

File Retention Compliance – Storage Architecture and Installation of KOM Networks KOMworx software



Executive summary.....	2
Introduction.....	3
Challenges of Storage Management	4
File retention compliance	4
Workflow logic.....	4
Cost control.....	5
Storage Architectures	7
The life-cycle of storage	7
Direct Attached Storage.....	8
Network Attached Storage.....	9
Storage Area Network.....	10
Hierarchical Storage Management.....	11
ILM file retention	12
Hardware and software requirements	14
Hardware requirements	14
Software requirements	15
Storage system hardware planning	15
KOMworx software configuration planning.....	17
Installing the KOM Networks KOMworx software.....	18
Installation overview	18
Software Installation procedures	18
Configuring the KOM Networks KOMworx software	22
Configuration overview.....	22
Software configuration procedures.....	22
HP Services.....	36
Glossary – Storage Management	37
Glossary – Storage Technology	37
Summary	40

Executive summary

Government regulations, specifically HIPAA, The Health Insurance Portability and Accountability Act, also known as Kassebaum-Kennedy requires that all healthcare institutions must comply with specific file retention policies. Informal procedures to delete old records may result in financial consequences. Other regulations such as Sarbanes-Oxley include criminal prosecution for intentional or even unintentional data loss.

KOM Networks KOMworx software utilized in conjunction with HP ProLiant servers provides the file retention compliance storage solution by which healthcare providers and other corporations can ensure cost-effective compliance to file retention regulations.

By utilizing an HP ProLiant server and an HP StorageWorks Modular Smart Array, initial acquisition and deployment costs and total cost of ownership is but a fraction of competing proprietary storage vendor alternatives.

This architecture and installation guide presents the advantages of a file-centric storage management system and the audited procedures for installing KOM Networks KOMworx software. A basic configuration, utilizing HP equipment was tested and documented within.

This file compliance solution addresses many storage challenges of the healthcare industry and enterprise computing by providing:

- File retention compliance
- Cost control through the utilization of a SATA storage subsystem
- Unified storage virtualization and management of separate storage subsystems
- Capacity expansion to hundreds of Terabytes
- Importation of any legacy managed or unmanaged file based data
- Utilization of highly reliable fast SCSI based storage subsystem
- Integration of extremely stable MO storage media and jukebox storage subsystem

NOTE: HP provides this document only as an introduction of a file retention compliance architecture offered by a software vendor that operated in accordance with the vendor's specifications on HP storage and server products. Compliance with laws and regulations, including HIPAA and Sarbanes-Oxley, remain your responsibility. Your own legal counsel should be consulted.

Introduction

Storage management software has traditionally been designed from the perspective of a specific storage device to be managed. As complex business needs have evolved and storage subsystems cost and performance properties have diverged, the need for an umbrella storage management system has become clear. New storage management systems enable cost effective automated management, consolidating different types of physical storage into a single unified storage system.

A simple test environment was configured and reviewed during the preparation of this material and is documented in this installation guide. This guide provides a proof point to the compatibility and value of file-centric storage subsystem. A step-by-step installation guide and bill of materials (BOM) are also included.

This installation guide presents the necessary details to install and configure KOMworx on an HP ProLiant DL380 G3 server running Microsoft Windows 2003 using an HP 300mx StorageWorks Magneto Optical jukebox and an HP StorageWorks Modular Smart Array 1500 (MSA1500), thereby overcoming the inherent file management challenges of traditional hardware-centric storage system.

Challenges of Storage Management

This section provides a review of business drivers for storage management. Multiple business drivers mandate efficient storage management: file retention compliance is but one of these drivers. If one addresses these requirements with a monolithic storage centric strategy of 'forklift' upgrades, each upgrade creates a new pool of separately managed storage and challenges in the migration and ongoing management of the separate storage subsystems.

Storage management, specifically file-centric storage management, provides automated best practices file-based management responding to different business drivers. These business drivers include:

- File retention compliance
- Workflow logic
- Cost control
- Performance and reliability
- Security
- New services methods and data formats

By bringing separate storage subsystems under a single umbrella, all these business drivers can be rationalized into an automated policy-based process, providing traditional access and concurrently fulfilling multiple complex policies.

Separate business drivers mandating storage policies are reviewed below.

File retention compliance

File retention legal requirements such as HIPAA and Sarbanes-Oxley mandate comprehensive file retention policies. These mandates can be met by the design and implementation of a fixed data-centric storage management system providing complete automated file lifecycle management. This system provides the management of all patient files from creation, ongoing use and modification, fixed reference access, including temporary or permanent archival.

Workflow logic

Workflow logic evolves from years of successful execution in an organization. Initially, informal procedures become formally documented and refined providing a solid foundation for business process data flow. Storage systems must accommodate these effective business processes to simplify the adoption and integration. Thus, flexible data-centric storage management systems that adhere to the individual business requirements are more effective than their hardware volume centric counterparts.

Cost control

Cost control may be addressed by leveraging the different physical cost characteristics of the various storage subsystems. Recently available SATA hard drives are considerably more cost effective than the traditional and currently more common SCSI counterparts. Storage subsystems built on SATA technology are challenging the traditional storage methods by providing vast capacity of low cost storage. The SATA storage subsystem alone only provides another physical repository; storage management software provides the virtualization and umbrella management of the various storage subsystem components.

Performance and reliability

Performance and reliability requirements may be addressed by creating policies for the migration of files between various physical storage subsystems. A broad continuum of performance, reliability, and cost exist. The MTBF (Mean Time Before Failure) of the SCSI-based storage of the MSA-1500 is 1.5 million hours at 100% utilization while the MTBF of the SATA drives is 500,000 hours with 20% utilization. By creating simple policies based on file-centric attributes, files will always be managed in a way to ensure the optimum physical storage subsystem is leveraged during the life of a given file.

Security

Security of patient records may traditionally be considered prevention of unauthorized access. Security from a broader prospective might be the prevention of unauthorized or unintentional deletion. Computer viruses have impacted healthcare infrastructures at an increasing rate each year. By managing the permissions on each file and by extending the storage management system to override the base logic of the operating system permissions, even the owner of a newly created file may not have the ability to delete the file, if the rules of the storage management system so dictate. By providing these controls in the storage infrastructure all systems, users, and applications are forced to comply with the security policies specified by the system architect.

New services, methods and data formats

New services as well as new methods and data formats continually expand the role of storage systems in all healthcare infrastructures. Consequently, as various departments are refreshed methods and formats change. If the resulting output is file-based and it is stored consistently in a migrating repository, the evolution of the departments will support consistent access and business continuity. These application format agnostic methods insure that new challenges may be met with comprehensive management achieving the business goals.

By providing a file based extensible management system, new file types, users, and groups can be efficiently incorporated. As digital formats impact more departments, storage management becomes even more critical. The system must be able to grow and change.

Some common examples include:

- Medical images, X-rays, EKGs
- Lab reports
- Scanned registration and consent forms
- Billing documents
- Insurance filing correspondence
- Security access logs

NOTE: Applications broadly use either database services or file systems to store and retrieve data. While one would not expect to store a database in a file-centric storage management system, many patient management applications create copious files to provide record retention and external file attachment capabilities. A file-centric storage management system is the logical system to address these file based requirements.

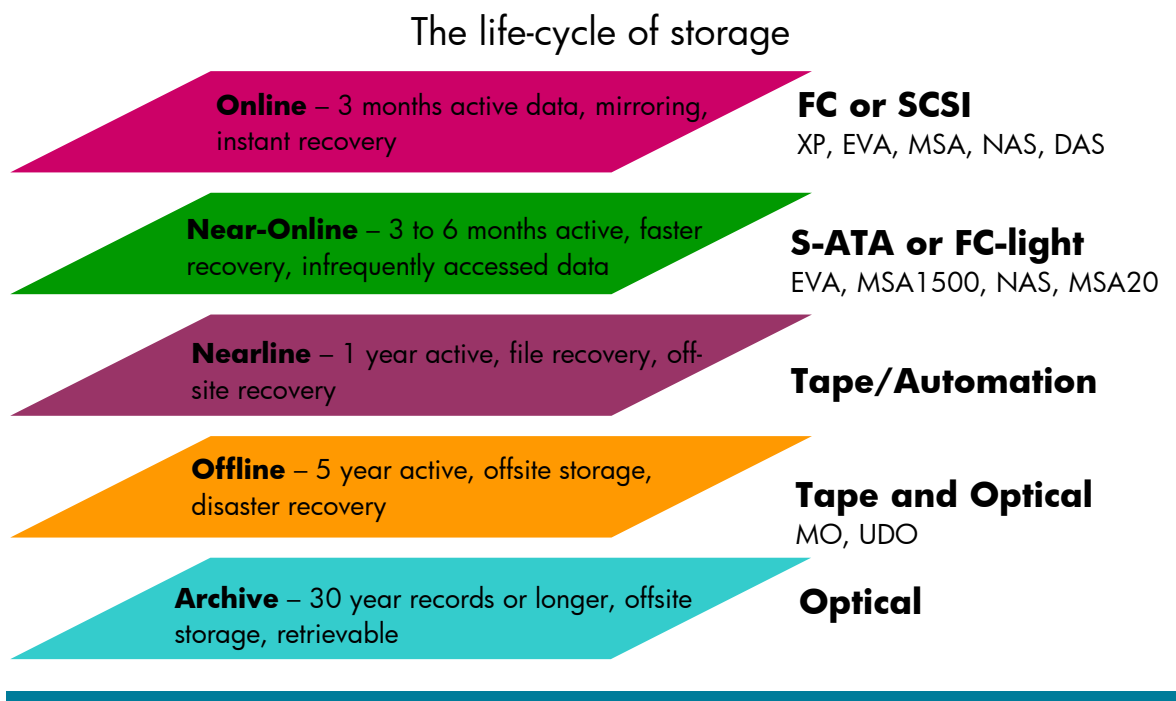
It is recommended that readers unfamiliar with individual industry acronyms refer to the glossary provided toward the end of this document as required.

Storage Architectures

This introduction to storage architectures contrasts the differences in each method of managing storage and is primarily intended to highlight the significant advantages of a data-centric storage management system, specifically KOM Networks KOMworx.

The life-cycle of storage(Figure 1) provides a taxonomy of storage subsystems, illustrating various types and applications. Depending on the business requirements and application performance requirements, various technologies are utilized during the life of a file.

Figure 1. The life-cycle of Storage



In the following section, both physical architectures and storage management systems provide a view of classic storage architectures to illustrate the evolution of storage management.

Example storage architectures and data management systems reviewed include:

- DAS – Direct Attached Storage
- NAS – Network Attached Storage
- SAN – Storage Area Network
- HSM – Hierarchical Storage Management
- ILM – Information Lifecycle Management

Direct Attached Storage

DAS examples are provided as an evolutionary starting point for storage architectures. Figure 2 shows examples of stand alone access devices that contain internal storage where data files may reside. Servers that utilize internal hard drives are also examples of direct attached storage. In these access devices, illustrated below, file management is often difficult and expensive. This may lead to unmanaged pools of data.

Figure 2. Direct Attached Storage

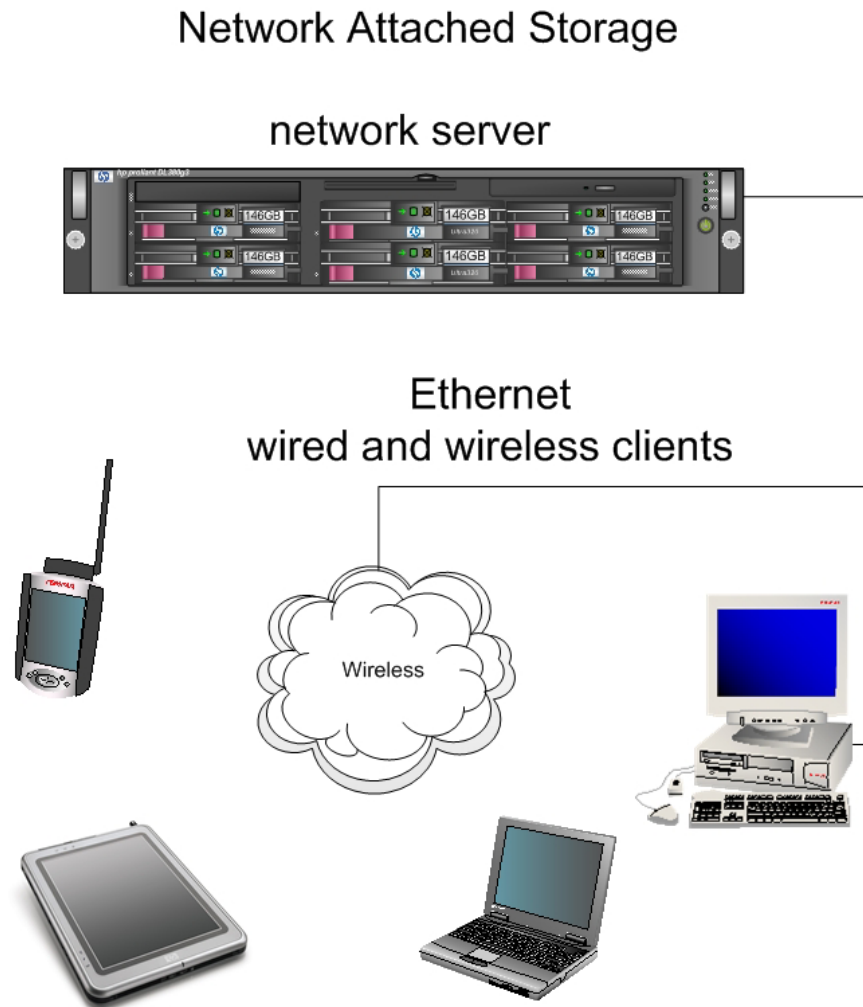


Unmanaged local storage subjects data to unacceptable risk while preventing unified access.

Network Attached Storage

NAS provides common storage and retrieval for client data throughout the network. Typically all files are stored in the same physical storage system. Storage capacity may be expanded by connecting external storage via SCSI or fibre channel (FC). Figure 3 shows an example of network attached storage system.

Figure 3. Network Attached Storage

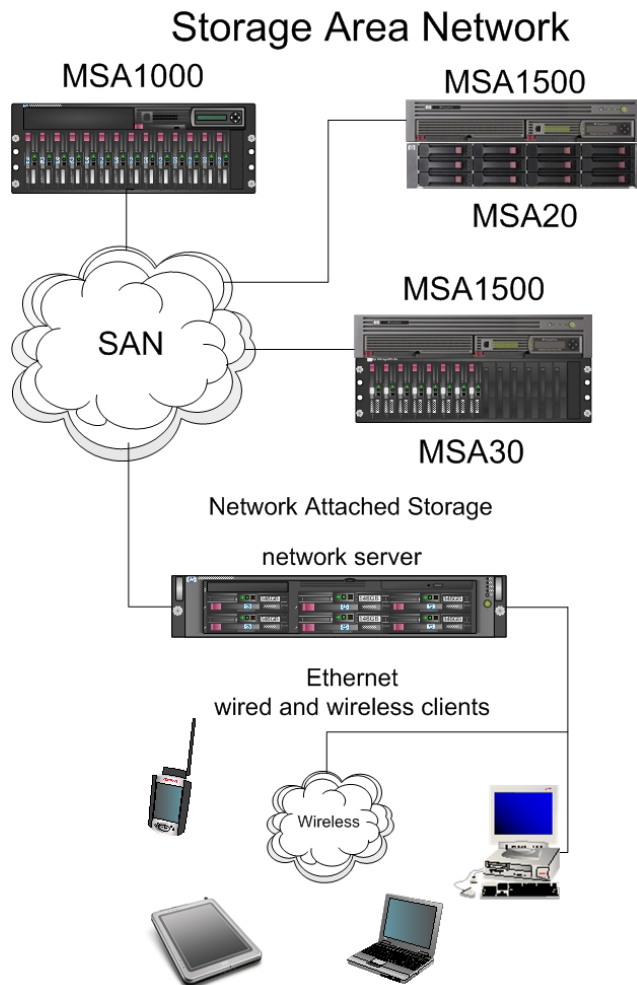


Storage management of a NAS is device or subsystem oriented; all files are treated equally and no file migration capabilities exist. As storage requirements approach system capacity, capacity expansion, or forklift upgrades are inevitable.

Storage Area Network

A SAN may start as a single fibre channel switch used to connect multiple storage arrays to a single server via a fibre channel HBA (Host Bus Adaptor). It can grow to encompass countless hosts and storage arrays. Figure 4 shows an entry level deployment. The SAN cloud might be replaced in this example with a HP StorageWorks SAN Switch 2/16, (HP part number 322118-B21).

Figure 4. Storage Area Network




Several independent small SANs such as the one shown here may provide a cost effectively method to scale separate storage subsystems as apposed to the creation of a single large SAN. This would simplify separate management policies for the various storage requirements in a single institution.

Many options exist in the management of SAN-based storage. These include volume migration, replication, snapshot, and backup. While these volume-centric tools may provide value in fault tolerance, they don't specifically address the challenges of file retention.

Hierarchical Storage Management

HSM provides migration of files based only on their age. While HSM does introduce manage of storage from the prospective of a file, it limits the options of the solution architect to compressively address the latest business goals. File retention compliance requirements, and technical performance requirements may not be thoroughly addressed by the capabilities of HSM. Figure 5 shows the addition of a large automated tape library to the previous example. HSM software must be added to the network server to provide the concatenation of both tape and the native file system.

Figure 5. Hierarchical Storage Management



As shown the HSM system utilizes both SAN storage and an automated tape library to provide age based file migration. Due to the sequential access of tape, unacceptably long delays may result during the retrieval of migrated data.

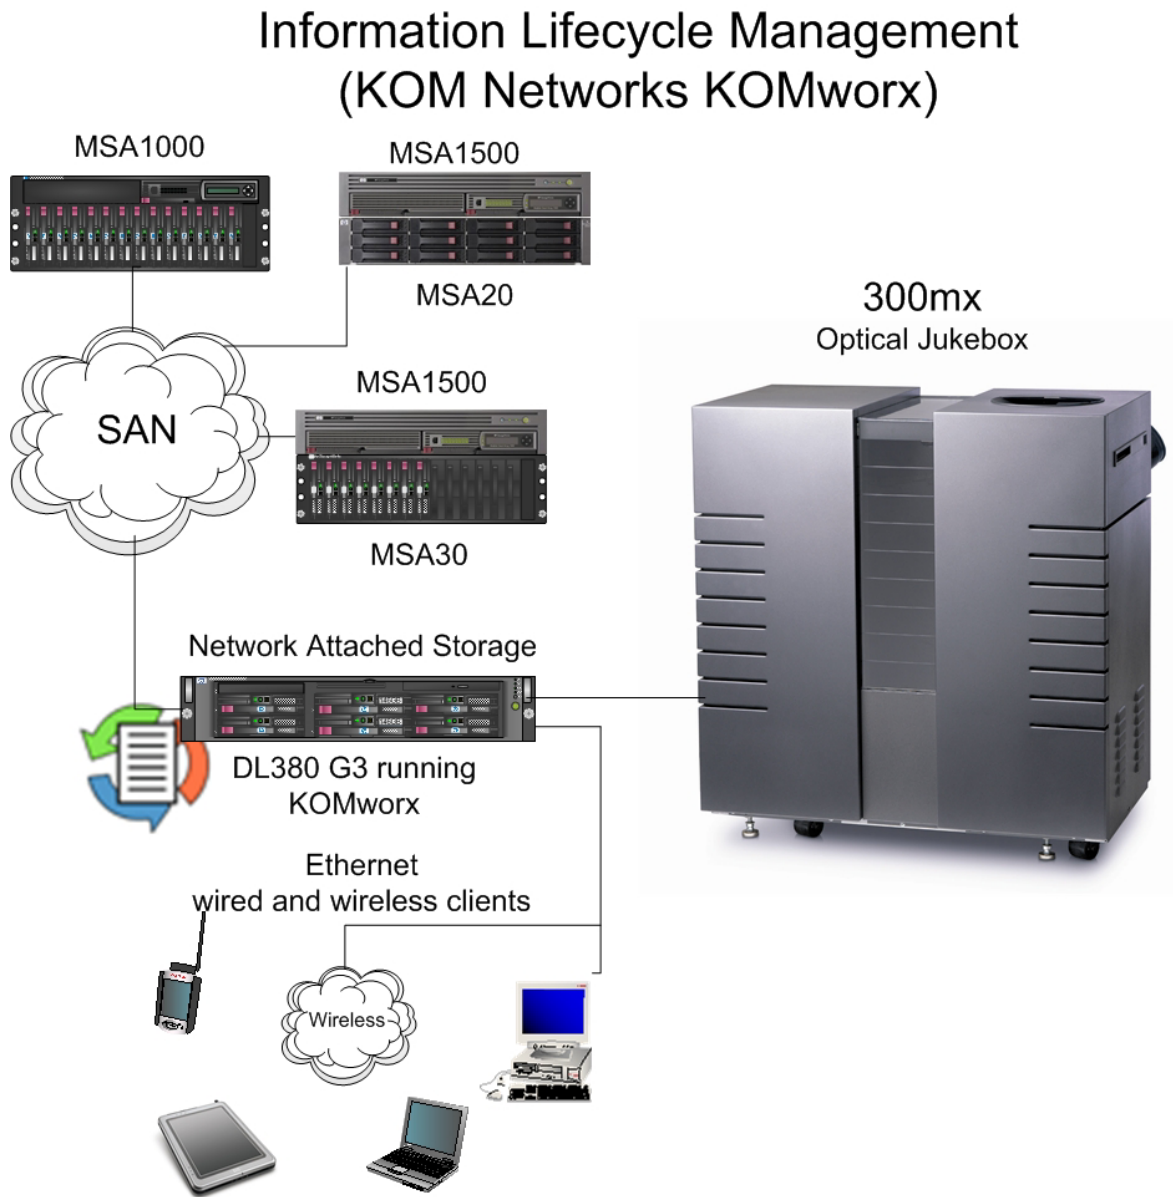
ILM file retention

ILM file retention builds on the common theme and adds a magneto optical (MO) jukebox as the nearline, offline, and archive storage. Depending of the solution architect's specific design, the resulting system can provide capacity growth and management of data as it moves through the lifecycle. The architect may choose to design a system that permanently archives all data as it moves through the lifecycle or the system may manage data only for the required time set forth by management and legal compliance requirements.

The resulting system may be designed to incorporate both SCSI and SATA drive technologies or may also be supplemented through the addition of an MO jukebox. By managing data for a specified period of time, the system may utilize a fixed amount of storage capacity and meet stable business conditions. As old data becomes eligible for deletion, new data enters the system. As the volume of business grows or as new data requirements are encountered new raw capacity and policies may be added to meet additional requirements.

Figure 6 shows a variety of physical components and connects the storage arrays through a SAN. The recommended initial design components are shown in the next section, Hardware and Software requirements.

Figure 6. Information Lifecycle Management



This example shows the design of incorporating KOM Networks KOMworx software. Unbeknownst to the end user or the network application, the familiar Windows network file share is actually a complex orchestration of business workflow logic, cost control procedures, and file retention compliance policies. Files are available through the network, yet the actual physical location of the data is managed automatically by the simple configuration of the storage management system. Files seamlessly migrate according to the specific rules defined.

Hardware and software requirements

Hardware requirements

HP and KOM Networks tested the HP ProLiant DL380 G3 server, the HP StorageWorks MSA-1500 storage array, and the HP StorageWorks 300mx Magneto Optical jukebox. The specific part numbers are shown in Table 1 below. Figure 7 provides a graphic of the tested configuration.

Table 1. Hardware BOM

Quantity	Description	Part Number
ProLiant DL380 G3 server components		
1	Intel Xeon Processor 3.2 GHz/533MHz - 512KB/2MB	352529-001
2	FCA2214 Host Adapter (Linux, Microsoft 32-bit)	281541-B21
6	36GB 10K U320 Pluggable Hard Drive WW	286713-B22
MSA-1500 Storage Array components SATA		
1	Modular Smart Array 1500	AA986A
1	MSA20 - Serial ATA enclosure	335921-B21
12	250GB Serial ATA hard drive, universal carrier	349239-B21
1	2 m LC-LC Multi-Mode Fibre Channel Cable	221692-B21
MSA-1500 Storage Array components SCSI		
1	Modular Smart Array 1500	AA986A
1	MSA30 single bus enclosure	302969-B21
6	72.8 GB Pluggable Ultra320 Universal Hard Drive,15,000 rpm (1")	286778-B22
1	2 m LC-LC Multi-Mode Fibre Channel Cable	221692-B21
300mx Optical jukebox		
1	HP StorageWorks Optical 300mx Jukebox - 9.1GB drive technology	C1150M
1	SCSI Controller, Ultra160, LVD, Dual Channel, VHDCI	284688-B21
1	SCSI Cable, HD68 to HD50, 2m C2906A / 5181-7707	5181-7707
SAN switch(optional)		
1	HP StorageWorks SAN Switch 2/16	322118-B21

Figure 7. HP Tested Configuration

Tested Configuration



DL380 G3

MSA1500

MSA20

MSA30

300mx



Other HP ProLiant servers might also be considered, including the HP ProLiant ML580 G2. While these servers were not tested for this audit, one would expect similar results. The specific configuration is flexible to the requirements. The use of latest version of the UDO jukebox would be logical; UDO equipment was unavailable at the time of testing.

Software requirements

The following software is required for the installation described in this guide:

- Windows 2003 Standard Edition -- CD ROM
- HP ProLiant essentials foundation pack HP SmartStart -- CD ROM
- KOM Networks KOMworx -- CD-ROM (Version 2.72.1159 tested)
- KOMworx Registration file -- provided to customer by KOM Networks

Storage system hardware planning

Initial hardware planning before ordering the hardware and software is recommended. Depending on the size of the storage subsystems and the performance required, multiple storage array controllers and disk subsystems may be required. Be sure to review best practices for deployment of the MSA1500. It is recommended that the physical storage subsystem is initially at only 50 percent of the physical capacity required.

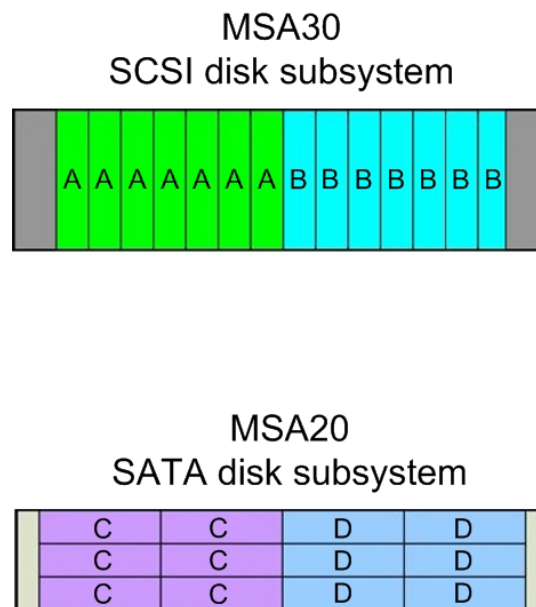
Consideration of the following should be included:

- Hardware scaling methods
- Performance requirements
- Server and hardware recovery methods

Having a simple plan for expansion will serve system administrators and technicians well. For example; add physical storage to the SATA system in full MSA20 increments. Partition two logical disks in each MSA20. This will provide approximately one terabyte per logical disk and two logical disks per MSA20, based on allocating storage with 6 SATA drives of 250 Gigabytes each. For the SCSI storage, add storage in full MSA30 disk subsystems. Configure a logical disk per 7 SCSI drives. This will provide approximately 600 Gigabytes per logical drive and two logical drives per MSA30.

Figure 8 illustrates the recommended logical disk configuration.

Figure 8. Recommended logical disk configuration



Performance may be addressed allocating more SCSI storage in a specific design and by balancing the number of logical arrays and logical disks that are configured on the arrays. In the largest of requirements complete standalone systems may be required for various departments. Configuration of storage arrays vary as the number of files, and the size files vary in a deployment.

Server and hardware recovery methods may be addressed by utilizing a tape backup on the local server storage and purchasing spare server hardware components and storage subsystem components. Various service contracts are available for purchase providing service response times to meet the most critical of requirements.

Versions of the KOMworx software with the capability of maintaining multiple instances of all files stored should be available at the time of publication. This provides a notion of distributed file replicas ensuring extensive recovery capabilities.

KOMworx software configuration planning

The system architect is responsible to provide a design for deployment that meets both the business goals defined by management and the technical detail required for necessary sizing and performance. Consultation with KOM Networks or HP Services personnel is recommended. This document is only an introduction; see KOM Networks documentation for necessary detail. General goals include: understanding software capabilities, designing specific policies, and outlining testing procedures to be used to confirm correct execution.

Installing the KOM Networks KOMworx software

It is assumed that the HP ProLiant DL380 server hardware has been installed. Two FC HBAs and the SCSI adaptor should be installed at this time. See the documentation included with the server for complete explanation.

Installation overview

A general installation overview follows. Review the KOM Networks documentation for a more complete review of the installation and configuration.

The installation procedure consists of the following general steps:

1. Install the operating system on the server without the storage arrays connected.
2. Connect the storage arrays and install the included support software.
3. Configure the storage arrays.
4. Partition and format the logical disks.
5. Connect the MO jukebox.
6. Install the KOMworx software package.

Software Installation procedures

These are the specific commands entered to install the tested software and hardware configuration. Required installation and configuration will vary and is specific to the design goals and management policies required.

1. Install the operating system.

Windows 2003 Standard Edition was installed by a simple drag and drop operation in a RDP management console. This automatically installed the operating system and added the HP ProLiant essentials management agents. Manual operating system installation may be accomplished by the use of the ProLiant essentials foundation pack HP SmartStart CD. In data centers where operating system installation is a common occurrence, RDP may provide time savings. Refer to <http://www.hp.com/servers/rdp>.

2. Connect the storage arrays and install the included support software.

Connect the MSA1500(s) via the FC cables to the FC HBA(s) and install the support software including the Qlogic HBA driver and the HP Array Configuration Utility. Use the CD and the installation documents provided with the MSA-1500.

3. Configure the storage arrays.

Each MSA1500 is connected individually to a FC HBA in the server. Each MSA1500 has homogenous storage connected to it via a SCSI cable provided with the MSA20(s) and the MSA30(s). If capacity requirements exceed that of a single MSA1500, SAN

expansion is possible through the use of the optional HP StorageWorks SAN Switch 2/16, HP part number 322118-B21.

3.1 Create each array through the use of the HP ACU (Array Configuration Utility).

Click **Start/All Programs/HP System Tools/HP Array Configuration Utility/HP Array Configuration Utility**

Security settings may need to be set on the initial access through Microsoft Internet Explorer. Expand the Internet Explorer window and click on the MSA1500 controller connected to the MSA30, shown at the left. The MSA1500 and associated MSA30 disk subsystem management lights will blink.

3.2 Create the array.

Click on **Create an Array**

Create an RAID ADG (Advanced Data Guard) disk array utilizing the SCSI disks. One array was created for the six SCSI drives. Clear the checkbox of any disks that are required to be assigned as spares later.

3.3 Assign the spare(s).

Click on **Spare Management** at the right and assign the spares as required.

3.4 Create the logical drive.

Click on **Create a Logical Drive**

One logical drive was created utilizing the entire ADG array. It is recommended to utilize a separate logical drive for each 1 terabyte of storage.

3.5 Configure the RAID Array ID for the controller.

Select **Controller Settings** and step through to the RAID Array ID and enter **SCSI** as the ID for the controller.

3.5 Save your changes.

Click on **Save**

3.5 Click on the remaining MSA1500 controller shown at the left. The MSA1500 and associated MSA20 disk subsystem management lights will be illuminated.

3.6 Create the SATA logical array.

Click on **Create an Array**

Create a STAT Logical array utilizing the SATA disks. One array should be created per MSA20. Clear the checkbox of any disks that are required to be assigned as spares later.

3.3 Assign the spare(s).

Click on **Spare Management** at the right and assign the spares as required.

It is recommended to assign at least one spare per MSA20.

3.4 Create the logical drive.

Click on **Create a Logical Drive**

The first logical drive was created utilizing 1,000,000 megabytes. The remaining logical drive was slightly smaller. This depends on the number of spares assigned.

3.5 Configure the RAID Array ID for the controller.

Select **Controller Settings** and step through to the RAID Array ID and enter **SATA** as the ID for the controller connected to the MSA20 Disk subsystem.

3.6 Save your changes.

Click on **Save**

4. Partition and format the logical drives.

4.1 Use Windows Disk Manager to initialize the logical drives as basic disks.

Click **Start/ALL Programs/Administrative Tools/Computer Management**

Click **Disk Management**

The Windows Initialize and Convert wizard will be automatically launched. Step through the wizard, but only initialize the drives. Do Not Convert any of the drives to Dynamic.

4.2 Partition and format the logical disks with a single primary partition that utilizes all the existing space.

Right-click the graphic of each new logical drive and select **New Partition**

NOTE: A single initial SCSI primary partition should be created with 50 Gigabytes to serve as the Metadata Volume. Create a second primary partition utilizing the balance of the first SCSI partition. A drive letter is required for the Metadata Volume.

The New Partition Wizard will be started. Step through the defaults of wizard but do not assign a drive letter to the partition. A volume name was assigned to aid in identification of the storage, for example: **MDV, SCSI1, SATA1, SATA2**. Repeat this procedure for each new logical drive added to the system. Allow the formatting process to complete before proceeding to the next step.

5. Connect the MO jukebox.

Use the 50 to 68 pin SCSI cable to connect the optical jukebox 50 pin connector to the SCSI HBA installed in the ProLiant DL380 G3 server. Reboot the server and power up the MO drive. Confirm that the controller recognizes the MO jukebox.

6. Install the KOMworx software package.

A CD was provided by KOM Networks that included a registration file. This registration file enables the required features and functions.

Navigate to the Setup directory of the CD

Click on **KOMworx_Installer.exe**

Step through the install and be sure to click on the **Add Hardware button** to add the hardware interface for the MO jukebox.

This concludes the software installation.

Configuring the KOM Networks KOMworx software

Special attention must be given to the following section; it is recommended that only personnel that have been trained on the KOM Networks software or have reviewed the software extensively perform this procedure. To ensure success a solution architect must design the policies to be implemented with both the knowledge of the capabilities of the KOMworx software and the specifics of the customer's requirement.

This section is only provided as an introduction to the user interface and the methods for configuration. While the concepts are introduced, the specific configuration required will be completely unique. This section provides exposure to the reader of the capabilities yet the reader should not imply limits to the capabilities as the tested configuration was only to review basic functionality.

Configuration overview

A configuration overview of the tested system follows. Review the KOM Networks documentation for a more complete review of the installation and configuration.

The Tested configuration procedure consists of the following general steps:

1. Start the Management Console and complete the software registration procedure.
2. Create the ownership tags and the storage resource pools.
3. Assign raw storage to resource pools.
4. Create a virtual volume.
5. Create the remaining storage groups.
6. Create the File Life Cycle.
7. Share the virtual volume to the network.

Software configuration procedures

These are the specific commands entered to configure the test system. The necessary configuration will vary and is specific to the design goals and management policies required.

1. Start the Management Console and complete the software registration procedure.

Click **Start/All Programs/KOM Networks/KOMworx/KOM Management Console**

2. Create the ownership tags and the storage resource pools.

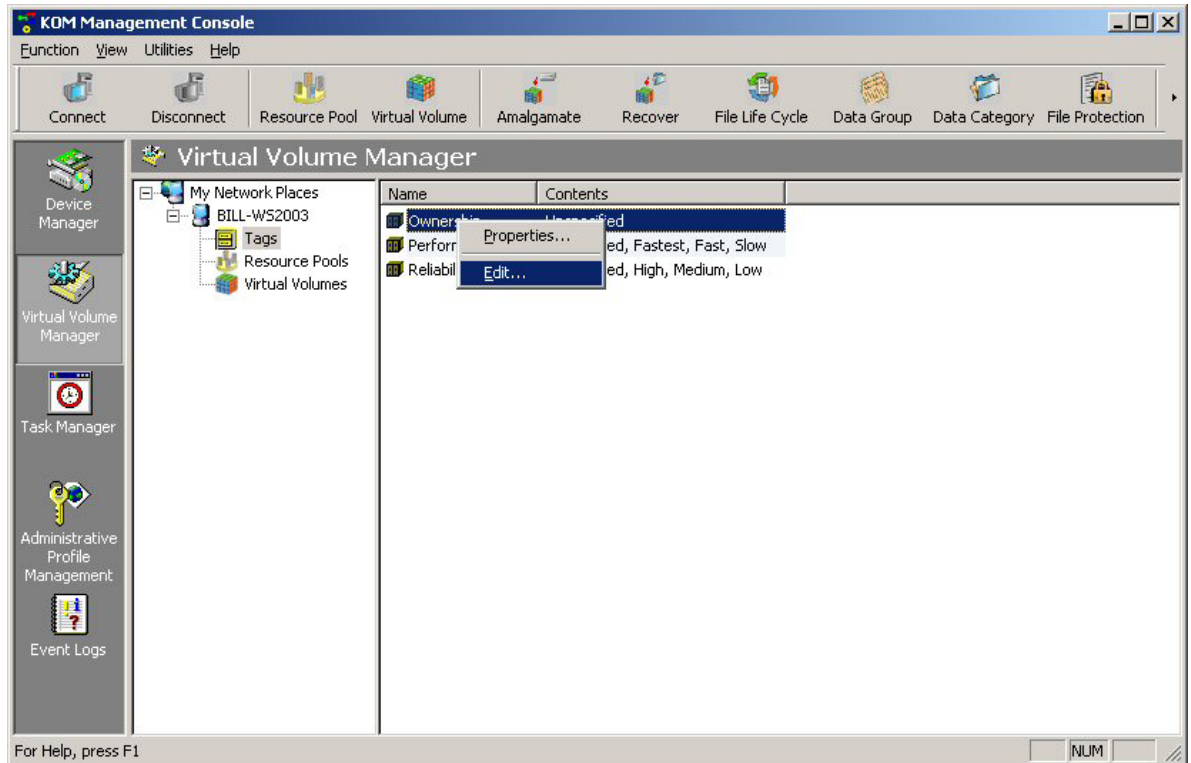
Click on the **Virtual Volume Manager** at the left.

Navigate to the **Tags** icon shown in under the server.

Click on **Ownership** at the right.

Figure 9 provides guidance to select ownership Tags.


Figure 9. Select Ownership Tags



Right-click **Ownership** and select **edit**.

Insert a Tag element for **SCSI**, **SATA** and **MO**. Figure 10 provides an example. The first icon to the right of **Tag Elements:** provides the insertion and editing function.

Figure 10. Insert Tag Collection Elements

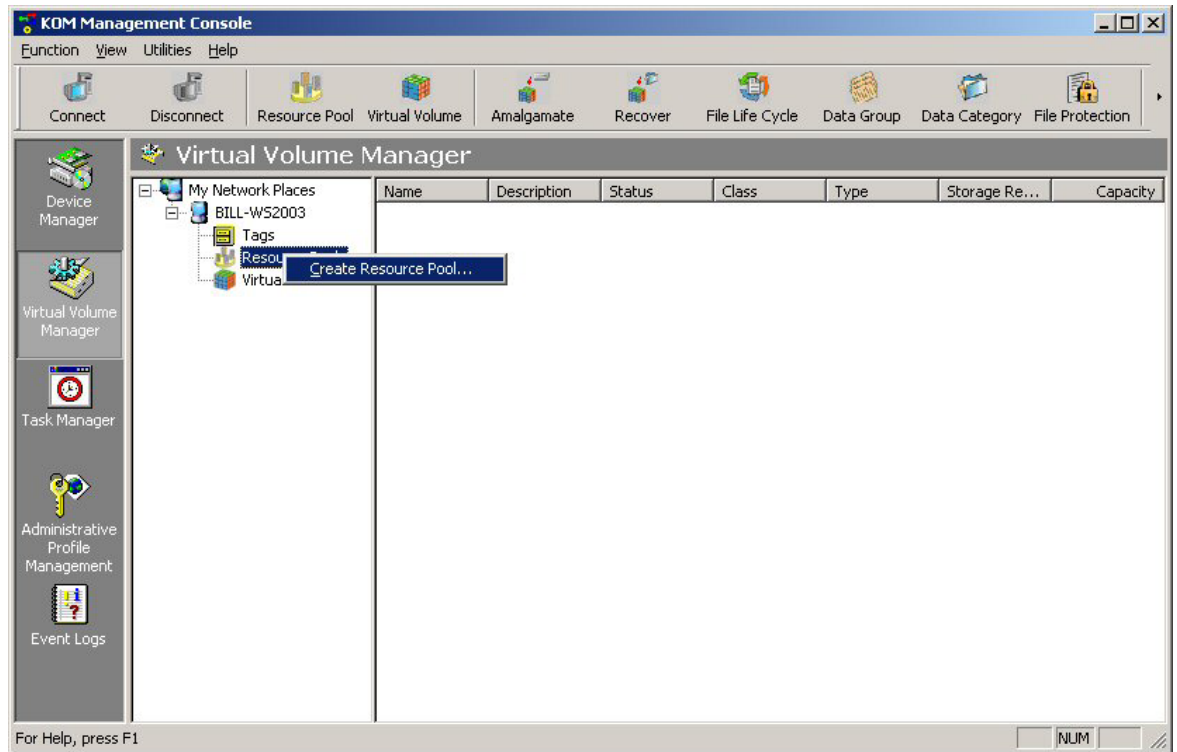


2.1 Create a Storage Resource Pool for each storage type.

Select the **Create Resource Pool**.

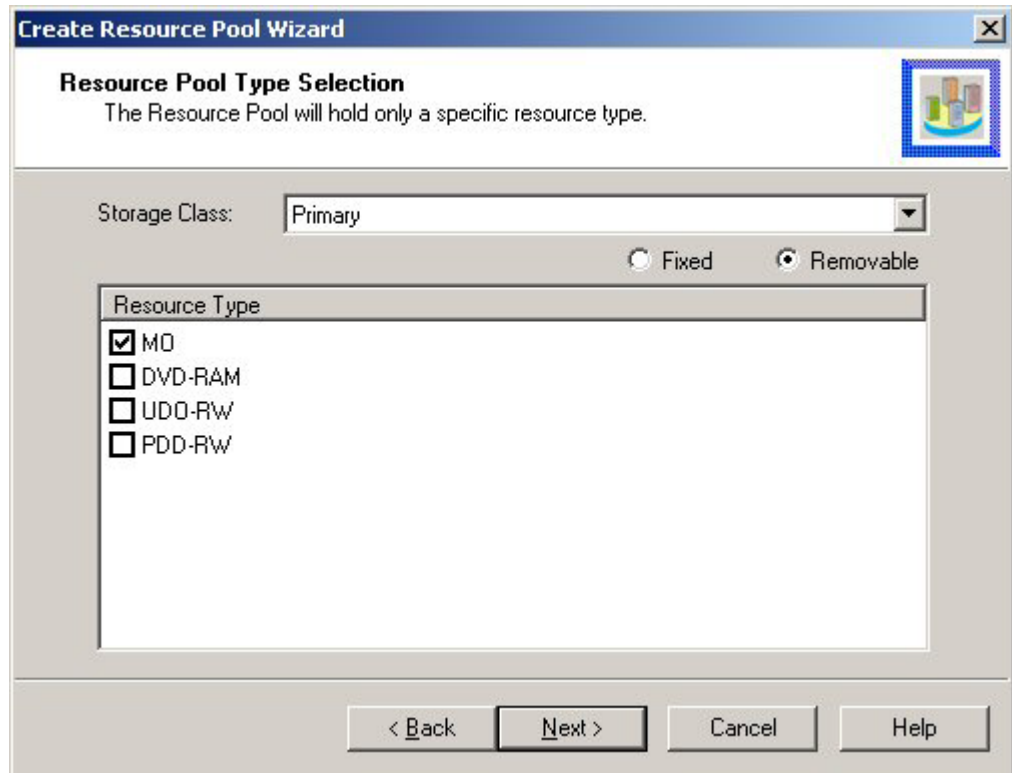
Figure 11 provides guidance.

Figure 11. Create Resource Pool



The Create Resource Pool Wizard will launch. Invoke the wizard three times; add three storage pools, one for each type of storage. Be sure to click **Removable** on the MO and clear the check for all but the **MO** as shown in figure 12.

Figure 12. Resource Pool Type Selection

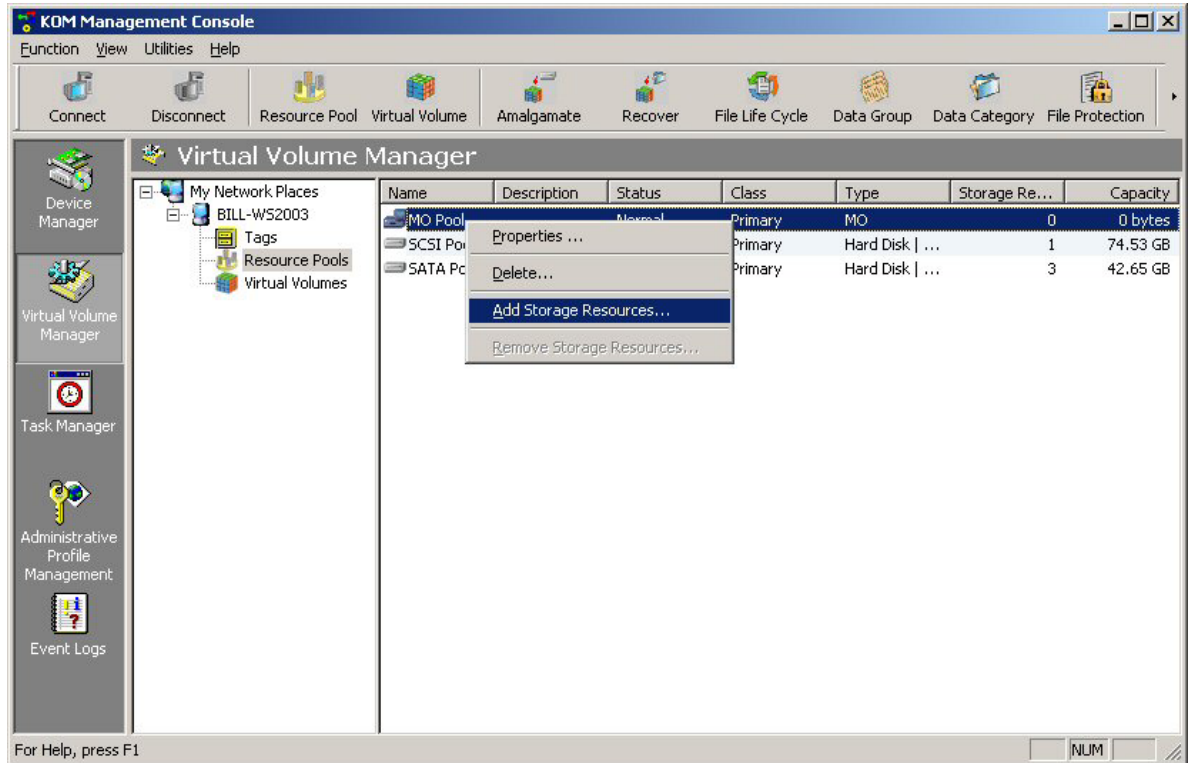


3. Assign raw storage to resource pools.

Right-click each storage pool individually and select **Add Storage Resources**.

This will launch the Add Storage Resource Wizard. Figure 13 provides guidance.

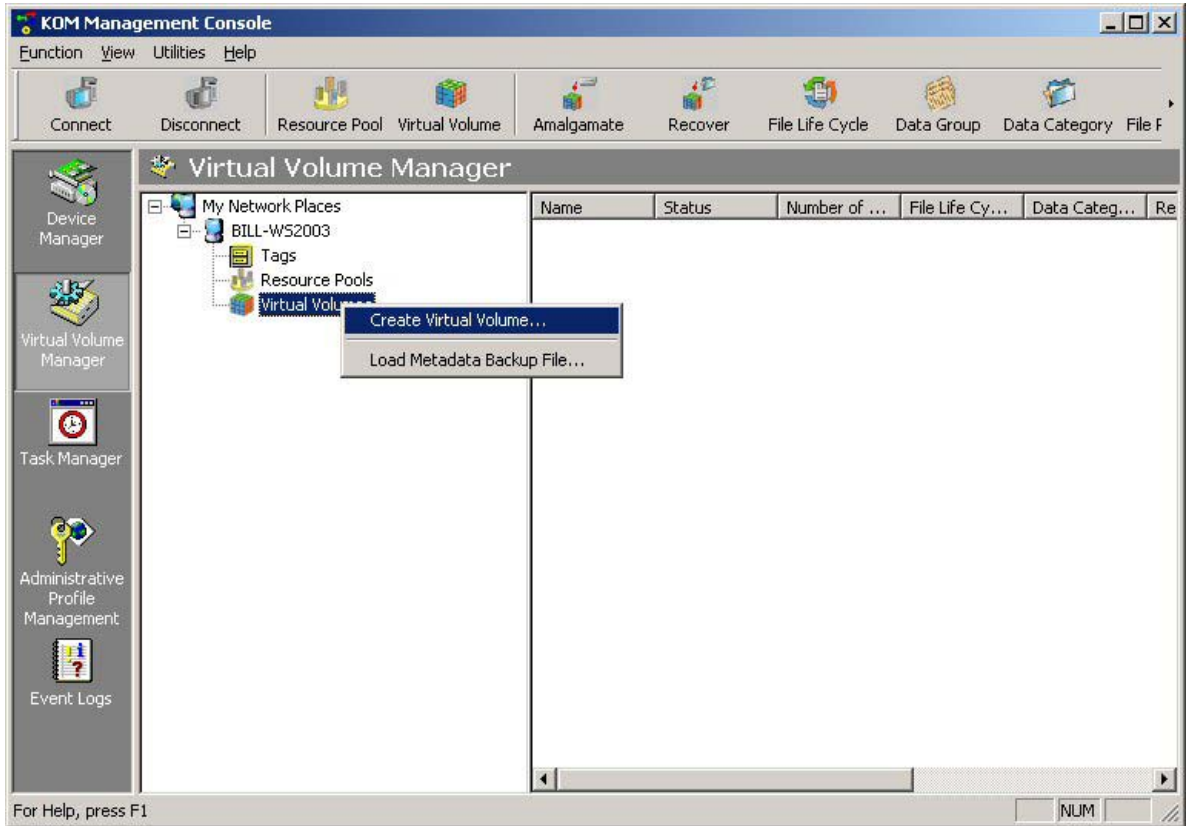
Figure 13. Add Storage Resources



4. Create a Virtual Volume.

Right-click the **Virtual Volume** and select **Create Virtual Volume**. Figure 14 provides guidance.


Figure 14. Create Virtual Volume



Provide a volume name. Use the Metadata volume created in step 4.1 of the Software Installation. Specify **SCSI** as a new default storage group. Enable the resource as **Shareable**, allow **Spillover**. Specify the SCSI group as the owner. **Create change log** was not enabled.

Figure 15 shows the final notice before the creation of the Virtual Volume.

Figure 15. Final notice



The creation of the Virtual Volume will destroy the contents of the volume used for the Metadata store.

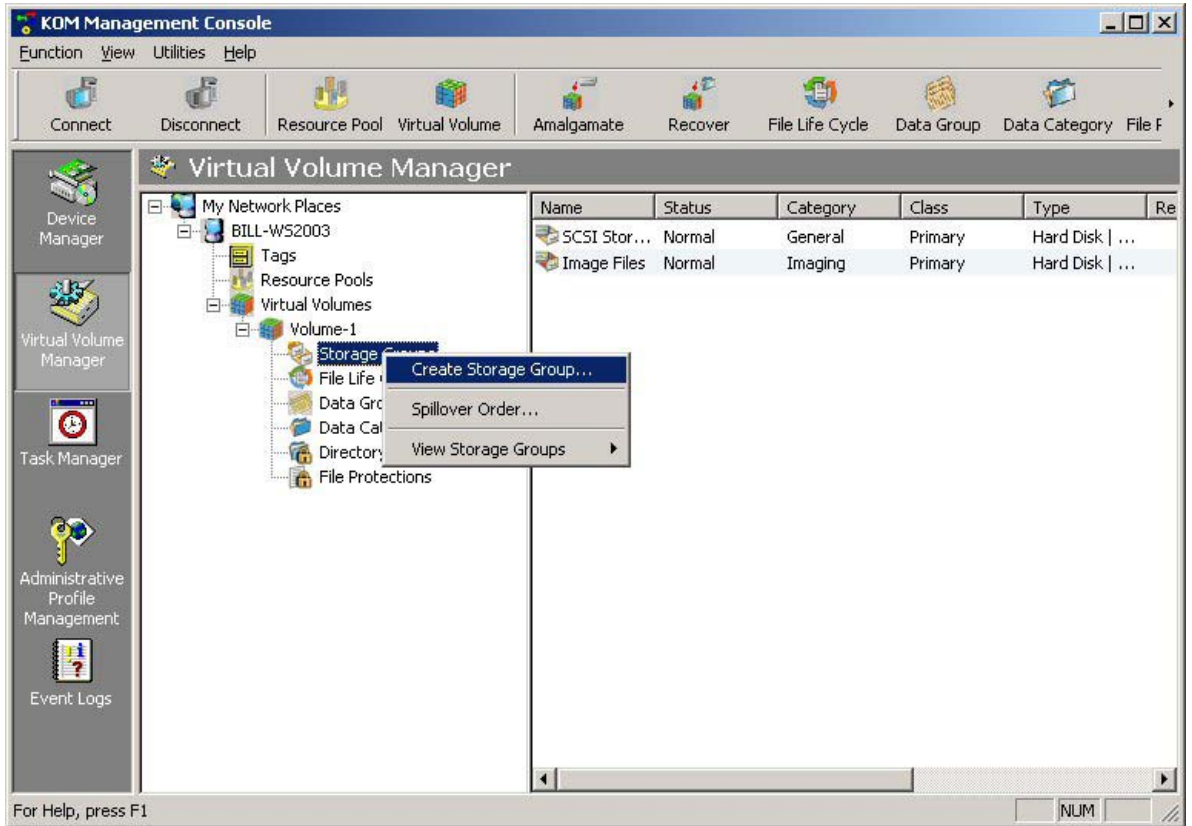
Click **OK**

5. Create the remaining storage groups.

Two additional storage groups were created, one for **SATA** and one for **MO**.

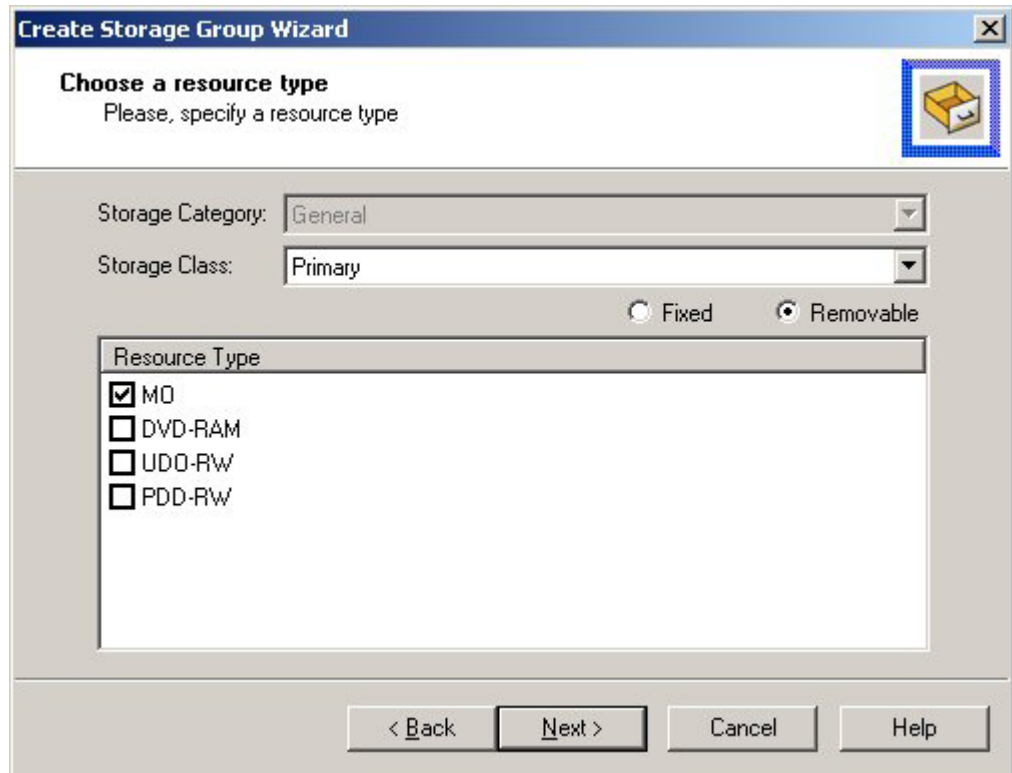
Right-click **Storage Groups** and select **Create Storage Group**. Figure 16 provides guidance.

Figure 16. Create Storage Group



The Storage Group Creation Wizard will be launched. Select the ownership of each respective group. Special care for the MO storage group is required to select only MO as the resource type and to specify removable storage. Figure 17 provides an illustration.

Figure 17. MO resource type



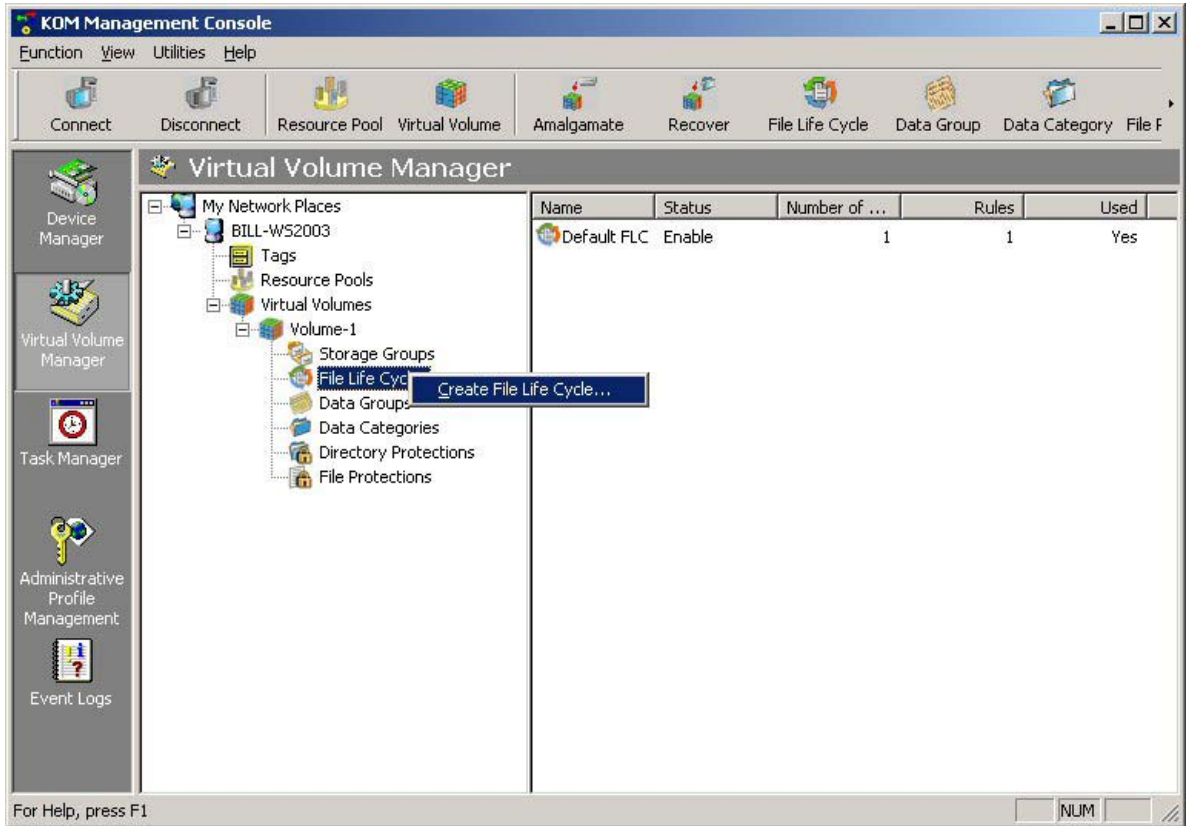
The defaults were accepted for the SCSI and SATA groups.

6. Create the File Life Cycle.

The simple File Life Cycle policy tested migrates files based on the elapsed time since they were last modified. Files that were modified less than 6 months ago will stay on the fastest SCSI storage. Files that were modified between 6 months and 5 years will be migrated to the SATA storage. Files that were modified 5 years or more ago will be migrated to the stable storage of the removable MO.

Right-click **File Life Cycles** and select **Create Life Cycle**. Figure 18 shows the procedure.

Figure 18. Create File Life Cycle



The File Life Cycle Creation Wizard will be launched. Specify a name **FLC-1**. Select **File Age** as the criteria for the storage rule. See Figure 19.

Figure 19. Choose Storage Rule Criteria

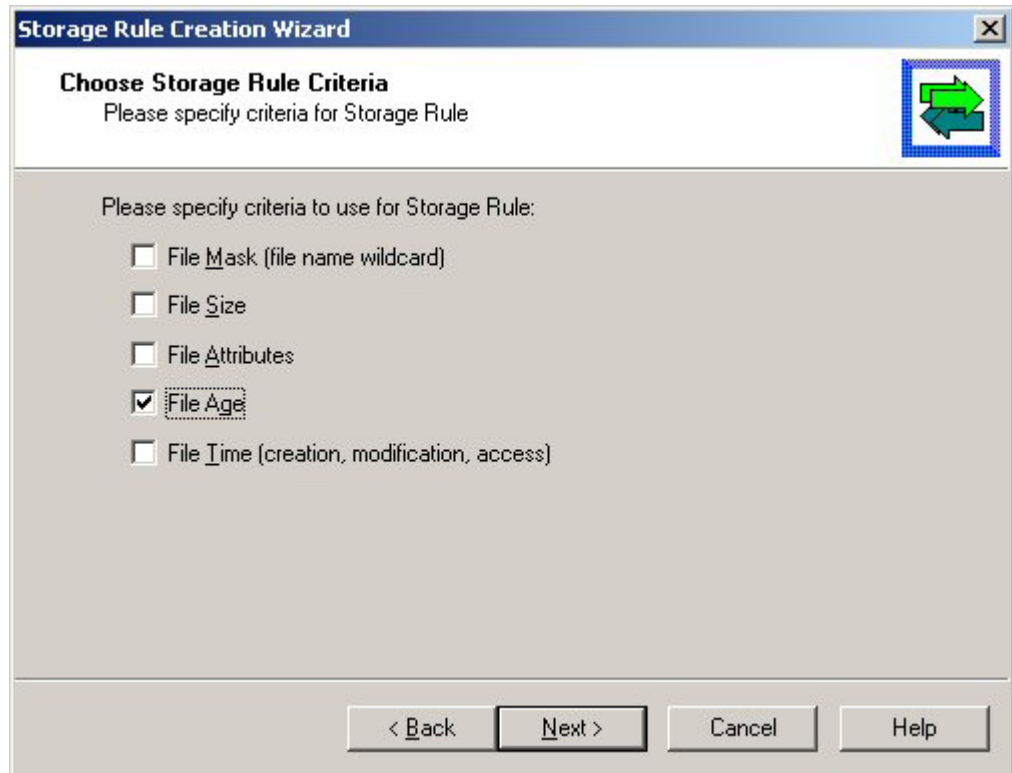
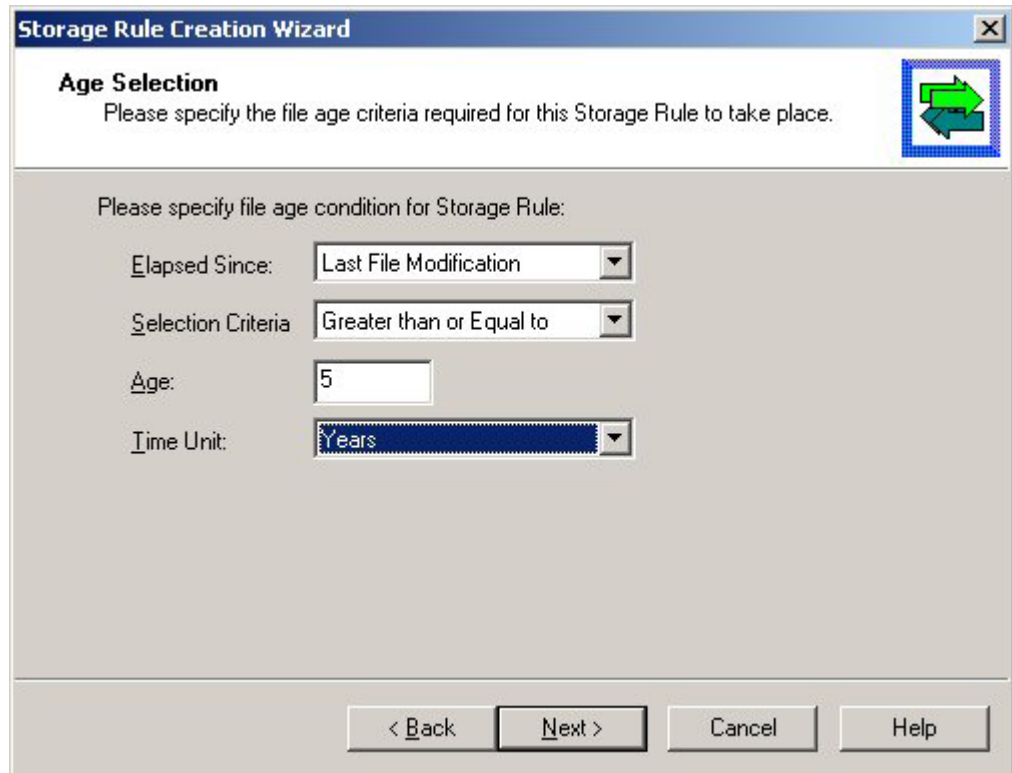


Figure 19 shows the criteria that may be specified for a specific policy one may choose to create. This shows the data centric nature of the storage management software.

The Age Selection was configured as shown in Figure 20.

Figure 20. Age Selection



The Target was specified as MO.

A similar policy was created for files that were last modified greater than six months earlier and **SATA** storage was the specified target.

Logical groups of files may also be created by the use of Data Categories. In this way security policies may be defined. The system can make use of group and end even user attributes. Refer to KOM Networks documentation for an exhaustive review of the system capabilities.

7. Share the Virtual Volume over the network

Click **Start/All Programs/My Computer**

Navigate to the volume created in step 4. Right click the virtual volume and select **sharing and security**. Share the volume utilizing common methods.

8. Test the Configuration.

Informal procedures were utilized by copying files and directories containing files into the virtual volume. Review of the physical storage subsystems, utilizing the KOMworx management console, reflected the expected file locations resulting from the policies created. The virtual volume responded quickly and from all practical expectations appeared to be a common NTFS volume but with extensible policy capabilities. Files migrated quickly to policy changes or to file attribute changes.

The importance of an optimized configuration can not be understated. While hardware and software alone have specific value the successful deployment depends on a balance between three areas: hardware, software, and professional services. Some companies leverage internal resources as an effective strategy for professional services. Others depend on service organizations to provide this aspect of the solution. Each business should review this balance to determine the most cost effective strategy for the long term total cost of ownership.

HP Services

Successful cost efficient technology deployment depends on a viable plan and effective execution. While every attempt to ensure the accuracy of this document has been taken, one cannot predict the complications that may be introduced into the system by the installation and configuration of the entire system. Issues that introduce unforeseen complications would include but not be limited to:

- New or different hardware selection
- Updated ROM, driver, application software and OS versions
- Different sizing and configuration requirements

Some technology may be adopted by simple installation. Other technology requires a blend of design and management oversight. A cost effective ILM system requires consideration of many aspects of the business it serves. A file retention system is but one component of the entire ILM solution a given business may require.

HP Services has helped thousands of enterprises meet business and IT goals by delivering high-quality global IT solutions with reduced risk at the right price. Whether customers are launching a major IT initiative or seeking ongoing support, HP can help them design, implement and manage a solution using proven tools and processes and flexible sourcing of skills from our extensive network of 65,000 service professionals in 160 countries.

Glossary – Storage Management

File retention

File retention addresses the management of a document at the file level. An institution may encompass many different departments including labs that create files as the output from their daily processes. File retention addresses the reliable storage of these documents. An important tenet of file retention is the concept of WORM.

HSM (Hierarchical Storage Management)

HSM has existed for over a two decades in the mainframe computing environment. This technology provided for the simple migration of files based on the age of the file stored. Multiple types of physical storage may be virtualized and presented as a single unified file system to the application or end user. These legacy systems typically simply concatenated disk and tape into the virtualized storage. HSM systems have limited capacity to incorporate any attribute other than file age into the migration logic storage management system.

Information Lifecycle Management

ILM, as a practice, refers to the discipline of creating procedures and architecting storage management systems that address comprehensive management of data. File retention is but one aspect of ILM.

ILM, as a requirement, results from the goal to manage data as it created, stored, retrieved, modified, and archived. ILM responds to the need to manage data to achieve various business goals. Management of ILM can reduce costs by efficiently utilizing the optimum physical storage medium for any given data that needs to be stored, retrieved and archived, while complying with legal requirements, performance, security mandates, and business requirements.

WORM

While physical media may be WORM (Write Once Read Many) this is true at a logical level as well. A storage management system maybe designed to provide WORM functionality at the system level.

Glossary – Storage Technology

FC

A FC (fibre channel) connection provides a fiber optical storage interconnect between computers and storage subsystems much the same way that SCSI is utilized. FC provides higher speeds, greater distances and increased functionality than available with SCSI.

HBA

A Host Bus Adaptor provides a fibre optical connection to the server much the way a SCSI host bus adaptor interconnects a server to a storage device.

MO storage media

A Magneto Optical storage media provides a stable media on which to store data permanently. Consumer technologies such as CD-ROMs or DVD-ROMs simply can not provide stable storage. MO media is protected by a hard plastic cartridge that surrounds the media, protecting it from dirt and normal handling in both automated and manual procedures. The method applied to write data includes the use of both thermal, red laser and magnetic techniques. This ensures that data is stable well beyond room temperature and that in the process of reading the data with laser alone will not alter data in any way. Physical media may be WORM (Write Once Read Many) or rewritable.

Optical Jukebox

An Optical Jukebox is a self contained automated optical media library. These systems can automatically load and unload optical media into the one or more internal drives that read and write the optical media. Optical jukeboxes may contain MO or UDO drives and media. Without storage management software that provides the necessary directives as to which media to load or unload the device must be manually instructed on which platter to load or unload.

WORM

Physical optical media may be WORM (Write Once Read Many) or rewritable. Policies may dictate a combination of both rewriteable and WORM media depending on strategy.

RAID

A RAID or Redundant Array of Independent Disks is a method of using a group of disks to scale a storage subsystem to increase the speed and/or the reliability of the group of disks. Concatenated capacity may be logically provisioned.

SA

A SA (Storage Array) physically and logically manages a RAID of disk drives providing higher performance and reliability than available from unmanaged hard drives directly connected to a computer. Separate servers may access a storage array through a SAN.

SAN

A SAN (Storage Area Network) is typically based on fibre channel connected storage devices and the servers used to access the storage subsystems. Other components that

provide storage management may also be connected. SAN switches are used to create fibre channel storage area networks.

SATA Hard drive

The Serial Advanced Technology Attachment Interface standard was ratified in November of 2000 and was based on the popular parallel Advanced Technology Attachment specifications. SATA hard drives are accepted as the newest low cost leader in the hard drive in the industry.

SCSI

The Small Computer Systems Interface standard has evolved over several decades and provides the physical and logical signaling protocols to interconnect storage devices and computers. Many different storage devices and computer peripherals depend on the SCSI interface.

SCSI Controller

The SCSI controller provides physical interface between a server and a SCSI storage device, such as an optical jukebox.

SCSI Hard drive

SCSI hard drives are a family of high performance hard drives based on the use of SCSI interface standard. Hard drives based on this interconnect are broadly accepted as high reliability and high performance.

UDO storage media

UDO (Ultra Dense Optical) is the latest evolution in optical storage. UDO has approximately three times the capacity of MO due to the use of the higher frequency blue LASER light. UDO storage media provides a much more stable media on which to store data than popular consumer CD-ROMs or DVD-ROMs. UDO media is protected by a hard plastic cartridge that surrounds the media protecting it from dirt and normal handling in both automated and manual procedures. Physical media may be WORM (Write Once Read Many) or rewritable.

Summary

This document provides an introduction to the concept of a file-centric storage management system providing file retention compliance using KOM Networks software on Microsoft Windows 2003 operating system and provides the reader with a proven HP and vendor tested configuration. The resulting system may be customized and refined to the specific requirements of the reader. HP ProLiant servers, HP StorageWorks Modular Smart Arrays, and HP StorageWorks storage area network switches provide a comprehensive proven hardware platform by which the reader can architect storage system that provides a simple, reliable, and cost effective solution.

The resulting file retention storage system, when viewed as a unit, provides the benefits and performance of proprietary monolithic systems that are much more expensive to acquire and maintain.

© 2003 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group.

[11/2003]-1

