



DXi-Series Configuration and Best Practices Guide

For IBM Spectrum Protect™ / TSM

Quantum: 20020436-011 Rev E

BPG00014A-05



Table of Contents

DXi-Series Configuration and Best Practices Guide for Tivoli® Storage Manager / Spectrum Protect™	6
How to Use This Guide	6
Shortcuts to Quick Start Activities	6
Documentation and References	7
<i>Online Documentation for Your Quantum Products</i>	7
DXi-Series Management Console & User's Guides	7
Virtual Tape Library Setup	7
Network Attached Storage.....	7
DXi Replication	8
TSM Documentation	8
<i>Summary of Tuning Parameters for IBM TSM</i>	9
Configuring TSM with the DXi-Series	10
<i>Configuring TSM with DXi VTL</i>	10
Robot/Media Changer Device Serialization Considerations.....	10
Install/Configure Library Device Driver(s) as Required by TSM.....	11
DXi Tape Drive Emulations and Recommend Drivers	11
Configure the DXi for VTL	11
Configure VTLs on a DXi.....	12
Configure TSM to Use VTLs.....	13
<i>Best Practices Guide with DXi VTL</i>	14
General Installation and Setup Considerations	14
Virtual Tape Drive Device Description and Recommendation	14
Device Driver and Firmware Level	14
Number of Concurrent Tape Drives in Use	14
Tape Drive LUN Mapping	15
Tape Cartridge Capacity Considerations.....	16
TSM Trigger Setting and Cartridge Size.....	16
VTL Fibre Channel Performance Tuning.....	17
Considerations for the Handling of Expired Media within TSM	17
Reclamation.....	18
<i>TSM Reclamation</i>	18
<i>DXi Reclamation</i>	18
Additional Best Practice Considerations.....	18
<i>Configuring TSM with DXi NAS</i>	19
NAS Device Path Considerations.....	19
Configure the DXi for NAS.....	20
Configure NAS Backups with TSM.....	20
<i>Best Practices Guide with DXi NAS</i>	21
Number of Shares Considerations	21
Network Share Access Control Considerations	21
Network Considerations.....	21
TSM NAS Settings and Tuning Considerations	22
Additional Best Practice Considerations.....	22
<i>Common Operational Considerations for TSM</i>	23
Data Deduplication Considerations	23

<i>Good Candidates for Data Deduplication</i>	23
<i>Not So Good Candidates for Data Deduplication</i>	23
Replication Considerations.....	23
Oversubscription of Space on the DXi.....	24
Space Reclamation.....	24
DXi Multiprotocol Guidance - NFS/VTL Scenario.....	26
Appendix A: Additional TSM Information	28
TSM Driver Information.....	28
The Transaction Byte Limit Client Option.....	29
RELABELSCRATCH Information.....	29
General Virtual Tape Library Considerations for Tivoli Storage Manager Servers.....	29
Appendix B: TSM Common Tasks - Command Line	30
Starting TSM to Get to the Command Prompt (UNIX).....	30
Commonly Used Command Line commands.....	30
Configure the Library with TSM.....	30
Register Client (Host) to TSM Server.....	31
Label Multiple Media via TSM.....	31
Inventory Library via TSM.....	31
Format/Label Single Media.....	31
Format/Label Multiple Media.....	31
Quick Erase Media.....	31
Back Up the Database.....	31
Local Backup.....	31
Set Up Scheduled Backup.....	32
Media Clone/Media-to-Media Copy.....	32
Extend Database and Recovery Log Files.....	32
Removing a Library to Replace It with a Different One.....	32
Helpful Resources.....	34

DXi-Series Configuration and Best Practices Guide for Tivoli® Storage Manager / Spectrum Protect™

Note: As of Version 7.1.3, IBM Tivoli® Storage Manager (TSM) is now called Spectrum Protect™. All references to Tivoli Storage Manager or TSM in this Best Practices Guide also refer to Spectrum Protect. On the IBM Web site, some references to this product, and some hyperlinks, may still refer to Tivoli Storage Manager / TSM, until the IBM rebranding effort is completed.

This guide seeks to help Quantum customers who own DXi-Series systems (DXi4000-Series, DXi6000-Series, DXi8000-Series, and DXi9000-Series), and who also use Tivoli Storage Manager, get the most out of their investment. It is also intended to help Quantum field sales teams by providing guidance to enhance the installation and integration of TSM with Quantum DXi-Series systems. This guide includes advice and best practices for using Quantum DXi-Series systems with TSM.

How to Use This Guide

This document assumes that the reader has basic expertise with TSM, as well as basic networking and SAN experience. It also assumes that the reader has a Quantum DXi installed in a working TSM environment.

This document provides key recommendations and useful information for quickly setting up a DXi system with TSM. It expands on these recommendations and discusses the features and performance tuning considerations relevant to various storage access methods.

This document is organized according to the various storage target access methods to be employed with TSM. Depending on the DXi model, the DXi can appear as a Virtual Tape Library (VTL) storage device over Fibre Channel (FC), or as a Network Attached Storage device (NFS or CIFS). These access methods are discussed in the following order.

- DXi VTL
- DXi NAS - NFS and/or CIFS

Shortcuts to Quick Start Activities

To go directly to any of the following sections, click that section's name.

- » [Online Documentation for Your Quantum Products\(p.7\)](#)
- » [Summary of Tuning Parameters for IBM TSM \(p.9\)](#)
- » [Configuring TSM with DXi VTL \(p.10\)](#)
- » [Best Practices Guide with DXi VTL \(p.14\)](#)
- » [Configuring TSM with DXi NAS \(p.19\)](#)
- » [Best Practices Guide with DXi NAS \(p.21\)](#)
- » [Common Operational Considerations for TSM \(p.23\)](#)

Documentation and References

The following is a list of documents, references, and links where you can find additional information regarding specific activities and products. Access to many of the documents below requires a valid serial number. Please have that available when following the hyperlinks to the documents.

Online Documentation for Your Quantum Products

<http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/Index.aspx#>

In the lists below, if multiple versions of a manual for your DXi are listed on the linked page, scroll down and select the version for your hardware/software combination. Please note that documents are not listed in alphabetical order, so two versions of the same manual will not necessarily be listed in consecutive lines.

DXi-Series Management Console & User's Guides

The DXi-Series Management console helps you configure and use your storage solution. Refer to the following documents for more information on DXi-Series Management:

- [DXi9000 User's Guide](#)
- [DXi8500 User's Guide](#)
- [DXi6900 / DXi6902 User's Guide \(go to **Previous Documentation > DXi6902 User Guide** -same doc covers both models\)](#)
- [DXi6800 User's Guide](#)
- [DXi6700 User's Guide](#)
- [DXi4800 User's Guide](#)
- [DXi4700 User's Guide \(go to **DXi4700 Print Documentation > Previous Documentation > DXi470x User Guide**\)](#)
- [DXi4000 User's Guide](#)
- [DXi V-Series User's Guide](#)

Virtual Tape Library Setup

Refer to the following documents for VTL setup:

- [DXi9000 User's Guide](#)
- [DXi8500 User's Guide](#)
- [DXi6900 / DXi6902 User's Guide \(go to **Previous Documentation > DXi6902 User Guide** -same doc covers both models\)](#)
- [DXi6800 User's Guide](#)
- [DXi6700 User's Guide](#)
- [DXi4800 User's Guide](#)
- [DXi4700 User's Guide](#)

Network Attached Storage

Refer to the following documents for NAS Share setup:

- [DXi9000 User's Guide](#)
- [DXi8500 User's Guide](#)
- [DXi6900 / DXi6902 User's Guide \(go to **Previous Documentation > DXi6902 User Guide** -same doc covers both models\)](#)
- [DXi6700 User's Guide](#)
- [DXi4800 User's Guide](#)
- [DXi4700 User's Guide \(go to **DXi4700 Print Documentation > Previous Documentation > DXi470x User Guide**\)](#)
- [DXi4000 User's Guide](#)
- [DXi V-Series User's Guide](#)

[DXi Replication](#)

Refer to the following documents for DXi-to-DXi Replication setup:

- [DXi9000 User's Guide](#)
- [DXi8500 User's Guide](#)
- [DXi6900 / DXi6902 User's Guide \(go to **Previous Documentation > DXi6902 User Guide** -same doc covers both models\)](#)
- [DXi6800 Installation and Configuration Guide](#)
- [DXi6700 User's Guide](#)
- [DXi4800 User's Guide](#)
- [DXi4700 User's Guide \(go to **DXi4700 Print Documentation > Previous Documentation > DXi470x User Guide**\)](#)
- [DXi4000 User's Guide](#)
- [DXi V-Series User's Guide](#)

[TSM Documentation](#)

TSM / Spectrum Protect documentation is available at the following Web sites:

[Tivoli Storage Manager documentation](#)

[IBM Spectrum Protect documentation](#)

IBM's documents include Installation, Upgrade and Migration, and Administrator's guides, and Administrator's References for AIX, HP-UX, Linux, Solaris, and Windows.

See the [IBM Spectrum Protect Supported Devices for AIX®, HP-UX, Solaris, and Windows®](#) Web page, or the [IBM Spectrum Protect Supported Devices for Linux](#) Web page, for further information.

Summary of Tuning Parameters for IBM TSM

For backup administrators who are well versed on IBM TSM and Quantum DXi systems, the following table offers a summary of suggested parameters/values. As with any modifications that impact performance and/or tuning, results will vary for each individual environment.

Parameter or Option	Setting
Compression	No
Encryption	No
Deduplication	No
TDP Multiplexing	No
Dynamic Tracking	Disabled (applies to AIX only)
TSM Tape Reclamation Trigger	50%-90%, depending upon backup load (See the TSM Trigger Setting and Cartridge Size section on p.16.)
ReUseDelay	0
TXNBYTELIMIT	25600 2097152 for backup direct to tape (For detailed information, see the section entitled The Transaction Byte Limit Client Option on p.29 in Appendix A: Additional TSM Information.)
Miscellaneous Options	Recommendations
Server Resources	Memory Requirements for All Operating Systems: <ul style="list-style-type: none"> • 12 GB. • 16 GB if you are using deduplication. • At least 32 GB for heavily used servers. Using 32 GB or more of memory enhances performance of the Tivoli Storage Manager Server database inventory. • For multiple instances, each instance requires the memory listed for one server. Multiply the memory for one server by the number of instances planned for the system.
Server Name	The server name is an internal name for Tivoli Storage Manager, and is used for operations that involve communication among multiple Tivoli Storage Manager servers. Examples include server-to-server communication and library sharing. The <i>Server Name</i> is different from the <i>host name</i> . However, the default server name that is suggested is the same as the host name of the system that you are using.
Database Name	The database name is always TSMDB1 , for every server instance. This name cannot be changed.

Configuring TSM with the DXi-Series

IBM Tivoli Storage Manager is based on client/server architecture. A TSM domain is a logical group of policies and clients within a TSM server. The TSM server controls and directs all TSM operations. Backup management information is stored or managed by the TSM server, including configuration information, parameters for client systems to be backed up, the index of files backed up from those clients, and media volumes where backed-up data is stored. Normally, a client operates under the control of the TSM server, according to the policies an administrator establishes.

When a backup occurs, a TSM client sends data to the TSM server, which directs the data to the appropriate storage pool. The TSM server receives metadata information about the files being backed up separately and makes appropriate entries in the TSM database.

In traditional TSM environments, the primary storage pool is a minimal amount of high performance disk, either direct-attached or SAN-attached to the TSM server. Once the amount of data in the primary storage pool reaches a threshold, the data is copied to a secondary storage pool, typically a tape library, for longer term storage.

Note that the DXi's storage is shared by all NAS shares and VTLs configured in the system.

Configuring TSM with DXi VTL

A Virtual Tape Library is a data storage virtualization technology used for backup and recovery, in which disk storage is virtualized as tape storage. A VTL presents itself as a tape library with tape drives for use with existing instances of TSM. Because VTLs emulate all the SCSI capabilities of tape hardware, VTL usage is not apparent to Tivoli Storage Manager Servers, regardless of the server level or platform.

Virtualizing disk storage as tape allows integration of VTLs with existing TSM software and existing backup and recovery processes and policies. The benefits of such virtualization include storage consolidation and faster data restores. Creating a backup image on a virtual tape is no different than creating a backup image on a physical tape — the backup functionality remains unchanged. The VTL is just standing in for a physical tape library as the secondary storage pool for TSM.

Robot/Media Changer Device Serialization Considerations

One of the key ways to ensure that SAN-connected physical and virtual tape libraries are detected properly by backup servers is *serialization*. Serialization provides a unique identifier for each device in a physical or virtual tape library, to automate device association from multiple backup servers. These identifiers, returned by the VTL devices, are separate from the *element addresses* that define the position of devices in the library. The element address is used by the library's robot or medium changer to manage the tape drives.

Serialization allows the servers running the data protection application (the media servers) to coordinate tape drive configuration by aligning the device serial number with the device's element address. This enables TSM device discovery to align these two addresses, reducing the potential for improper configuration.

If the device or configuration manager does not serialize the devices listed, **DO NOT** commit the changes, and be sure to check the VTL's online state. The DXi VTL partition must be online for this to function properly. The Quantum-recommended device identification for each DXi system is as an i6000.

Install/Configure Library Device Driver(s) as Required by TSM

Always ensure that the latest drivers and firmware have been installed for the tape drives, and for the media changer and HBA. For best performance, drivers should be loaded for the tape drive.

Tivoli provides its own device driver for the media changer. Use the 'for Tivoli' driver for the media changer, but not for the tape drives. Non-IBM tape drive emulations will also use the Tivoli Storage Manager driver. For more information, see the IBM Web page [Attaching devices for the server](#).

For OEM-Branded IBM tape drives, use the regular IBM 'non-Tivoli' driver — not the one included in Tivoli.

Beginning with Tivoli Storage Manager Version 6.3, the Tivoli Storage Manager device driver uses persistent reservation to reserve the following physical and emulated tape drives:

- HP LTO-3
- HP LTO-4
- HP LTO-5

By default, the IBM tape device driver uses persistent reservation for IBM tape drives.

Note: AIX displays a message that it “needs **devices.fcp.changer** or **devices.scsi** Fileset” when trying to use a non-IBM device with Tivoli Storage Manager Server running on AIX. The message is generated because the OS is seeing a device, but does not recognize what it is because there is no driver for it loaded into the AIX Kernel.

The Tivoli Storage Manager device driver does not load into the AIX Kernel, so you can expect this message for any devices that are being used with Tivoli Storage Manager that the OS does not already have its own driver for. You can safely ignore and bypass this message.

When configuring devices on a different SCSI or FC adapter, run 'cfgmgr' with the '-l' option to restrict it to looking at a single bus, rather than searching the entire system, which it does by default. e.g. 'cfgmgr -l scsi0'.

DXi Tape Drive Emulations and Recommend Drivers

If you are using IBM tape devices with Tivoli, IBM device drivers must be loaded. Tivoli device drivers will not work with IBM tape devices. This includes DXi virtual IBM drives. For specific driver information, refer to the [TSM Driver Information](#) section on p.28 in [Appendix A: Additional TSM Information](#).

Note: Quantum does not support mixing drive types in a single partition.

Configure the DXi for VTL

Installing and configuring the DXi and TSM for VTL operation consists of the following steps:

1. [Configure VTLs on a DXi](#)
2. [Configure TSM to Use VTLs](#)

Configure VTLs on a DXi

The following steps describe how to configure a DXi with VTLs:

1. Create a virtual tape library on the DXi, specifying the emulation type, number of slots, number of tape drives, and their emulation type. If desired, multiple VTL's can be created to service different backup servers.
2. Create virtual tape cartridges on the DXi for the VTL, specifying the cartridge type, quantity, and capacity.
3. Assign the VTL to a client on the SAN to which the DXi is attached.

In the DXi Remote Management Console, under the **Configuration** tab, the VTL page allows you to configure a DXi to present its storage capacity as VTL partitions that are compatible with TSM and other standard backup applications. You may add virtual tape drives and storage slots to VTL partitions to create and work with virtual tape cartridges, then map partitions to hosts.

Partitioning lets you divide the DXi virtual tape drives and storage elements into separate VTLs, usable by separate host computers. The **Partitions** page contains a list of assigned tape drives, as well as listing all user-defined VTL partitions that are currently configured on the system. This page also lets you add, edit, and delete partitions.

The **Summary** page displays the maximum number of VTL partitions, the total number of tape drives, and the number of assigned tape drives. The **Summary** page also provides a list of configured partitions on the system. Click the link in the **Name** column to edit a specific partition.

Caution: You can use a VTL for any virtual tape library only when the following conditions are true:

- There is no mixed media involved in the VTL. Only one type and generation of drive and media is emulated in the library.
- Every server and storage agent with access to the VTL has paths that are defined for all drives in the library.

If either of these conditions is not met, any mount performance advantage from defining a VTL library to the Tivoli Storage Manager server may be reduced or negated.

Note: If you are planning to replicate partitions to another DXi system, you must ensure that every partition name and barcode number on the system is unique. You **CAN NOT** have duplicate partition names or barcode numbers on a DXi system, or on a system receiving a replicated partition.

The **Create Media** page allows you to create virtual media for a specific partition. Once created, these virtual cartridges are available for backing up data. You can configure the media type, capacity, starting barcode, and initial location on this page.

Note: It is possible to oversubscribe space on the DXi system. The sum total of capacity for all media could be more than the capacity of the system. Be careful to ensure that this does not happen.

Configure TSM to Use VTLs

The following steps describe how to configure TSM with VTLs:

1. When a DXi is configured as a VTL, it will present itself to TSM as a standard tape library with drives and cartridges. Scan for devices at the TSM server. The VTL and its tape drives will be discovered.
2. In the Tivoli Integrated Portal Web site on the media server, go to **Storage Devices**. Then:
 - a. Create a library, specifying the Library Type of the VTL.
 - b. Add drives to the library one at a time.
3. Label and check in volumes:
 - a. Choose whether to automatically label volumes, or to do it manually.
 - b. Add volumes by searching for all eligible volumes in regular slots in the library, and accept all defaults in the dialog box. Then, discover volumes.
 - c. Create the appropriate Device Class (drive type), and select the library created in step 2.
 - d. Create a Storage Pool, specifying the following:
 - Name: <name of storage pool>
 - Type: Sequential access
 - Device Class: <use device class created in step c>
 - Scratch volume: <use maximum value>

Refer to the section on “Sequential Access Storage Pool Volumes” in the IBM Tivoli Storage Manager Administrator’s Guide for more information on creating storage pools of this type.

4. Activate the policy on the storage pool just created.
5. Perform a test backup to verify the configuration.

Best Practices Guide with DXi VTL

General Installation and Setup Considerations

- Ensure that all non-existent tape drives are cleared from the ODM database.
- If installing TSM on AIX, ensure that it is acceptable to reboot the AIX server at least one time during implementation.
- Configure the VTL as a Scalar library.
- Ensure that a device addressed at LUN0 is present, and that **there are NO GAPS** in the LUN numbering schema following LUN0.
 - If there is no device at LUN0, device discovery will abort.
 - If there is a gap in the LUN numbering, device discovery will not find devices with LUN numbers beyond the gap.

Virtual Tape Drive Device Description and Recommendation

Quantum DXi products support multiple tape drive emulations and allow users to define the capacities of tape cartridges, to support the drive type. The drive type selection does not internally impose any throttling of the ingest rate, or any limitations on the virtual cartridge capacity. The host system uses tape device drivers to communicate with the virtual tape drives and the TSM application. The qualified tape drive emulations published on the hardware compatibility list are IBM LTO2, LTO3, LTO4, and LTO5, as these drives have the widest support across platforms.

Quantum recommends using the IBM tape drive drivers for backup hosts running Windows 32-bit operating systems. For backup hosts running 64-bit versions of the Windows operating system, the OEM tape device drivers are recommended. The complete tape drive driver recommendations from IBM can be located in IBM's hardware compatibility list ([IBM TSM Hardware Compatibility List](#)).

Quantum recommends against using the Windows update function to get the latest tape drive drivers, as it is possible that Windows Update can find unexpected driver matches for a tape drive. These unexpected matches often are not the same as those supplied by the tape drive manufacturer and may introduce unexpected incompatibilities. This can result in a change to the device in the Windows Device Manager. Tape device drivers from IBM can be found at the following IBM TSM Information Center Web page: [Attaching devices for the Server](#).

Device Driver and Firmware Level

Ensure that the latest drivers and firmware have been installed for the tape drive, robotic library, and HBA. Drivers should be loaded for the tape drive for best performance. See this IBM Web page for more information: [Attaching devices for the Server](#).

Number of Concurrent Tape Drives in Use

Each DXi model has a maximum number of virtual tape drives that can be configured. Each model also has a maximum aggregate throughput rate, which will be divided relatively equally between the virtual tape drives in use, assuming there are no server or network bottlenecks. However, this does not prohibit a single tape drive from using all available bandwidth.

Keep in mind that increasing the number of concurrently active virtual tape drives does not increase the aggregate DXi bandwidth. It could result in a failed backup job, due to a timeout from a bandwidth-starved operation. The individual TSM servers typically determine individual tape drive performance.

It is not a good idea to configure the maximum number of virtual tape drives onto a single server and then perform I/O through all of them concurrently. Better performance can be achieved by defining a subset of all the virtual tape drives across multiple TSM servers.

Quantum also recommends that backups be staggered so that only a subset of drives is in use at any time. Staggering backups is recommended *IF* you would otherwise exceed the recommended maximum number of concurrently active VTDs. During a backup, the data transfer rate is primarily controlled by the server, because the DXi system does not restrict the ingest data rate. This creates the opportunity for one or more TSM servers to burst data at a higher rate, leaving less bandwidth for the remaining virtual tape drives.

The following table lists the recommended maximum number of concurrently active virtual tape drives for various maximum aggregate bandwidths.

DXi Model	Max VTDs*	Max # of Concurrently Active VTDs	Max Aggregate Bandwidth
DXi470x	64	32	1,650 MB/s (5.9 TB/Hr)
DXi4800			
DXi6700	80	80	972 MB/s (3.5 TB/Hr)
DXi6701 / DXi6702	256	80	1,580 MB/s (5.7 TB/Hr)
DXi6802	256	80	3,299 MB/s (11.9 TB/Hr)
DXi690x 256GB	512	256 **	6,277 MB/s (22.6 TB/Hr)
DXi690x-S 256GB	512	512 **	9,722 MB/s (35.0 TB/Hr)
DXi8500***	160	160	1,777MB/s (6.4 TB/Hr)
DXi8500****	512	160	3,047MB/s (11.0 TB/Hr)
DXi9000			

* Virtual Tape Drives; max # defined in the system

** Quantum suggests that you contact Spectrum Protect / TSM about possible limitations or other considerations related to the maximum number of concurrently active VTDs

*** DXi8500 w/64GB RAM & 2TB drives

**** DXi8500 w/128 GB RAM & 3TB drives

Tape Drive LUN Mapping

Quantum recommends the following:

- Map the device starting with LUN 0 on each port and **DO NOT** skip any LUNs.
- As a best practice, zone the VTL devices and the TSM media servers to prevent other servers from taking control of the VTL resources. Single-initiator style zoning is recommended and is a common best practice in Fibre Channel SANs.
- Additionally, Quantum recommends using the persistent binding feature in the HBA driver to bind the devices to a specific address. This helps keep devices in the same order after a reboot.
- Set the **WWNN = WWPN** for DXi systems. This allows for binding on the HBA to use either WWNN or WWPN.

Quantum DXi-Series VTL devices support “reserve and release” to accommodate sharing drives. This option allows devices to be shared between TSM clients with the Storage Agent installed and TSM servers, thus creating a pool of drives available to each of those systems. Other SAN architectures assign drives to each system and eliminate this shared function.

For both conditions, it is a good practice to keep the TSM server installed on machines separate from other production servers, to eliminate downtime from maintenance activities. Wherever possible, the TSM servers should have fast network connections to their source data.

Tape Cartridge Capacity Considerations

Space on a given tape cartridge cannot be reused until after all backup data on that cartridge has expired. The greater the capacity of a cartridge, the longer it will typically take for all data on that cartridge to expire. Expired data continues to take up space on the virtual tape cartridge, as well as in the DXi, until that cartridge is overwritten, relabeled, or erased.

This means that lower cartridge capacities are more desirable, so that tapes will be returned to the TSM scratch pool for reuse and overwritten sooner. Quantum's general guidance is to specify a smaller virtual tape cartridge capacity, such as 50GB to 100GB, for the reasons mentioned below.

There is virtually no relationship between the configured capacity of a virtual tape cartridge and the tape drive emulation that has been configured for the partition:

- Backup/restore operations will span the number of tapes required, based on the configured capacity.
- Vaulting/duplicating operations performed by the backup application will ignore the virtual capacity when writing to another cartridge, whether virtual or physical. Duplication operations are performed on a set of backup data, not on a per-cartridge basis.
- DXi-Series devices limit the maximum capacity permitted by the tape drive emulation; the minimum is 5GB.

Capacity utilization is tracked in COMPRESSED GB, and data is stored in compressed form. That is, 100GB of data that is 2:1 compressible will be reported as occupying 50GB of virtual tape cartridge space.

TSM Trigger Setting and Cartridge Size

Always re-label cartridges that are returned to the scratch pool. As of version 5.5.3, TSM provides an automatic way to do this. When you re-label cartridges, you are informing the DXi that it can reclaim the space occupied by those cartridges. This will minimize the amount of space consumed by backup data on the DXi. If you do not re-label cartridges, the DXi will continue to protect old expired data until such time as the cartridges are reused. That will artificially increase the amount of space required by the DXi and can result in early or unnecessary capacity upgrade costs.

Setting the TSM Tape Reclamation Trigger is a balancing act. If set too low, more performance load is induced on the DXi, due to increased space reclamation activity. If set too high, an excess of capacity is occupied on the DXi.

Start with the TSM Tape Reclamation Trigger at 50% to balance these factors. In addition, keep in mind that:

- TSM environments with a very low backup load can have the trigger set lower, inducing more load on the DXi to reduce capacity demand.
- TSM environments with a very high backup load might need a higher trigger value, sacrificing capacity to make more "I/O budget" available, to achieve a better backup window.

Unless there are overwhelming reasons not to, use a small virtual cartridge size. 100 GB is recommended.

- *A reason to use a larger cartridge size:* Limitations on the number of cartridges per VTL. (i.e. the end user wants a single VTL to manage all backup data).

- *A reason to use a smaller cartridge:* This minimizes the amount of TTR I/O for files that span multiple cartridges. For example:
 - Assume that a file spans 3 cartridges, occupying part of A, all of B, and part of C. If that file does not expire, but data on A and C expires and triggers TTR, only A and C will be rewritten, and B will be left as-is. This reduces TTR load.
 - If that file had been written to one super-size cartridge, all that data would have been transcribed again and again.

VTL Fibre Channel Performance Tuning

To enhance performance for TSM environments, consider using the tuning parameter indicated by the article mentioned below to eliminate interference from the host system.

According to the Microsoft knowledgebase article [Windows Server 2003 cannot perform backup jobs to tape devices on a storage area network](#), you may encounter the following problem:

“A conflict in Windows Server 2003 causes a Test Unit Ready (TUR) request issue on SCSI-attached, fiber-attached, or iSCSI-attached devices. When this issue occurs, an overflow of TUR requests causes the storage unit not to respond or to respond slowly to SCSI commands. In a fiber or iSCSI SAN environment, any Windows Server 2003-based computer that is zoned to detect the TBU hardware can send TUR requests. If you can see the devices in Device Manager on the server, the devices will send TUR requests even if drivers are not installed for the devices.”

The article referenced above lists the cause of the problem, and a workaround for it.

Considerations for the Handling of Expired Media within TSM

When a tape is expired by TSM, this is not directly communicated to the DXi. The result is that a tape may be displayed as empty or SCRATCH to the user by the TSM console, but the same tape will show on the DXi GUI as containing data. This indicates that the data on the expired tape is still using space on the DXi.

To reclaim this space, Quantum recommends that you label the media when TSM expires the tape cartridge. Any data after the label is truncated on the virtual cartridge (similar to the data after the label being no longer accessible on a physical cartridge.) Space reclamation can then be initiated via the Scheduler, or started manually from the DXi GUI. Additional information about the space reclamation process can be found in the DXi user manual.

Note: In version 5.5.2 and later, TSM can re-label volumes automatically, using the RELABELSCRATCH parameter. (Also available in TSM 5.5.1.0 with a patch - see APAR IC58862).

Example of using RELABELSCRATCH:

```
CLI %> define library <libName> libtype=scsi relabelscratch=yes
```

For more detail, see the [RELABELSCRATCH Information](#) section on p.29 in [Appendix A: Additional TSM Information](#).

Reclamation

TSM Reclamation

Backups and archives in TSM are policy-based. These policies govern:

- Which files on a server will be backed up or archived
- To which media the files will be backed up or archived
- How many versions of various files will be retained
- How long the versions will be retained

Because TSM is policy-based, backup objects, or versions, are expired, rather than the entire backup or individual piece of media. As backup objects on a tape cartridge expire, the cartridge becomes more and more logically empty. Eventually, the amount of free space remaining on the tape drops below a predefined threshold, and space needs to be reclaimed. The TSM reclamation process mounts the tape to be reclaimed, mounts another tape with free space, and copies all the remaining valid data from the source cartridge to the destination cartridge. The source cartridge is now empty and can be reused as a scratch tape.

This TSM reclamation process occurs automatically and is usually scheduled to occur once per day. Reclamation ensures that data is stored efficiently, and that the media is optimized for restoration. If active data pools are used, reclamation must run regularly to remove inactive versions from the active data pool. A version becomes inactive whenever a new version of that object is received from the client.

DXi Reclamation

There is another process, completely internal to the DXi system, known as “reclamation”. This DXi process has nothing to do with “TSM reclamation.” DXi reclamation is part of the internal space management process that occurs on DXi systems. DXi space reclamation only runs when invoked by schedule or initiated manually by the user.

Space reclamation performs two primary functions:

- Reconciliation: The list of expired data is processed to reduce the reference count for the unique data blocks that they referenced.
- Deletion: All unique data blocks that have zero reference counts are deleted from the deduplication pool.

Additional Best Practice Considerations

Several operational considerations are common to both access methods (VTL and NAS). See the [Common Operational Considerations for TSM](#) section on p.23 for more information on Deduplication, Encryption, Compression, Backup Streams and Replication.

Configuring TSM with DXi NAS

A Network Attached Storage (NAS) unit is essentially a self-contained computer connected to an Ethernet network, with the sole purpose of supplying data storage services to other devices on the network. Several DXi models can present themselves as a NAS appliance for backup purposes.

Before you can use a DXi system as a NAS appliance, you must first configure a NAS share on the DXi. The Quantum Network-Attached Storage (NAS) appliance is intended to act as a target for backup applications. This includes Network-Attached Storage or shares. TSM can use a NAS share as a Backup-to-Disk Target.

A DXi system can serve as a NAS backup system where the following protocols are supported:

- **CIFS Protocol** -The CIFS (Common Internet File System) protocol defines a standard for remote file access using many computers at a time. This protocol allows users with different platforms to share files without installing additional software. This protocol is used with Windows networks.
- **NFS Protocol** - The NFS (Network File System) protocol was originally designed by Sun Microsystems and allows all network users to access shared files stored on computers of different types. NFS provides access to shared files through an interface called the Virtual File System (VFS) that runs on top of TCP/IP. Users can manipulate shared files as if they were stored locally on the user's own hard disk. With NFS, computers connected to a network operate as clients while accessing remote files, and as servers when providing remote users access to local shared files. This protocol is used with UNIX/Linux networks.

NAS Device Path Considerations

Network segmentation is the process of splitting a single network into several sub-networks or segments. The advantages of a segmented network are improved performance and security. Performance is improved because there are fewer hosts on the segmented network, which in turn minimizes local traffic. Security is improved because the data traffic is contained on this segment and is not visible to the outside network.

Note: If you are using network segmentation and Microsoft's Automated Deployment Services (ADS), you must use the DXi data segment IP information for ADS management, **NOT** the DXi management segment. ADS uses the Server Message Block (SMB) data protocol to manage the NAS shares on your system, which requires that the management traffic use the data segment.

DXi systems allow you to configure your network for separate segment types. The three primary segments are defined by the type of network traffic that can be used on that segment. The three types of network traffic are:

- **Replication traffic** - This segment is used exclusively for replication data movement.
- **Management traffic** - This segment is used exclusively for remote management (Web page access).
- **Data traffic** - This segment is used exclusively for NAS data movement.

Each network segment has its own network interface (IP address, network mask, and default gateway). In this way, the segment is separated from other network segment traffic.

Note: Regardless of which type of bonding is used on the DXi (Round Robin or LACP), the ports on the Ethernet switch that the DXi are connected to must be in a matching group type as the bond on the DXi.

Installing and configuring the DXi and TSM for NAS operation consists of the following major steps, which are covered immediately below.

- » **Configure the DXi for NAS**
- » **Configure NAS Backups with TSM**

[Configure the DXi for NAS](#)

The DXi system allows you to configure it to present its storage capacity as NAS shares that are compatible with TSM. You can create NAS shares for use with Windows or Linux networks. You can also join the DXi to a Windows domain or workgroup, and manage users.

In the DXi Remote Management Console GUI the **Configuration** page allows you to configure many of the features of the DXi, including storage presentation

Configuring the DXi for NAS lets you choose which network protocol will be used as the transport method for backing up data from client machines to the TSM media server. CIFS (Windows) and NFS (UNIX/Linux) are available on the **NAS > Summary** tab. After NAS Shares have been configured on the DXi, TSM can be configured to use these shares as storage resources.

[Configure NAS Backups with TSM](#)

In this configuration, we are using the NAS share(s) on the DXi as the primary storage pool. With the available capacity of the DXi, and using deduplication, you can create a configuration where data is copied to tape, as a secondary storage pool, only for offsite long-term storage.

In the simplest LAN backup-to-disk configuration, there is a single TSM client, a backup server, and DXi's storage device, all connected via an Ethernet LAN.

Setting up this environment involves the following steps:

1. Mount (or map) the share on the TSM server:
 - a. Create a mount point for the NFS share to be used by TSM, and mount it, using the commands for your particular operating system.
 - b. Map the CIFS share to a drive in Windows Explorer.
2. Using TSM's Integrated Solutions Console, select **Storage Devices**.
3. Create a device class, specifying the following properties:
 - a. Device type: **FILE**
 - b. Name: <name of device class>
 - c. Path: /<mount point>
 - d. Mount limit: <use default value>
Note: Ensure that the mount limit is set to a high enough value to support all concurrent sessions if the default is not sufficient.
 - e. Maximum volume capacity: <desired size of volume you wish to use>
4. Create a storage pool, specifying the following:
 - a. Name: <name of storage pool>
 - b. Type: Sequential access
 - c. Device class: <use device class created in step 3>
 - d. Scratch volume: **1**

Refer to the section on “Sequential Access Storage Pool Volumes” in the IBM Tivoli Storage Manager Administrator’s Guide for more information on creating storage pools of this type.

After you have configured the storage pool:

1. Activate the policy on the storage pool that was just created.
2. Perform a test backup to verify the configuration.

During a backup operation, the NFS or CIFS share on the DXi designated as the storage pool volume receives the backup objects directly from the client/server.

Best Practices Guide with DXi NAS

Number of Shares Considerations

Each DXi system can support multiple NAS shares simultaneously, with a maximum of 128 shares. It is recommended that users create only the required number of shares for each media server. DXi systems can support concurrent NFS and CIFS shares, and can support Fibre Channel VTLs concurrently with those NFS and CIFS shares.

When using NAS shares on DXi systems, Quantum recommends creating at least one share for each TSM server to use. TSM servers should not share NAS shares during normal backup operations. Root access to an NFS share is not allowed, and the access rights will be changed to **nfsnobody** as a security precaution. This does not impact the access to the share from the Tivoli Storage Manager.

Network Share Access Control Considerations

In Windows Active Directory environments, DXi CIFS shares act as targets for TSM. Shares are not intended as primary storage or drag-and-drop storage for other Windows users. Quantum recommends that you create a new account and workgroup, as opposed to joining the domain, to limit access and prevent accidental file deletion by other users. We recommend that you **DO NOT** reconfigure or delete NAS shares while data is being written. There is no mechanism to detect the I/O and provide a warning to the application or administrator.

Network Considerations

Some network considerations include:

- Use a dedicated network for backup data, or use QoS features that guarantee network bandwidth. Another option would be to use virtual networks (VLANs) to segregate backup from production network traffic.
- Configure network interface cards (NICs) in the server and clients, and set routers to full duplex.
- Cabling:
 - Use only CAT5e or CAT6 cables (1Gb/s rated cables).
 - Use only OM3 or OM4 (Aqua) Fibre Optic cables (10Gb/s rated FC).
- If you are using a DNS server, verify that the DNS server configuration settings are correct and don’t conflict, by using **nslookup** on the host name, as well as the IP address.
- It is also a good idea to add the HOSTNAME and IP Address to the host file.
- Use multiple DXi ports when connecting to the network. The more DXi Ports used, the better the performance capability will be across the ports.
- For redundancy, connect at least two DXi ports to an Ethernet switch.

- Leverage the DXi's ability to set up multiple networks. The DXi network configuration allows for integration into nearly any networked environment.
- Set each switch port used by the DXi to auto-negotiate/auto-sensing. The DXi network interface cards are preset to auto/auto and cannot be changed.

TSM NAS Settings and Tuning Considerations

Quantum recommends creating one share for each backup server, when possible. With the DXi NAS attached, there are several advantages to not sharing the NAS B2D shares between backup hosts. The DXi Series can simultaneously support up to 128 shares configured as CIFS and/or NFS shares, allowing concurrent backup jobs to occur to their own shares. This results in faster and more reliable backups.

For TSM, field experience has shown that NFS shares perform better than CIFS shares.

NAS shares created on the DXi Series should be hidden from network browsing by using the **Hide this share from network browsing** option.

When using a DXi as NAS for a backup-to-disk target, consider the following when you create a backup-to-disk folder:

- Set the maximum size for backup-to-disk files to an appropriate size. If you create small but numerous files, performance may be slow, since the computer must still process each file. However, if you create large files, file system limitations can cause memory allocation problems or network issues. These issues can be a problem if you store files across a network.
- Specify fewer backup sets in a backup-to-disk file, to allow TSM to reclaim disk space faster. Fewer backup sets may allow the overwrite protection period to expire sooner.
- Enable the **Allocate the maximum size for backup-to-disk files** option, to reduce disk fragmentation. To ensure that backup data fills each backup-to-disk file to capacity, increase the append period of the media set that you associate with the backup-to-disk files. The backup data's overwrite protection period may also increase, because the overwrite protection period starts at the end of the last append job.
- Backup performance may be affected when you enable the **Allocate the maximum size for backup-to-disk files** option. To find what works best for your environment, enable this option for a job. Then, compare the performance with a job that does not allocate the maximum size for backup-to-disk files. Experiment with the options for buffered reads and buffered writes. Enabling these options may increase backup performance.

Additional Best Practice Considerations

Several operational considerations are common to both access methods (VTL and NAS). See the [Common Operational Considerations for TSM](#) section on p.23 for more information on Deduplication, Encryption, Compression, Backup Streams, and Replication.

Common Operational Considerations for TSM

Data Deduplication Considerations

Deduplication results are negatively impacted by compression, encryption, software deduplication, and multiplexing. These functions all change the data stream in a way that obscures patterns in the data content. They will reduce the performance and deduplication from any downstream appliance, including DXi systems. To obtain effective deduplication rates on the DXi, you should **NOT** encrypt, deduplicate, compress, or multiplex your backup data before sending it to a DXi appliance.

Multiplexing was intended to aggregate multiple streams from slow data sources to meet or exceed the minimum transfer rate required by physical tape drives. This provided more efficient use of a limited number of expensive physical tape drives. Since the virtual tape drives in DXi systems are not susceptible to performance losses from slow data transfer rates, and the number of virtual tape drives can easily be increased in quantity with no cost penalty, there is no reason to use multiplexing with a DXi.

Good Candidates for Data Deduplication

Data deduplication can work well with VMware, large databases (note exceptions below), PowerPoint presentations, Word documents, Excel spreadsheets, Microsoft SQL Server, Oracle (note exceptions below), Exchange databases, and source code repositories.

Not So Good Candidates for Data Deduplication

Data deduplication does not work well with in-line compressed data, SQL with LiteSpeed (in-line compression), Oracle with multi-channel RMAN (in-line multiplex), and compressed video, audio, and JPG images. These data types will not get good FIRST PASS deduplication, but deduplication is not compression. The value of deduplication builds over time with repetition. If there are a large number of files in these formats that don't change, and they get backed up every week as part of a full backup, they will achieve excellent deduplication rates.

For long-term archiving, Quantum recommends that you vault the data to a physical tape device. The DXi NAS shares are the TSM primary storage pool, and physical tape devices act as the secondary storage pool.

Replication Considerations

For first-time replication setups, it is important to manually replicate the name space once the target system is configured and is online. This facilitates the first replication following the first backup to that share/partition.

Keep in mind that:

- The replication is only available to NAS shares with deduplication enabled.
- The DXi supports user-selectable 128-bit or 256-bit AES encryption for replication.
- Data is only encrypted while in transit between the replication source and replication target.
- Data is decrypted upon arrival at the replication target.
- Encryption may affect replication performance. You should disable encryption if your WAN is already secured.

For more information, please refer to *Quantum DXi-Series Best Practices for Data Replication*.

Oversubscription of Space on the DXi

Deduplication will reduce the amount of space used on the physical system. However, Quantum recommends that you monitor for Low Space conditions on the DXi and free up virtual media and disk targets before reaching this threshold. A best practice would be to trigger the Space Reclamation process *before* the DXi reaches approximately 80 percent full.

The **Disk Usage** overview on the **Home** page of the DXi Management GUI displays the following information about disk usage on the system (Note: Values are displayed as an amount, and as a percentage of the total capacity in the system.):

Disk Capacity - The total usable disk capacity of the DXi.

Available Disk Space - The disk space available for data storage (free space).

I/O Write Low Threshold state (Yellow) - Free disk space is equal to or less than 500GB + [10GB * (Total system capacity in TB)]

Stop Write state (Red) - Free disk space is equal to or less than 250GB.

Stop I/O state (Red) - Free disk space is equal to or less than 10GB.

Note: For optimal system performance, Quantum recommends keeping the amount of Available Disk Space (free space) at 20% or more.

Note: When disk capacity is low, target replication to the system is paused. In addition, space reclamation is automatically started to free up disk space.

Space Reclamation

Changes were made to the space reclamation process in DXi 2.2 software for DXi6700 systems. These changes were also made available to DXi4000 and DXi8500 systems, beginning with the DXi 2.2.1 software version.

Space reclamation is the process used to delete blocks that are no longer referenced by metadata, returning capacity to the free space pool (within the blockpool) for future reuse. Starting with DXi 2.2.1 software, there are now two space reclamation modes:

- Legacy Mode
- New Mode

Legacy mode space reclamation is the only method of space reclamation used in pre-DXi 2.2 software. It has four stages:

- Stage 1 – Compaction
- Stage 2 - Calculating Deletion Candidates
- Stage 3 - Deleting Candidates
- Stage 4 – Compaction

New Mode space reclamation was introduced in DXi 2.2 to increase performance and is now the default mode. With New Mode, when space reclamation is manually started using the Remote Management Console **Space Reclamation** page, or is initiated as a scheduled event, only Stage 2 (Calculating Deletion Candidates) and Stage 3 (Deleting Candidates) are run. This improves performance and avoids running Stage 1 and Stage 4 (Compaction), which require a lot of disk I/O.

When disk capacity is low, space reclamation automatically starts, so it can free up disk space on the system. In this case, all four stages of space reclamation are run, which is the same in Legacy mode.

With new Mode space reclamation, free bytes are automatically consumed by the Blockpool. This impacts the byte counts exclusively used by Compaction:

- File system 'disk free disk space' will appear fuller, and will fluctuate less than in pre-DXi 2.2 software versions.
- Unless Compaction is manually run, the "logical free space" in the blockpool will not appear in the file system free disk space pool.
- The logical free space will be reused by the blockpool even if Compaction is not run.

Note: You can still manually execute Compaction using the **syscli --start compaction** command.

Lifecycle of Bytes

The Pre-DXi 2.2 / Legacy Mode Lifecycle of Bytes describes the lifecycle of bytes in pre-DXi 2.2 (or Legacy Mode) space reclamation. It flows like this:

1. Bytes are ingested into the system and are in use (Used Disk Space).
2. When data is deleted, the bytes are still considered in use, but the blobs associated with them are now unreferenced.
3. During space reclamation Stage 3, reference counts are dropped in the blobs (Reclaimable Space). This does not actually free space; it simply makes the blocklets unreferenced.
4. Space reclamation Stage 4 (Compaction) takes all unreferenced blobs and moves them out of cluster bodies and into free space in the file system. Compaction occurs in both Stage 1 and Stage 4 of Legacy Mode space reclamation.

In the New Mode (DXi 2.2), the Lifecycle of Bytes space reclamation bypasses the Compaction stages.

During ingest, the Blockpool takes unreferenced blocklets, reads in the parts of a cluster body that are needed, and then adds new data to it. Reclaimable Space is automatically consumed (reclaimed) by the Blockpool during ingest. The system is not writing data out to Free Space in the file system, which saves on disk I/O.

With New Mode space reclamation, the Reclaimable Space will be reused by the blockpool even if Compaction is not run.

As noted above, Compaction can still be run manually using the **syscli --start compaction** command. Refer to the *DXi-Series Command Line Interface (CLI) Guide* for more details on compaction commands.

Legacy Mode Space Reclamation in DXi 2.2 Software

As previously mentioned, New Mode space reclamation is the default mode in DXi 2.2 software. However, there is a situation in DXi 2.2 when Legacy Mode space reclamation is still run automatically by the system. When free space is low, the system will automatically run Legacy Mode space reclamation to free disk space. In this situation, all four stages of space reclamation are run.

Space Reclamation Status Page

The **Space Reclamation Status** page of the Remote Management Console application can be used to report the status of New Mode space reclamation. After you run space reclamation manually from the **Space Reclamation Status** page (**Utilities > Space Reclamation > General**), the system skips **Stage 1 – Compaction** and starts running **Stage 2 - Calculating Deletion Candidates**.

After Stage 2 completes, Stage 3 – Deleting New Candidates runs.

When Stage 3 completes, Stage 4 – Compaction is skipped and the space reclamation process completes. The **Reclaimable Space** field displays an **Advanced Reporting** link.

To view the **Capacity** report, click the **Advanced Reporting** link to launch DXi Advanced Reporting. Review the **Disk Usage** graph on the **Capacity** report to identify the space reclamation statistics, including **Reclaimable Space**.

The **Disk Usage** graph is new and was added to the **Capacity** report in the most recent version of DXi Advanced Reporting (2.1.2). Use the Disk Usage graph to monitor the amount of disk space that can be freed during space reclamation.

The following statistics display on the graph:

- **Used Disk Space** (light blue) is the amount of disk capacity that is currently being used to store data, including deduplicated data, data that has not been deduplicated, and data allocated to the blockpool that is available to be reused after space reclamation. Used Disk Space increases as data is ingested by the system. Use this information to see the amount of used and free disk capacity on the system.
- **No Deduplication** (dark blue) is data not intended for deduplication and is the amount of disk space used by shares or partitions that are not configured for deduplication.
- **After Reduction** (dark green) is the amount of data that remains after the ingested data has been deduplicated and compressed. This is the amount of data consumed by the blockpool. Use this information to see the amount of deduplicated, compressed data stored on the system.
- **Reclaimable Space** (light green) is the amount of disk space available for new deduplicated data after space reclamation has been run. For non-deduplicated shares and partitions, the DXi may free reclaimable space to create more needed free space when all free space is allocated for Used Disk Space. As previously described, with New Mode space reclamation, reclaimable space is reused by the blockpool even if compaction is not run.

Quantum recommends that you schedule daily space reclamation to manage the available space. Set the scheduled time to start the reclamation process after daily backups are complete. The default schedule is weekly, and the default time for the space reclamation is set to 12:00 AM on Sunday. These parameters are user configurable and should be configured as to not interfere with the normal backup window.

[DXi Multiprotocol Guidance - NFS/VTL Scenario](#)

The NFS **Synchronous** setting requires all data written to be committed to physical storage before the command is acknowledged as complete. This ensures that when a backup completes, all the data resides on disk.

Asynchronous mode allows the system to acknowledge receipt before having the data fully written to disk. This mode allows backups to be completed faster (from the TSM point of view) with the possibility of having an incomplete backup if the system fails (e.g., power is lost) before all the data gets flushed to disk.

The default setting is **Synchronous**. This setting can be altered through the CLI.

Simultaneous inline deduplication of VTL **and** NFS traffic represents the mixing of a heavy, intensive I/O payload with an out-of-order, bursty and response time sensitive protocol.

In a mixed VTL and NFS environment, you should change the configuration for NAS shares from the default mode of **synchronous** to run in **asynchronous** mode. You can do this in the DXi Command Line Interface:

```
syscli --nfscommit async {--share <sharename>} | --all
```

Additional notes:

- All other multi-protocol combinations work well together.
- This recommendation applies to all operating systems and applications.
- Reduced VTL traffic may lessen the frequency of NFS timeouts.

Appendix A: Additional TSM Information

TSM Driver Information

OS	Device Type	Driver	Notes
MS Windows	Quantum Media changer	Native driver or Tivoli driver	Starting from TSM 6.2, the native driver is recommended if available.
	Quantum LTO (SAS Only)	Native driver or Tivoli driver	The native driver is recommended if available.
	OEM Branded IBM LTO (ULTRIUM-TDx)	IBM Driver	IBM Driver must be used.
	OEM Branded HP LTO	Native driver or Tivoli driver	The native driver is recommended if available.
Linux	Quantum Media changer	Linux 'sg' with TSM passthru driver	
	Quantum LTO (SAS Only)	Linux 'sg' with TSM passthru driver	
	OEM Branded IBM LTO (ULTRIUM-TDx)	IBM Driver	IBM Driver must be used.
	OEM Branded HP LTO	Linux 'sg' with TSM passthru driver	
Solaris	Quantum Media changer	Tivoli driver	
	Quantum LTO (SAS Only)	Tivoli driver	
	OEM Branded IBM LTO (ULTRIUM-TDx)	IBM Driver	IBM Driver must be used.
	OEM Branded HP LTO	Tivoli driver	
HP-UX	Quantum Media changer	Tivoli driver	
	Quantum LTO (SAS Only)	Tivoli driver	
	OEM Branded IBM LTO (ULTRIUM-TDx)	IBM Driver	IBM Driver must be used.
	OEM Branded HP LTO	supported with TSM driver	
AIX	Quantum Media changer	Tivoli driver	Tivoli device driver must be used.

	Quantum LTO (SAS Only)	Not supported	
	OEM Branded IBM LTO (ULTRIUM-TDx)	IBM Driver	IBM Driver must be used.
	OEM Branded HP LTO	Supported with TSM driver for non-SAS	

The Transaction Byte Limit Client Option

For information on the TXNBYTELIMIT client option, see the “TXNBYTELIMIT” discussion “ on the following IBM Web page, which is accurate as of [TSM 7.1](#):

https://www.ibm.com/support/knowledgecenter/en/SS36V9_3.2.0/com.ibm.itsm.perf.doc/r_opt_client_txnbytelimit.html

RELABELSCRATCH Information

For information on the RELABELSCRATCH client option, see the “RELABELSCRATCH parameter supported for DEFINE and UPDATE library commands (for Virtual Tape Libraries)” discussion on the following IBM Web page:

https://www.ibm.com/support/knowledgecenter/en/SS36V9_3.2.0/com.ibm.itsm.perf.doc/r_opt_client_txnbytelimit.html

General Virtual Tape Library Considerations for Tivoli Storage Manager Servers

For information on VTL considerations for TSM servers, see the following IBM Web page:

<http://www-01.ibm.com/support/docview.wss?rs=663&uid=swg21425849>

Appendix B: TSM Common Tasks - Command Line

Starting TSM to Get to the Command Prompt (UNIX)

```
>cd /usr/tivoli/TSM/server/bin
>./stopserver (if it doesn't stop shortly, make note of the PID it is trying to stop and type)
>kill -9 (PID)
>./dsmserve (this will start the TSM server in command line mode)
```

Commonly Used Command Line commands

```
help                /*main help menu */
query log           /*show activity log*/
query status       /*show general server information*/
query request      /*see if server needs a response*/
query session      /*show sessions*/
query schedule     /*show schedules*/
query association  /*show nodes associated with schedules*/
query event * * begindate=-7 enddate=today endtime=now exceptiononly=yes
                  /*show missed schedules for last 7 days*/
query node         /*show client accounts*/
query admin        /*show administrative accounts*/
query occupancy   /*show client file space information*/
query db          /*show database statistics*/
query log         /*show recovery log statistics*/
query vol         /*show disk volumes*/
query libvol      /*show volumes in a tape library*/
query dbvol       /*show database volumes*/
query logvol     /*show recovery log volumes*/
disable events nteventlog error /*don't log general errors*/
enable events nteventlog anr2578 /*log missed schedules to nt event log*/
query enabled nteventlog /*showevents that will be logged to nt event log*/
query library     /*show library definitions*/
query drive      /*show drive definitions*/
query devclass   /*show device classes*/
query mount      /*show mounted devices*/
query process    /*show active processes*/
query spacetrigger db /*show settings for extending the database*/
query spacetrigger log /*show settings for extending the log*/
```

Configure the Library with TSM

```
CLI %> define library <libName> libtype=scsi relabelscratch=yes
```

The RELABELSCRATCH parameter is supported in TSM 5.5.1 and higher for the DEFINE and UPDATE library commands (for Virtual Tape Libraries). Use this to blank media that is being recycled. This will free up the space being used by the cartridge.

```
CLI %> define path <serverName> <libName> srctype=server desttype=library
autodetect=yes device=lbx.x.x.x
```

```
CLI %> define drive <libName> <driveName>
```

```
CLI %> define path <serverName> <driveName> srctype=server desttype=drive  
library=<libName> Autodetect=yes device=mtx.x.x.x
```

```
CLI %>define devclass lto devtype=lto library=<libName> mountlimit=x (number  
of drives) estcapacity=400G
```

Add the parameter DRIVEENCRYPTION=OFF for LTO-4 drives. Default is "ALLOW".

```
CLI %>define stgpool tapepool lto pooltype=primary access=readwrite  
maxsize=nolimit maxscratch=50 reusedelay=1
```

```
CLI %>update stgpool backuppool highmig=0 lowmig=0 nextstgpool=tapepool
```

Register Client (Host) to TSM Server

```
CLI %>register node <clientName> <password> (I use the host's root password)  
domain=standard archdelete=yes backdelete=yes maxnummp=x (number of drives)
```

Label Multiple Media via TSM

```
CLI %> label libvolume <libName> search=yes checkin=scratch overwrite=yes  
labelsource=barcode "reply messageID" Check in media via TSM
```

Inventory Library via TSM

```
CLI %> audit library <libName> checklabel=barcode
```

Format/Label Single Media

```
CLI %> label libvolume <libName> <tapeName> overwrite=yes "reply messageID"
```

Format/Label Multiple Media

```
CLI %> label libvolume libName overwrite=yes labelsource=barcode search=yes  
volrange=<tapeName1>,<tapeNameX>
```

Quick Erase Media

```
CLI %> delete volume <tapeName> discarddata=yes (deleted tape is returned to scratch pool)
```

Back Up the Database

```
CLI %>backup db devclass=lto type=full scratch=yes wait=no
```

(This must be done before a file backup can occur.)

Local Backup

```
CLI %> update stgpool diskpool highmig=0 lowmig=0
```

(This will cause all the files to be written to tape as soon as each is written to the diskpool.)

Edit **/usr/tivoli/TSM/client/ba/bin/dsm.opt.smp** – remove "*" in front of "servername" and overwrite words to the right with TSM server name (set in the [Starting TSM to Get to the Command Prompt \(UNIX\)](#) section on p.30.). Save as **dsm.opt**.

Edit **/usr/tivoli/TSM/client/ba/bin/dsm.sys.smp** – change "server_a" to the same name used in the above step. Overwrite "node.domain.company.COM" with the name of your server, (example- svtaix2.adic.com). Save as **dsm.sys**.

```
>/usr/tivoli/TSM/client/ba/bin/dsmj (This starts the client backup gui.)
```

Set Up Scheduled Backup

```
CLI %> define schedule standard <schedule name> type=client action=selective
objects='/<directory path>/'
```

```
options='-subdir=yes' duration=1 durunits=hours period=2 perunits=hours
```

```
CLI %> define association standard <schedule name> <client name> From OS
prompt
```

```
>CD /usr/tivoli/TSM/client/ba/bin <ENTER>
```

```
>./dsmc schedule <ENTER>, <ENTER> on default user ID, type root (TSM Client)
password <ENTER>
```

Media Clone/Media-to-Media Copy

```
CLI %> move data <tapeName>
```

Remove tapes from I/E station.

Extend Database and Recovery Log Files

```
CLI %> define dbvolume db2.dsm formatsize=30 (30 MB)
```

```
CLI %> q db (look for Maximum Extension amount)
```

```
CLI %> extend db <Maximum Extension amount>
```

```
CLI %> define logvolume log2.dsm formatsize=20 (20 MB)
```

```
CLI %> q log (look for Maximum Extension amount)
```

```
CLI %> extend log <Maximum Extension amount>
```

Removing a Library to Replace It with a Different One

```
CLI %> Q volume
```

```
CLI %> Delete volume <volser> discarddata=yes (Do this for each tape listed by q volume.)
```

```
CLI %> Delete volhistory todate=today totime=now type=all
```

```
CLI %> Q libvolume (to get a listing of tapes in the library that TSM knows about)
```

```
CLI %> Checkout libvolume <library name> volrange=<first volser,last volser>
checklabel=no remove=no (there are no spaces between the volsers and the comma)
```

```
CLI %> Delete path <server name> <drive name> srctype=server desttype=drive
library=<library name> (do this for each drive)
```

```
CLI %> Delete drive <library name> <drive name> (do this for each drive)
```

```
CLI %> Delete path <server name> <library name> srctype=server
desttype=library
```


CLI %> Delete library <library name>

Once you have defined your new library and path, give the following command:

CLI %> Update devclass lto library=<library name> (assuming you are using devclass lto)

Helpful Resources

The following is a list of documents, references, and links where you can find additional information regarding specific activities and products.

Retired Best Practice Guides – superseded by this document:

- From Quantum – *Best Practice Guide – Tivoli® Storage Manager with DXi-Series*

Quantum Web Site

<http://www.quantum.com>

StorageCare Guardian Web Site

<http://www.quantum.com/ServiceandSupport/Services/GuardianInformation/index.aspx>

StorageCare Vision Web Site

<http://www.quantum.com/Products/Software/Storagecarevision/Index.aspx>

Quantum Service Web Site

<http://www.quantum.com/ServiceandSupport/Index.aspx>

Worldwide Telephone Support Numbers

For worldwide telephone support numbers, see [this Web page](#).

Call Center Americas

In North America, to contact Quantum's world-class support representatives, please refer to the information below:

- Telephone (toll free): 800-284-5101
- Telephone (toll): 720-249-5700
- Hours of operation (subject to change without notice): 7 days a week, 24 hours a day with valid contract
- 7x24x4 or 7x24x2 coverage available.
- Users with all other contracts can contact Quantum during normal business days from 5 AM to 5 PM US Pacific Time.

View Quantum's Service-Level Objective:

<http://www.quantum.com/ServiceandSupport/ServiceLevelAgreement/Index.aspx>