

Image Database Development Project

IDD Archival Hardware Architecture and Workflow

Version 1

Don Mendonsa, Faranak Nekoogar, Harry Martz
Lawrence Livermore National Laboratory
Livermore, CA 94551

Work performed on the
Science & Technology Directorate of the
Department of Homeland Security
Statement of Work
PR RSEN-08-00066

September 11, 2008
LLNL-TR-411399



This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Table of Contents

INTRODUCTION	4
COMPONENTS	5
IDD WORKFLOW.....	8
IMAGE PROCESSING.....	9
DATA RETRIEVAL AND EXPORT	10
APPENDIX A	11
APPENDIX B	13
APPENDIX C	16
APPENDIX D	19
APPENDIX E	22
APPENDIX F.....	23
APPENDIX G	27

Introduction

This document describes the functionality of every component in the DHS/IDD archival and storage hardware system shown in Fig. 1. The document describes steps by step process of image data being received at LLNL then being processed and made available to authorized personnel and collaborators. Throughout this document references will be made to one of two figures, Fig. 1 describing the elements of the architecture and the Fig. 2 describing the workflow and how the project utilizes the available hardware.

Components

Fig. 1 represents various components of the storage and archival hardware with real pictures of each device.


- 1: 
10 Gig Network Switch
- 2: 
Application Server
(Current Processing Server)
- 3: 
Input Server
- 4: 
I/O Server
- 5: 
4 Gig SAN Switch
- 6: 
60 Terabyte Active Storage
- 7: 
40 Terabyte Long Term Storage
- 8: 
Tape Drive

Figure 1: LLNL's storage and archival system architecture with actual hardware in place

The following section explains each device's functionality in IDD hardware architecture with respect to the device name in Fig. 1.

1) 10 Gig Network Switch All of the computer equipments are connected to the network through a network switch for communication with other network devices, including servers and the internet. All network components of the IDD hardware including servers and workstations are connected at 10 Gigabits. This high speed connection is ten times faster than what is standard in industry¹, therefore making data throughput much faster and more efficient for IDD project. The 10 Gigabit connection speed is optimal for the handling of large amounts of data transfer and processing in IDD project. [Appendix A](#) includes detail technical information of this switch.

2) Application Server The application server is a workgroup class server and has served as the original processing and IDD data storage server. This server is the current application server and will be installed with Red Hat Linux and run the IDD custom application for cataloging and tracking all IDD data. Exact technical specifications are not available for this server due to the server being an existing server and not part of the IDD hardware acquisition.

3) Input Server The input server is an entry-level server that serves as a computer to receive media, scan for viruses and malware, register data received and transferred to production data storage. [Appendix B](#) detail technical information about this server.

4) I/O Server The I/O server is an enterprise Sun server that houses the file system software and serves as the physical throughput to the storage system. [Appendix C](#) describes the in depth technical information about this server.

5) 4 Gig SAN Switch The SAN switch is a switch specifically for connecting all fiber components of the architecture together, Such as the storage units and the tape drive. At the time of purchase this was the fastest fiber technology available for the components used. [Appendix D](#) provides technical information about this switch.

¹ The standard network connection speed is 1 Gigabit and has an average data transfer rate.

6) Storage Arrays There are currently three Nexsan SataBeasts storage arrays. These are fiber connected storage units with 100 Terabytes usable storage capacity. The storage units are logically combined with the file system software and space allocation is configured through the file system software. Allocation of short term and long term storage is dynamic with FileTek. See [Appendix E](#) for technical information.

7) Tape Unit The tape unit is a workgroup class tape unit. It is fiber connected and has three tape drives and one robot arm. The tape unit uses DLT4 tapes that have a capacity of 800Gigs per tape. See [Appendix F](#) for technical information.

8) File System Software The file system software is Stor House RFS from FileTek. It is a policy based archival file system that allows policy management of data. See [Appendix G](#) for technical information.

IDD Workflow

This section describes the workflow of IDD data from when it arrives at LLNL to when it's processed to when customers retrieve processed data. The flow of this diagram starts on the right in the 'Outside LLNL' pane. Each step is outlined below figure 2.




Figure 2: IDD Hardware Archival System Workflow

Workflow

Step 1. Image data is acquired outside LLNL and placed on a hard drive or optical drive.

Step 2. The media is sent to LLNL for image processing.

Step 3. The IDD Administrator receives the media and copies it to the 'Input Server'

Step 4. The IDD administrator will scan the data for viruses and malware and launch the IDD custom application for the data ingest process to the archival system

Step 5. The custom application will catalog the data in the IDD database registering all the meta data that is available at that time.

Step 6. The data is copied from the 'Input Server' over the 10 gigabit network to the 'I/O Server'

Step 7. The 'I/O Server' will the automatically, using filetek software, copies the data to an 'Active Storage' partition, 'Long Term Storage' partition and to the tape unit , all at once.

NOTE: The filetek policy is what maintains the integrity of the data. Policies are predefined with filetek so that new data is automatically saved to a 'Long Term Storage' partition for up to a year. With this policy it is not possible to erase or change the data, not even from the system administrator or the data owner. The data is simultaneously sent to tape for indefinite storage. There is never a need to perform a off hours backup, data will be automatically moved to tape based on policies. Image processing output can be stored in a data set or to a particular partition upon a policy that could send the data to both 'Long Term Storage' and the tape.

Image Processing

The image processing workflow is not indicated in a figure but utilizes all elements of the environment. Image processing workstations are connected at 10 Gigabits to the same network switch indicated in Figure 2. The image processing process is listed in the following steps.

Step 1. IDD User accesses custom software to process a test plan. The user workstation will connect at 10 Gigabit to the ‘Application Server’ for data lookup and retrieval.

Step 2. Once data is identified, the ‘I/O’ server is accessed and will path the image processing software on the users workstation to the data set on the ‘Active Storage’ partition.

Step 3. The IDD image processing application on the IDD workstation can then process the data.

Step 4. The results of the image processing will result in additional metadata that will be stored in the IDD database. The metadata will be cataloged in the IDD database via the ‘Application Server’ and processed data will be sent directly to tape.

Data Retrieval and Export

The data retrieval and export process will be performed by the IDD administrator. The steps are as follows.

Step 1. A request is submitted to the IDD system administrator

Step 2. The IDD system administrator utilizes the custom application to retrieve requested data from the archive system.

NOTE: The custom application validates that the requestor has authorized access to the data being requested.

Step 3. The requested data is copied to the ‘Input Server’ and then transferred to media for distribution.

NOTE: All data coming to and leaving LLNL must be password protected. This process is completed at the ‘Input Server’.

Appendix A

Cisco Switch

Technical Information

Overview

The Cisco Catalyst® 4900M series is part of the Cisco® data center switching portfolio. The Cisco Catalyst 4900 Series Switches are designed as a top-of-rack access layer switch for use with rack-optimized servers. The Cisco Catalyst 4900M is purpose built as a top-of-rack switch and delivers the buffering, latency, and performance characteristics necessary to excel in that role. The Cisco Catalyst 4900M is a two-rack-unit (2RU) top-of-server rack Ethernet switch optimized for 10/100/1000 and 10 Gigabit Ethernet access devices. It is unique in the top-of-rack space because it is semi-fixed. This flexibility provides a cost-effective and modular migration path from Gigabit Ethernet to 10 Gigabit Ethernet.

The design of the Cisco Catalyst 4900M addresses several common challenges data centers face, offering these benefits:

- Offers hot-swappable flexibility, providing high availability in a dynamic environment along with investment protection in environments with a constantly changing mix of servers
- Provides optional half-cards that can be changed as new servers require different media and port speeds while maintaining the investment of the Cisco Catalyst 4900M base unit
- Designed with 16 MB of shared buffers to help prevent performance degradation under heavy traffic loads and help ensure that revenue-producing packets are not lost due to insufficient buffering
- Eliminates the need to deal with multiple OS and management systems simply to provide flexible access switch transport options

The switch has 8 fixed wire-speed 10 Gigabit Ethernet ports typically used to uplink to an aggregation switch. The switch also features two half-slots that can be filled with any combination of the following:

- 20-port wirespeed 10/100/1000 (RJ-45) half-card
- 4-port wire-speed 10 Gigabit Ethernet (X2) half-card
- 8-port (2:1) 10 Gigabit Ethernet (X2) half-card (Cisco TwinGig Converter Module compatible)
- Infrastructure Simplification and Flexibility

- Provides flexible 10 Gigabit Ethernet configurations ranging from 8 to 24 ports
- Offers more than nine combinations of media types, including Gigabit Ethernet copper and fiber interfaces and 10 Gigabit Ethernet copper and fiber interfaces
- Delivers an investment protection path for 10GBASE-T access upgradeability
- Performance and Scalability
- Offers wire-speed Layer2+ Ethernet switching with 320 Gbps and 250 million packets per second (mpps) of aggregate capacity in 2RU
- Eliminates access port-to-uplink port oversubscription bottlenecks with up to 40 ports of wire-speed 10/100/1000 at the access to 8 wire-speed 10 Gigabit

Ethernet uplinks

- Achieves a latency of less than 5 microseconds with 64-byte packets
- Provides Cisco TwinGig Converter Module support on the 8-port X2 half-cards,
- allowing up to 32 Gbps over fiber optics per switch
- Allows control of transmit queue sizing: default size of TX queue is 8000 packets
- (irrespective of packet size); this can be configured into up to 8 queues with different
- sizes and graded priorities
- Continuous Operation
- Provides online hot addition and deletion of half-cards for in-service

hardware upgrades

- Provides 1+1 AC or DC power supplies and hot-swappable fan tray with redundant
- fans for high availability for mission-critical applications

Appendix B

Input Server

Technical Information

Dell PE2950

Base Unit:

Quad Core Xeon X5450 Processor 2x6MB Cache, 3.0GHz, 1333MHz FSB, PE2950

Processor:

Quad Core Xeon X5450 Processor 2x6MB Cache, 3.0GHz, 1333MHz FSB, PE2950

Memory:

8GB 667MHz (4x2GB), Dual Ranked Fully Buffered DIMMs

Video Card:

LOM NICs are TOE Ready

Video Memory:

Riser with 3 PCIe Slots for PowerEdge 2950

Hard Drive:

400GB 10K RPM Serial-Attach SCSI 3Gbps 3.5-in HotPlug HardDrive

Hard Drive Controller:

PERC6i SAS RAID Controller, 2x4 Connectors, Int, PCIe, 256MB cache, x6 Bkpl

Floppy Disk Drive:

No Floppy Drive for x6 Backplane

Operating System:

Windows Server 2003 Standard Edition

NIC:

Two (2) Integrated Broadcom NetXtreme II 5708 Gigabit NICs TOE Capable

Modem:

Dell Remote Access Card, 5th Generation for PowerEdge Remote Management

CD-ROM or DVD-ROM Drive:

24X IDE CD-RW/DVD ROM Drive for PowerEdge 2950

Sound Card:

Bezel for PE 2950

1x6 Backplane for 3.5-inch Hard Drives

Documentation Diskette:

Electronic Documentation and OpenManage CD Kit, PE2950

Additional Storage Products:

400GB 10K RPM Serial-Attach SCSI 3Gbps 3.5-in HotPlug HardDrive

Feature

Integrated SAS/SATA RAID 0, PERC 6/i Integrated/SAS6/iR

Feature

Universal Sliding Rapid/Versa Rails, includes Cable Management Arm

Service:

Dell Hardware Warranty Plus Onsite Service Initial YR

Service:

Dell Hardware Warranty, Extended Year(s)

Service:

Mission Critical Package: 4-Hour 7x24 On-Site Service with Emergency Dispatch, 2 Year Extended

Service:

ProSupport for IT: 7x24 HW / SW Tech Support and Assistance for Certified IT Staff, 3 Year Service:

Mission Critical Package: 4-Hour 7x24 On-Site Service with Emergency Dispatch, Initial Year

Service:

Thank you choosing Dell ProSupport. For tech support, visit <http://support.dell.com/ProSupport> or call 1-800-945-33

Service:

MISSION CRITICAL PACKAGE: Enhanced Services, 3 Year

Misc:

Energy Smart Redundant Power Supply with Dual Cords

Misc:

400GB 10K RPM Serial-Attach SCSI 3Gbps 3.5-in HotPlug HardDrive

Misc:

400GB 10K RPM Serial-Attach SCSI 3Gbps 3.5-in HotPlug HardDrive

Misc:

400GB 10K RPM Serial-Attach SCSI 3Gbps 3.5-in HotPlug HardDrive

Misc:

400GB 10K RPM Serial-Attach SCSI 3Gbps 3.5-in HotPlug HardDrive

Misc:

Intel 10GBase-SR XFP Single Port 10GbE Optical NIC, PCIe-8

Appendix C

I/O Server

Technical Information

Processor	
Processor	4, 6, and 8 core 1.2 GHz or 8 core 1.4 GHz UltraSPARC T2 processors; one processor per system, maximum 64 threads
Architecture	SPARC V9 architecture, ECC protected
Cache per processor	4 MB integrated L2
Other key features	- Under 95 W (nominal), and fewer than 2 W per thread - Dual, multithreaded, 10 Gb Ethernet and PCIe integrated onto chip - Eight floating-point units - Onboard cryptography supporting 10 embedded security industry-standard ciphers including: DES, 3DES, AES, RC4, SHA1, SHA256, MD5, RSA to 2048 key, ECC, CRC32
Main Memory	
Main Memory	Sixteen FB-DIMM slots, system maximum of 64 GB; support for 1 GB, 2 GB, and 4 GB DIMMs; minimum of 4 GB (4x 1 GB)
Standard/Integration Interfaces	
Network	Four 10/100/1000 Mbps Ethernet Up to 2 x optional 10GbE XAUI connections
Expansion bus	2x 8 lane PCI-E, 2x 4 lane PCI-E, 2x 4 lane PCI-E or XAUI
USB	Four 2.0 USB ports
Mass Storage and Media	
Internal disk	Up to eight 73- or 146-GB SAS disks Internal hardware RAID 0 (striping) and RAID 1 (mirroring) supported
Internal DVD	One slimline CD-RW/DVD-RW
Software	
Operating system	Solaris 10 Operating System 8/07 Ubuntu Linux distribution certified
Preloaded software	The latest Solaris 10 8/07

	Sun Java Enterprise System Release 5 (90-day evaluation) Logical Domains Manager 1.0.1 Logical Domains MIB 1.0.1 Sun Studio 12 Cool Tools for CoolThreads servers
Power Supply	
Power Supply	Two hot-swappable AC 750 W power supply units (PSUs), providing N+1 redundancy
Environment	
Operating Temperature	Sea level to 900 m (3,000 ft): 5° C to 35° C (41° F to 95° F) Above 900 m (3,000 ft): decrease maximum temperature as altitude increases by 1° C/300 m (1.6°F/1,000 ft)
Nonoperating Temperature	-40° to 65° C
Relative Humidity	10 to 90 percent RH, noncondensing, 27° C wet bulb
Nonoperating Relative Humidity	93 percent RH 35° C, noncondensing
Altitude	0-3,000 m (0-10,000 ft)
Nonoperating Altitude	0–12,000 m (0–40,000 ft)
Declared Acoustics	Operating/idling acoustic noise 7.4B (LwAd, 1B=10 dB) Operating/idling acoustic noise 63 dB (LpAm, bystander positions)
Regulations	
	Meets or exceeds the following requirements:
Safety	UL/CSA-60950-1, EN60950-1, IEC60950-1 CB Scheme with all country deviations, IEC825-1, 2, CFR21 part 1040, CNS14336, GB4943
Ergonomics	EK1-ITB-2000
RFI/EMC	EN55022 Class A, 47 CFR 15B Class A, ICES-003 Class A, VCCI Class A, AS/NZ 3548 Class A, CNS 13438 Class A, KSC 5858 Class A, GB9254 Class A, EN61000-3-2, GB17625.1, EN61000-3-3
Immunity	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

Regulatory Markings

CE, FCC, ICES-003, C-tick, VCCI, GOST-R, BSMI, MIC, UL/cUL, UL/DEMKO, UL/S-mark, CCC

Dimensions and Weight

Height

88 mm (3.49 inches); two rack unit

Width

425 mm (16.75 inches)

Depth

714 mm (28.125 inches)

Weight

Approximately 23.58 kg (52 lbs)

Appendix D

4 Gig SAN Switch

Technical Information

Ports per chassis

Eight to sixteen 4-Gbps ports (upgradeable in 4-port increments)

Four 10-Gbps XPAK MSA-compliant ports

Multi-switch fabric support

All topologies including: stack, cascade, cascaded loop, and mesh

Multiple 10Gb or 4Gb links between switches

Adaptive trunking, intelligent path selection

Port types

All ports are universal, auto-discovering, selfconfiguring, and assume the following states: F_Port, FL_port, E_port, G_Port, GL_Port

Performance Features

Fabric Port Speed

4-Gbps, 10-Gbps full-duplex. 4Gb ports autonegotiate with slower devices.

Fabric latency less than 0.4 μ s

Cut-through routing

Fabric Point-to-Point Bandwidth

850 MBps full-duplex on 4-Gbps ports

2550 MBps full-duplex on 10-Gbps ports

Aggregate Bandwidth

238-Gbps per chassis, end-to-end

Non-blocking architecture

Maximum Frame Sizes

2148 bytes (2112 byte payload)

Per-port Buffering

ASIC-embedded memory (non-shared)

Guaranteed 16 credit zero wait state buffer for full performance up to 13km @ 2-Gbps and 2km @ 10-Gbps

Extended distance via credit donation

Media

Hot-pluggable, industry-standard 3.3 volt Small Form Pluggable (SFP) for 4-Gbps ports

Hot-pluggable, industry-standard XPAK optics or copper stacking cables for 10-Gbps ports

Supported SFP types

Shortwave/longwave (optical), active copper

Media transmission ranges (2-Gbps)

Shortwave optical: 500 m (1,640 ft.)

Longwave optical: 10 km (6.2 mi.)

Cable types

50/62.5 micron multimode fiber optic

9 micron single-mode fiber optic

Interoperability

Compatible with FC-SW-2 compliant switches including Brocade®, Cisco®, and McData®
Management interoperability with leading SAN management applications

Fabric Managements

Management methods

QuickTools web applet

Enterprise Fabric Suite 2007

Command Line Interface

API, GS-4 Management Server (including FDMI), SNMP, RADIUS, FTP, TFTP, SMI-S.

Access methods

10/100 Ethernet BaseT (RJ45), serial port (RS-232 with DB9), in-band (Fibre Channel)

Diagnostics

Power-On Self Test (POST)

Optional SANdoctor fabric diagnostics software

Fabric servicess

Simple name server, hardware based zoning, Registered State Change Notification (RSCN), I/O StreamGuard, multi-chassis in-order delivery, automatic path selection, FDMI support

User interface

LEDs, command-line console, and web utilities

Mechanical/Power/Cooling

Single integrated power supply model

Rubber mounting feet, mounting brackets

Optional rail mount kit

Dual hot-swappable power supply model

Hot-swap power supplies, integrated fans

Optional rail mount kit

Front-to-back air pattern (both models)

RoHS compliant

Dimensions

H x W x D: 43.2 x 432 x 305 mm (1.7 x 17 x 12 in) – single power supply

H x W x D: 43.2 x 432 x 508 mm (1.7 x 17 x 20 in) – dual power supply

Weight

4.08 kg (9 lbs) – single power supply

6.80 kg (15 lbs) – dual power supply

Environmental and Equipment Specifications

Operating

Temperature: 0° to 40°C (32° to 104°F)

Humidity: 15% to 80% non-condensing

Altitude: 0 to 10,000 feet

Vibration: 5-500 Hz, random, 0.21 Grms, 10 min

Shock: 4g, 11ms, 20 repetitions

Non-Operating

Temperature: -40° to 70°C (-40 to 158 °F)

Humidity: 5% to 90% non-condensing

Altitude: 0 to 50,000 feet

Vibration: 5 to 500 Hz, random, 2.09 Grms, 10 minutes
Shock: 30g, 13 msec, trapezoidal

Electrical

Voltage: 100 to 240 VAC; 50 to 60Hz

Power load: 120 VAC

Heat output: 100W maximum

Appendix E

Storage Unit

Technical Information

SATABeast Technical Specifications

External Host Interface / Channels:

Dual 4Gb Fibre Channel SFP LC host ports per controller

Dual iSCSI ports per controller

Single or Dual active/active controllers

Internal Device Interface / Channels:

Forty-two SATA/100 or 133 device channels

Power

Current Voltage.....	Nominal	Peak
	350 watts	750 watts
	4.5 A @ 115VAC	6.5 A @ 115VAC
	2.5 A @ 230 VAC	3.5 A @ 230 VAC
	1200 BTU/hour	2500 BTU/hour

Dimensions and Weight

Height	7 in. (4U)
Width	fits in standard 19" EIA rack with supplied mounting rails
Depth	29.5 in.

Appendix F

Tape Unit

Technical Information

Product Description:

Qualstar's TLS-8000 Series offers superior performance and value for LTO Ultrium tape users. With capacities from 4.4 TB to over 422 TB, the TLS-8000 Series economically and reliably automates tape-based applications, whether backing up a single server or as an integral part of an enterprise-wide storage management solution.

The TLS-8000 family comprises six models, providing a broad range of cost-effective solutions. Housing one to eight LTO technology tape drives and up to 264 tapes. TLS-8000 multi-drive library models readily support concurrent reading and writing at data transfer rates exceeding 6.9 Terabytes per hour to meet the throughput needs of today's rapidly-growing networks.

Direct Fiber Attach (DFA) LTO drives connect each drive to maximize SAN performance.

All TLS-Series libraries include [Q-Link](#), Qualstar's web browser accessed remote library manager that enables administrators to configure, upgrade, and monitor any TLS-Series tape library via a company intranet or over the internet from anywhere in the world. Q-Link extends the library's comprehensive front panel menu system and adds additional functionality such as automatic e-mail notification based on rules set by the administrator, and firmware updates can be implemented remotely.

Developed for the RLS-Series libraries, [VIOP](#) is now incorporated in every TLS model. An exclusive Qualstar design feature, VIOP, Variable I/O Port, allows users to match the number of I/O slots to their operational requirements. Applications are not constrained by limits imposed by libraries with a fixed number of I/O slots. Users can assign the number of slots to best match their requirements. VIOP compliments the hardware I/O Ports which continue to be included or available in all models.

Logical Library partitioning is built into every TLS-Series SCSI library. It allows up to four servers to share a single library via a SCSI interface, even if each server is running a separate application. Systems that you couldn't afford to upgrade can now be protected automatically. Support for multiple SCSI busses is also built-in to maximize throughput in direct-attach

architectures.

Qualstar's N2Power subsidiary has pioneered and patented hi-efficiency power technology for years. Storing data in our [tape libraries uses much less power than disk](#), saving money by substantially reducing electrical and cooling requirements.

The TLS-8000 Series also includes barcode scanning and an I/O Port for importing and exporting individual tapes under application control.

Qualstar is known worldwide for quality and reliability. All electrical and mechanical adjustments have been eliminated for easy installation and enhanced reliability. Tape drive reliability is improved by providing individual fans for each Ultrium drive to insure proper cooling. Filtered, positive air pressure is maintained throughout the library to reduce contamination.

Providing built-in capacity and performance growth is a hallmark of all Qualstar TLS-Series automation families. All drive bays are pre-wired: users can easily install additional tape drives for increased throughput. Field upgrade kits are also available to double the tape capacity of some models.

Founded in 1984, Qualstar has been designing and manufacturing tape automation products since 1994. Offering 30 models for AIT and LTO formats, our TLS-Series families consistently set industry standards for quality, performance and value.

Features & Benefits:

Feature	Benefit
Up to eight LTO tape drives	High throughput; supports concurrent read and write
Direct Fiber Attach (DFA) LTO tape drives	Individual fibre channel connections to the SAN
Q-Link	Web browser based remote library manager
Wide capacity range	Common format throughout the enterprise
Automatic alignment	No special tools or procedures
Logical Library Partitioning (SCSI only)	Supports up to four separate partitions.
No electrical or mechanical adjustments	Easy, fast installation
Barcode scanner	Fast media inventory
I/O Port	Application controls tape inventory.
VIOP	Bulk Tape import/export
Positive-pressure, filtered air system	Reduces contamination, protects media
Individual drive cooling fans	Enhanced drive reliability
Brushless motors	Exceptional reliability

Specifications:

7 Models to Chose From		8211	8222	8433	8466	88132	88264
Number of Drives		1-2	1-2	1-4	1-4	1-8	1-8
Number of Tapes		11	22	33	66	132	264
Native Capacity (TB)	LTO 3	4.4	8.8	13.2	26.4	52.8	105.6
Native Capacity (TB)	LTO 4	8.8	17.6	26.4	52.8	105.6	211.2
Compressed Capacity (TB)*	LTO 3	8.8	17.6	26.4	52.8	105.6	211.2
Compressed Capacity (TB)*	LTO 4	17.6	35.2	52.8	105.6	211.2	422.4
Max. Full-height Native Data Rate (TB/hr.)	LTO 3	.57	.57	1.1	1.1	2.3	2.3
Max. Half-height Native Data Rate (TB/hr.)	LTO 4	.80	.80	1.7	1.7	3.4	3.4
Max. Full-height Comp. Data Rate (TB/hr.)	LTO 3	1.1	1.1	2.2	2.2	4.6	4.6
Max. Half-height Comp. Data Rate (TB/hr.)	LTO 4	1.7	1.7	3.4	3.4	6.9	6.9
Avg. Exchange Time (sec.)		14	15	16	16	18	18
Time to Scan Barcodes (sec.)		12	26	41	91	92	108
Barcode Reader		Opt.	Opt.	Std.	Std.	Std.	Std.
I/O Port		Std.	Std.	Std.	Std.	Std.	Std.
Field Expansion Kit		Yes		Yes		Yes	
Height (in./cm)		36/91		45/114		62/157	
Width (in./cm)		13/33		19/48		37/94	
Depth (in./cm)		25/63		30/77		30/76	
Net Weight (lb./kg)		82/37		105/48		359/163	
Power Consumption (watts, avg.)		150		250		470	

Additional Features:

- Individual Fans for Each Tape Drive
- Positive Filtered Air Pressure
- All Leadscrew Driven
- Self-calibrating, Digital Servo Control System
- Lockable Front Door for Media Security
- Inventory Sentry Minimizes Offline Time

Menu Control System:

- 80-Character, Back-lit Display
- Easy Setup, Configuration and Operation and Diagnostics
- Built-in Emulations
- Host-controlled Security Lockout
- Downloadable Firmware Update

General Specifications:

- Interface: LVD/SE SCSI
- Voltage: 100-120, 220-240 VAC
- Frequency: 48-62 Hz
- MEBF: > 2,000,000 Exchanges
- MTTR: < 30 minutes

Appendix G

File System Software

Technical Information

FileTek's StorHouse□ Archiving for File System Applications is a low-cost, easy-to use solution for archiving and rapidly retrieving today's escalating volume of e-mail, voice mail, medical images, spreadsheets, customer statements, word processing documents, and other digitized data in file format. The system provides unique benefits that distinguish it from hierarchical storage management (HSM) products and file system-based archiving alternatives, including:

- **Performance** – Regardless of file size, retrievals of requested data from any storage layer occur in seconds with no need to first restore entire files to disk.
- **Scalability** – Archives can grow to billions of files with no performance degradation due to a unique relational method for storing and indexing all file locator data (metadata about individual files).
- **Reduced data ownership costs** – StorHouse provides affordable, optimal storage of application data on a blended media hierarchy of RAID, MAID, and tape.
- **Built-in compliance support** – Configurable parameters and compliant media types ensure retrievals satisfy industry-based retention rules and archive data remains accessible throughout its required lifespan.
- **Data sharing across the enterprise** – StorHouse promotes secure access to archived data from mainframe, UNIX, and Windows platforms through support for NFS, NTFS, and CIFS file systems.
- **Simplified system administration** – Storage management is completely automatic, including data migration between storage levels, backup, recovery, and replication.
- **Support for Information Life Cycle Management (ILM)** – StorHouse architecture and features complement ILM strategies for managing data optimally from creation and use through archive and deletion.

The solution consists of a *StorHouse* data management system with *StorHouse/RFS* for presenting a standard NFS, NTFS, or CIFS file system interface to applications and *StorHouse/SM* for providing automatic storage management features.

How StorHouse/RFS Works

The StorHouse/RFS file system interface presents itself to an application as a virtual drive (an NFS, NTFS, or CIFS mount or share point). Applications read and write files to the virtual drive just as they would to any mapped Windows drive or NFS mount point.

Collecting and writing files. When an application writes a file to the virtual drive, StorHouse/RFS writes that file to a temporary staging area. Then, it groups multiple files from the staging area into a container, or collection, thereby providing highly efficient storage for small application files. At the same time, StorHouse/RFS generates file locator information such as file path, name, retention period, and size for each application file in the collection. When a collection reaches a pre-set (configurable) size or at user-specified intervals, StorHouse/RFS copies the collection to StorHouse/SM for permanent file storage. To StorHouse/SM, each collection is one file composed of many individual application files. After creating a collection, StorHouse/RFS inserts the associated file locator information into a relational table. On retrieval, the file locator data enables StorHouse/RFS to quickly identify the StorHouse file containing the requested application file.

Retrieving files. StorHouse/RFS uses block level retrieval to dramatically reduce network traffic and improve application throughput. An application opens a file from the virtual drive just as it opens a file from any NTFS, remotely mounted CIFS, or NFS drive. StorHouse/RFS intercepts the request, interrogates the locator data, and determines where the collection resides. Then, it returns only the requested data, which the application opens/reads using native NTFS, NFS, or CIFS I/O.

Summary

StorHouse Archiving for File System Applications is an easy-to-use, survivable, scalable, highly accessible, and cost-effective archiving solution for file system applications. It supports compliance requirements as well as ILM processes and procedures for storing data affordably according to its changing value over time. StorHouse/RFS provides the standard interface, while StorHouse/SM supplies the storage management features. For more information, visit www.filetek.com or contact your FileTek account representative.