation program for TCP/IP networks. The Telnet program runs on your computer and connects your PC to the StorEdge NAS server on the network. You can then enter commands through the Telnet program and they run as if you were entering them directly on the server console.
Transfer Rate	The rate at which data is transferred, usually measured in megabytes per second (MB/sec.).
Ultra320 SCSI	A standard for SCSI data transfers. It allows a transfer rate of up to 320 Mbytes/sec over a 16-bit SCSI bus.
Unicode	Unicode is a standard for representing letters that allows the language of computer messages and commands to be displayed in a variety of languages without rewriting the underlying programs.
URL	Stands for Uniform Resource Locator. An address system used by servers and clients to request documents. See also IP Address.
User Credentials	The information containing the user, account data, and the user’s group membership.
Volume	A volume is a virtual disk into which a filesystem, a DBMS, or an application can place data. A volume can be a single physical disk or a virtual disk mapped from one or more underlying extents. Applications that use volumes do not need to be aware of their underlying physical structure. Software handles the mapping of virtual partition addresses to physical addresses.
Web Browser	A web browser is a software application designed to search for and retrieve information from the Internet and the world-wide web.
WINS	Stands for Windows Internet Naming Service. A WINS server resolves NetBIOS names to IP addresses, allowing computers on a network to locate other NetBIOS devices more quickly and efficiently. WINS performs a similar function for Windows environments as DNS does for Unix environments.

Materials Abstract

Collateral	Description	Purpose	Distribution	Token # or COMAC Order #
References				
<ul style="list-style-type: none"> <i>Sun StorEdge™ 5310 NAS, Just The Facts</i> 	Reference Guide (this document)	Training Sales Tool	SunWIN, Reseller Web	425812
<ul style="list-style-type: none"> <i>Sun StorEdge 5310 NAS Sales Training Presentation</i> 	Sales/Training Presentation	Training Sales Tool	SunWIN, Reseller Web	426304
<ul style="list-style-type: none"> <i>Sun StorEdge 5310 NAS Intro's</i> 	Introductory E-mail	Training Sales Tool	SunWIN, Reseller Web	427205
<ul style="list-style-type: none"> <i>Sun StorEdge 5310 NAS Technical Customer Presentation</i> 	Customer Presentation	Training Sales Tool	SunWIN, Reseller Web	426305
<ul style="list-style-type: none"> <i>Sun StorEdge 5x10 NAS "What Works With What"</i> 	Compatibility/Support Guide		SunWIN, Reseller Web	414203
<ul style="list-style-type: none"> <i>Sun StorEdge 5310 NAS "How to Order"</i> 	Reference	Training Sales Tool	Sun WIN, Reseller Web	426615
<ul style="list-style-type: none"> <i>Sun StorEdge 5310 NAS Beat Sheet</i> 	Competitive	Training Sales Tool	Sun WIN	426619
Product Literature				
<ul style="list-style-type: none"> <i>Sun StorEdge 5310 NAS Data Sheet</i> 	Data Sheet	Training Sales Tool	SunWIN, Reseller Web	426022
External Web Site				
<ul style="list-style-type: none"> <i>Sun StorEdge 5310 NAS Web Site</i> 	http://www.sun.com/storage/5000/5300/5310/index.xml			
<ul style="list-style-type: none"> <i>Warranty terms, SunSpectrum, Installation, Software Support and Implementation information</i> 	http://www.sun.com/service/warrantiescontracts			
<ul style="list-style-type: none"> <i>User Guides</i> 	http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/nas/5310/index.html			
Internal Web Site				
<ul style="list-style-type: none"> <i>Sun StorEdge 5310 NAS Internal Site</i> 	http://webhome.ebay/networkstorage/products/5310/			
<ul style="list-style-type: none"> <i>Sun StorEdge NAS PTS Internal Site</i> 	http://pts-americas.west/nws/products/SE5210/			
<ul style="list-style-type: none"> 				
iForce Partner Web Site				
<ul style="list-style-type: none"> <i>Sun StorEdge 5310 NAS Partner Site</i> 	http://partner.sun.com/products/storage/5310.html			

Collateral	Description	Purpose	Distribution	Token # or COMAC Order #
<ul style="list-style-type: none"> <i>Sun StorEdge NAS Sharespace</i> 	http://sharespace.sun.com/gm/folder-1.11.62027			

Competitive Information

The Sun StorEdge 5310 NAS Appliance targets the “midrange” (higher end of \$25,000-\$99,000 and lower end of) NAS market space. Also in this space are Network Appliance, EMC, and HP:

	Sun StorEdge 5310 NAS	NetApp FAS940	EMC NS700
Capacity (raw)	32.7TB (FC) GA	12TB (FC)	21.9TB (FC)
Internal (Filer)	65.4TB (FC) 60 day lead-time		
	0TB	2TB	0TB
External	65.4TB (448x146GB)	10TB	21.9TB
HDD type	FC-AL (S-ATA Q3FY05)	FC-AL, S-ATA	FC-AL, S-ATA
RAM (Max)	4GB	3GB	4GB
RAID type (HW/SW)	Hardware	Software	Hardware
RAID Level	5	4	1, 5
Network ports, embedded	Dual 10/100/1000	Up to twelve 10/100/1000	Six 10/100/1000
Operating System	Sun StorEdge NAS	ONTAP	DART
File Access Protocols	NFS, CIFS, FTP	NFS, CIFS, HTTP, DAFS	NFS, CIFS, FTP
Filer Failover	Yes (Q3FY05)	Yes (FAS940C)	Yes (Requires 2 Data Movers)
File Sharing	UNIX, NT	UNIX, NT	UNIX, NT
Lock Type	CIFS & NFS	CIFS & NFS	CIFS & NFS
Authentication	ACL's user & group IDs, AD	ACL's user & group IDs, AD	ACL's user & group IDs, AD
Snapshots	16 per volume	255 per volume	?
Remote Copy	Yes	Yes	Yes
NDMP Support	Yes	Yes	Yes
Integrated Antivirus Detection	No	Yes	Yes
Quota management	Yes	Yes	Yes
Automatic File System Extension	Yes	Yes	Yes
Admin interface	Web GUI, CLI, Telnet, SSH	Web GUI, CLI, Telnet, SSH	Web GUI, CLI, Telnet, SSH
SNMP reporting	Yes	Yes	Yes
Performance (SPECsfs)	17,871	17,419	36,335 (Failover cluster)

Sun StorEdge 5310 NAS

Just the Facts



	Sun StorEdge 5310 NAS	NetApp FAS940	EMC NS700
Support and Warranty	2 yr. Limited warranty - Year 1 - NBD parts delivery Bus. Hrs/8 hr response phone - Year 2 - Return to Sun	3 yr. limited warranty: - Next Bus. Day parts delivery - 24x7 Phone - AutoSupport (phone home) - 90 day software (media)	3 yr. limited warranty: - Next Bus. Day parts delivery - 24x7 Phone - AutoSupport (phone home) - 90 day software (media)

Network Appliance

When comparing to NetApp it is important to know that NetApp's pricing model is a la carte. Drive enclosures, disk drives, NFS and CIFS licenses are all sold separately. In order to make an “apples to apples” comparison with NetApp filers, be sure to:

- Include the costs for the NFS and CIFS licenses as they are included with the OS of the Sun StorEdge NAS family.
- Capacities are similar. NetApp uses drive enclosures that house up to 14 drives. The head also, typically, houses up to 14 drives.

There are some questions that should be asked of customers either using or considering NetApp's NAS products:

- How does the NetApp system handle old “snapshots” vs. Sun StorEdge 5310 NAS when the volume becomes full? The StorEdge 5310 NAS automatically deletes the oldest “checkpoint”; avoiding system shutdown.
- Would you rather recover whole volumes or individual files? Do they really want millions of objects on a volume restored to a prior state? The StorEdge 5310 NAS enables the customer to recover individual files from a “snapshotted” state.
- How many “snapshots” do you need? NetApp supports up to 255 “snapshots” per volume. The Sun StorEdge 5310 NAS supports 16 snapshots per volume. How many “snapshots” would you like to manage? Do you understand the overhead imposed on the system with so many “snapshots”?
- With respect to mirroring, would you rather have all changes in your data replicated or only the most recent changes in a specified period? The purpose of mirroring is largely to ensure data availability with all changes to files being communicated. The Sun StorEdge File Replicator will send **all** changes to the customer's data that occur within a scheduled period.
- Journaling helps to ensure file system consistency and fast recovery in the unlikely event of a system crash. Would you prefer to have the journal on RAID protected disk or on unprotected NVRAM (single point of failure)?

Q & A

Q. Where have NAS devices been deployed most effectively -- in terms of best applications and size of IT environment or company?

NAS has been predominantly deployed in departmental or workgroup environments. Probably greater than 90% of NAS devices are in these environments. The size of company isn't the issue - it's the application or work environment that dictates what is deployed. Development (engineering CAD or programming) was one of the first big users of NAS. Now we're seeing departments in companies using NAS even if they have a central IT organization. Remote offices such as insurance agents have also been a success story for NAS.

Q. What are NAS "heads" or "gateways"?

Gateways are the NAS controller function (head) without the storage. All NAS resolves down to block storage and most use direct-attached storage (DAS). A NAS Gateway uses the storage in a SAN rather than direct-attached for the block I/O to be done.

Q. Does the StorEdge 5310 NAS support DFS?

DFS (distributed file system) is a hierarchical file system that allows files to be stored across multiple servers and managed as a single group. The StorEdge 5310 NAS can serve as a DFS target. This means that DFS referrals can redirect clients to the StorEdge 5310 NAS, but the StorEdge 5310 NAS does not provide referrals and cannot be configured as a root replica.

Q. What is "port aggregation"?

Port aggregation gives you the flexibility to scale your network I/O in port aggregation or to provide NIC port redundancy in high availability.

Port Aggregation is also known as "channel bonding" or "trunking." This type of bonding lets you scale network I/O by joining adjacent NIC ports. It forms a single network channel of high bandwidth from two or more channels of lower bandwidth. You must have a minimum of two available NIC ports for port bonding, and they must be of the same interface type (e.g., Fast Ethernet with Fast Ethernet).

High Availability port bonding provides NIC port redundancy or failover. More than one NIC port is bonded to a primary port as backup ports. If the primary port fails, the StorEdge switches over to the backup port that is first on the list of "high availability" bonded ports. If that port also fails, the port next on the list is used and so on.

Q. Can more than one default gateway be set?

No. The default gateway is the gateway used when a TCP/IP client needs to send data to a network which it does not have a specific route to. After checking the destination network against the routing table and finding no match, the data is sent to the default gateway. There is no provision within TCP/IP to choose between default gateways.

Some operating systems allow the administrator to configure a second default gateway to be used in the case of failure of the primary default gateway. StorEdge does not currently support this feature.

Q. How many StorEdge 5300 RAID Expansion Units (EU) and 5300 EU's may be attached to the StorEdge 5310 NAS Appliance?

1 or 2 StorEdge 5300 RAID EU's may be attached to a StorEdge 5310 NAS Appliance. Each StorEdge 5300 RAID EU supports up to 7 StorEdge 5300 EU F's. Diskless StorEdge RAID EU's may have up to 8 StorEdge 5300 EU S's attached. Attaching 2 StorEdge 5300 RAID EU's requires a StorEdge 5310 NAS Appliance configuration with two 2-port FC HBA's.

4 StorEdge RAID EU's may be attached to the StorEdge 5310 NAS Appliance. Configurations with 4 StorEdge 5300 RAID EU's will be processed via the "Get to Yes" program and have a 60-day lead-time.

Q. Are other arrays supported with the StorEdge 5310 NAS Appliance?

Only the StorEdge 5300 RAID EU and StorEdge 5300 EU are supported for attachment & capacity expansion on the StorEdge 5310 NAS Appliance.

Q. Are the StorEdge 5300 RAID EU and 5300 EU F the same as the StorEdge 6130 and the CSM100?

No. The StorEdge 5300 Expansion Units are based on the same technology as the StorEdge 6130 and the CSM100 and have been optimized for connection to the StorEdge 5310 NAS Appliance.

Q. Are rackmount kits included with the StorEdge 5310 NAS Appliance and StorEdge 5300 RAID EU and StorEdge 5300 EU?

Rackmount kits must be ordered separately for the StorEdge 5310 NAS (XTA-5200-2URK-19U) and the StorEdge 5300 RAID EU and StorEdge 5300 EU (XTACSM1-RK-3RU-19U).

Q. What configurations are the StorEdge 5300 RAID EU and 5300 EU shipped in?

The StorEdge 5300 RAID EU and 5300 EU are shipped in 14 drive configurations. EU's are configured as follows:

*14 drive StorEdge 5300 RAID EU: RAID 5 with 2 RAID groups of 5+1 with 2 hot spares

*14 drive StorEdge 5300 EU F and EU S: RAID 5 with 2 RAID groups of 6+1

Configurations up to 65TB and 179TB(raw) for FibreChannel and S-ATA ,respectively, are available with a 60-day lead-time via the "Get to Yes" program.

Q. Can the StorEdge 5300 RAID EU and 5300 EU be attached to hosts other than the StorEdge 5310 NAS Appliance?

Attaching the StorEdge 5300 RAID EU and 5300 EU to hosts other than a StorEdge 5310 NAS Appliance is not supported and may void the warranty.

Q. What are the storage capacities per StorEdge 5300 Expansion Unit?

Each StorEdge 5300 RAID EU holds 14 146GB FC drives. There are 2 RAID groups (5+1) and 2 hot spares. This leaves 10 usable drives with 4 for "overhead". Each 146GB drive yields 133GB after formatting.

Each tray has 10 "usable" drives @ 133GB each = 1.33TB/RAID EU (usable)

Each StorEdge 5300 EU F holds 14 146GB FC drives. There are 2 RAID groups of 6+1. This leaves 12 usable drives with 2 for "overhead". Each 146GB drive yields 133GB after formatting.

Each tray has 12 "usable" drives @ 133GB each = 1.596TB/EU F (usable)

For S-ATA config's, a diskless StorEdge 5300 RAID EU is attached to the StorEdge 5310 NAS Appliance with up to 8 StorEdge 5300 EU S's attached to the diskless 5300 RAID EU.

Each StorEdge 5300 EU S has 14 400 GB S-ATA drives. There are 2 RAID groups (5+1 and 6+1) with 1 hot spare. This leaves 11 useable drives with 3 for overhead. Each 400GB drive yields 364GB after formatting.

Each S-ATA tray has 11 useable drives @ 364GB each = 4.004TB/EU S (useable)

Q. What is StorEdge File Replicator?

StorEdge File Replicator is an available software option that allows a StorEdge NAS Appliance volume to be replicated on another remote StorEdge NAS Appliance.

Q. Why do I need StorEdge File Replicator?

Anyone that needs to have an online copy of data available at all times can benefit from StorEdge File Replicator. It can be used to simplify the following data administration tasks:

- *Disaster Recovery

- *Backup

Q. How does StorEdge File Replicator work?

Any update (disk write) made on the source volume is automatically duplicated on the target volume on another StorEdge NAS Appliance using an interconnecting network that is responsible for propagating the data between the nodes. The target StorEdge NAS Appliance can be in the same room as the source, or on the other side of the world, as long as the interconnecting network has sufficient bandwidth to carry the data across.

Q. How does the data move from the source StorEdge NAS Appliance to the target StorEdge NAS Appliance?

All communications between the data source and target utilize the TCP/IP protocol, so the target may be located on any reachable subnet. There has to be a reliable and persistent interconnection to ensure that data is always moving.

Q. How would one use StorEdge File Replicator?

StorEdge File Replicator can be used to help address many data management challenges facing IT professionals today:

- *Disaster Recovery

Without a reliance on slow tape media, StorEdge File Replicator helps to relieve the need for lengthy tape restores. StorEdge File Replicator enhances recovery time in case of a complete loss of data, as businesses can now access mission-critical data from an online backup on a mirror filer. With StorEdge File Replicator properly configured, the target, or mirror, is assured

to be an accurate, near real-time representation of the offline source filer. The StorEdge NAS Appliance target file volume can be brought online quickly to help enable uninterrupted operations.

***Backup**

A StorEdge File Replicator target volume may be dedicated for backing up source volumes. Without affecting production operations, replicated data can be backed up on the target: StorEdge File Replicator enhances operations by moving backup I/O to the remote volume. This shadow processing capability reduces CPU load on the production filer, streamlining operations.

Q. Does StorEdge File Replicator mirror entire files?

No. With StorEdge File Replicator, data replication is performed at the block level of files. This means that individual blocks that compose a file are sent over the network.

This block-level implementation dramatically reduces the required bandwidth on the interconnecting network by transmitting only the data blocks that have changed. For example, if only a single 4 KB block of a multi-megabyte (MB) file is updated, only that 4 KB block is transferred, not the entire file. As a result, StorEdge File Replicator efficiently utilizes StorEdge NAS Appliance resources and available network bandwidth

Q. Can StorEdge File Replicator software be used to replicate between StorEdge 5210 and 5310 NAS Appliances?

Yes. StorEdge File Replicator may be used to replicate data from a StorEdge 5210 NAS Appliance to a StorEdge 5310 NAS Appliance (or vice versa). Both StorEdge NAS Appliances must be using the same level of operating system.

Q. When does the eval period begin with the StorEdge File Replicator evaluation licenses?

The 90-day evaluation period begins on the day the customer faxes their license card information to Sun's Licensing Center for their enablement key.

Q. May CIFS and NFS clients share the same file system at the same time?

Yes, you may access the same file system using both CIFS and NFS. It works as you would expect; files created by e.g. Windows users may be read or modified by Unix users, etc.

As with other multiprotocol NAS servers, there are some initial configuration questions that the administrator needs to answer to allow the system to relate access rights in one domain to user credentials in the other (is a Windows Admin the same as Unix 'root'? What [uid, gid] corresponds to \\ENG-DOMAIN\MERHAR?)

Also, there is no need to purchase additional license for heterogeneous environments. The StorEdge 5310 and 5210, NAS Appliances come with both CIFS and NFS support as a standard feature.

Q. Can a Global Hot Spare for the NAS system reside on any StorEdge 5300 RAID EU or StorEdge 5300 EU F in the System or is a dedicated hot spare for each "EU" required?

When a RAID group goes critical because of a drive failure, any drive designated as a hot spare within any enclosure behind the same set of RAID controllers can be automatically brought online to heal the failure. In other words, it is possible (although rather foolish) to have exactly

one hot spare drive supporting e.g. sixteen RAID 5 groups, spread across a StorEdge 5300 RAID EU plus seven StorEdge 5300 EU F expansion enclosures.

For customer convenience, we will be pre-building two RAID 5 sets plus one designated hot spare for each 14 drive StorEdge 5300 RAID EU and EU F storage enclosure ordered. Customers having specific needs or preferences may modify these configurations.

Q. When implementing StorEdge File Cluster, are 2 dual-port FC HBA's required in the StorEdge 5310 NAS head?

Yes. To insure an appropriate level of reliability and availability, "best practices" require that each server in the StorEdge 5310 cluster must be dual attached, in an appropriate manner, to each StorEdge 5300 RAID EU in the system.

Q. Are the NAS heads in a StorEdge 5310C (cluster) configuration active-active or active-standby? Does each head provide its own service and failover each other or one head provide service only and the other standby?

During normal operation, each server in a StorEdge 5310C cluster owns its own set of LUNs, exports its own set of shares, and services clients accessing those shares. In that regard, the cluster is active/active.

On a failure, the surviving server takes on ownership of its failed partner's LUNs, adds its partner's network addresses as aliases to its own LAN interfaces, and begins to advertise its partner's shares along with its own. This process is quite fast, in general completing within a few seconds after the failure condition has been declared. Clients that were being serviced by the failed device will in general transition transparently to the new server, although some (quite old) CIFS applications reconnect only after requesting user confirmation.

Q. When a StorEdge 5310C node fails over in the cluster pair, will in-flight data be lost for CIFS as well as NFS clients?

NFS writes from a given application thread are directed to a single server. If that server or its supporting infrastructure fails, its network addresses are proxied on the other server, causing the retried NFS operation (failed operations including the Write are retried if not acknowledged or failed) to occur through the other server. If the original write succeeded, the returned ack confirms that the write data was committed to the external array controller (thus, not subject to loss due to failure of the server). In the arrays used in the 5310, cached write data is mirrored between the two controllers within the array, and is thus not subject to loss due to failure of a controller within the array either.)

The CIFS/SMB protocol does not make any guarantees about the disposition of write data on acknowledgement. NFS mandates that an ACK to write data confirms it has been stored "on persistent media", meaning on disk or in nonvolatile memory. Thus, any write sent to a clustered server that is ACKed has been stored on disk and is thus not lost, even if that server subsequently dies. By specification, a write operation that is not ACKed is indeterminate; it is not known whether the write occurred and the response lost, or the write was lost and no change occurred. Thus, the sending application is required to retry the write, and can make no assumptions about the value that would be seen from a read of the data in question.

Because of that preciseness of definition, we can state that "in flight" data is NOT lost in a NFS failover in a StorEdge 5310C cluster, assuming that the NFS client is behaving correctly as per the NFS specs.

Because CIFS does not make that guarantee, the state of written data is indeterminate at the CIFS protocol level, if the CIFS server fails. Thus, no promises can be made about 'in flight' data in a clustered server failover, or standalone server failure.

Q. Does StorADE support the StorEdge 5310 and 5210 NAS Appliances?

The StorEdge 5310 is a self-contained NAS Appliance, of which the 5300 RAID EU is one component. All event logging for all subsystems in the NAS Appliance, including the storage, is maintained in a single system log. All external event notifications such as email alerts and SNMP traps, are generated by the StorEdge 5310, including events it detects within its storage subsystem.

Please note that StorADE does support collection and integration of StorEdge 5310 (and 5210) alerts and events. Customers not already using StorADE may choose to handle notifications directly, for example by directing email alerts to their IT response team.

Q. Does the StorEdge 5310 (or 5210) NAS Appliance support HTTP WEB DAV?

This is asking for Distributed Authoring for web pages, a feature not currently supported. The web server embedded in the StorEdge 5000 NAS products is there for administration purposes only, and not exposed as a general purpose web server.

Q. Is StorEdge 5000 NAS a good fit in an "Exchange" environment?

Not yet. In order to support MS Exchange, or SQL, one needs block level access. This access is expected when iSCSI is supported on StorEdge 5000 NAS Appliances.

We have a great storage consolidation solution in the Sun StorEdge 5000 NAS, and we do not have a solution that will enable you to back up and restore an Exchange Server effectively. To back up Exchange, you need a specialized application like NetApp's SnapManager for Exchange.

In a nutshell, there are several versions of Exchange Server, the most recent of which is Exchange Server 2003. Exchange 2003 can have up to 20 separate databases, every 5 of which have a transaction log associated with them (includes distribution lists, attachments, etc.). This transaction log is used mainly for restoring the Exchange Server, and must be sync'd with the databases for proper restore. To keep the log backups up to date, you also need a means of quiescing the Exchange Server to take a snapshot of it; SnapManager does this by using the VSS (volume shadow service) of Exchange.