



SGI® 400 VTL for DMF
Quick Start Guide

007-5713-004

COPYRIGHT

© 2011–2013 Silicon Graphics International Corp. All Rights Reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

LIMITED RIGHTS LEGEND

The software described in this document is "commercial computer software" provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

TRADEMARKS AND ATTRIBUTIONS

COPAN, SGI, SGI OpenVault, SGI InfiniteStorage, the SGI logo, Supportfolio, and XFS are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and other countries.

Celerra is a registered trademark of EMC Corporation. IBM is a registered trademark of International Business Machines Corporation. Linux is a registered trademark of Linus Torvalds in the U.S. and other countries. All other trademarks mentioned herein are the property of their respective owners.

New Features in this Guide

This guide contains clarifications.

Record of Revision

Version	Description
001	January 2011 Original publication
002	Oct 2011 Revised with ISSP 2.5
003	April 2012 Revised with ISSP 2.6
004	April 2013 Revised with ISSP 3.0

Contents

About This Guide	xiii
Related Publications	xiii
Obtaining Publications	xiii
Conventions	xiv
Reader Comments	xiv
1. Introduction	1
DMF in a SGI 400 VTL Environment	1
System Overview	1
DMF Overview	4
SGI 400 VTL as Permanent Storage	4
SGI 400 VTL as Fast-Mount Cache	7
Backups and DMF	10
Naming Convention	10
Configuring OpenVault	10
Summary of the Completed Configuration	10
Preliminary Requirements	11
Set Up the DMF Server and SGI 400 VTL System	12
Ensure that OpenVault Is Not Active	12
Set Up the OpenVault Server and <code>dmf</code> Application	13
Copy the Support Files to the VTL Application Servers	13
Create the Directories that DMF Requires	13
Determine the Backup Requirements For Your Site	15
Site-Specific Factors to Consider for Backups	15
007-5713-004	vii

Default Backup Policy for DMF and SGI 400 VTL	15
Values You Must Determine For Backups	16
Space Required for the Daily Backup	16
Number of Dump Tapes Required	16
Overview of the Steps in this Guide	17
2. Configure DMF	21
Install the DMF License	22
Set Up the Environment for DMF	22
Create the DMF Configuration File	23
Summary of Configuration Parameters	23
DMF Configuration Procedure	24
3. Create and Configure Each Virtual Library	31
Invoke the Console and Attach to the VTL Server	31
Adding the VTL Server at First Use	31
Subsequently Connecting to the VTL Server	34
Determine the Migration Tape Size for the Library	36
Create the Library	40
Configure Virtual Migration Tapes for the Library	50
Create Migration Tapes	50
Relabel the Migration Tapes	54
Free Space in Library SGI-C00 for Dump Tapes	55
Delete Tapes with the 4.5 VTL GUI	58
Delete Tapes with the 7.5 VTL GUI	60
Export Library Access to the DMF Server	63
Configure OpenVault for the Library	74

Add an LCP for the Library	75
Add the Set of DCPs for the Library	80
Import the Migration Tapes for the Library	85
Test that OpenVault Can Mount a Migration Tape from the Library	88
4. Configure Virtual Dump Tapes	89
Create the Dump Tapes	89
Relabel the Dump Tapes	93
Import the Dump Tapes into OpenVault	94
Test that OpenVault Can Mount a Dump Tape	97
5. Test the DMF System	99
Validate the DMF Configuration	99
Enable Audits of the DMF-Managed Filesystems	100
Test the DMF Configuration	104
Appendix A. Using an Extremely Large Number of Tapes	107
Number of Slots	107
Creating Tapes and Relabeling Barcodes in Batches	107
Appendix B. Enabling Compression	109
Appendix C. Reclaiming Space for Dumps	111
Index	115

Figures

Figure 1-1	DMF with SGI 400 VTL as Permanent Storage	2
Figure 1-2	DMF with SGI 400 VTL as a Fast-Mount Cache	3
Figure 1-3	SGI 400 VTL as Permanent Storage: Migrating File Data	5
Figure 1-4	SGI 400 VTL as Permanent Storage: Freeing and Recalling File Data	6
Figure 1-5	SGI 400 VTL as a Fast-Mount Cache: Migrating File Data	8
Figure 1-6	SGI 400 VTL as a Fast-Mount Cache: Freeing and Recalling File Data	9
Figure 1-7	Sample Configuration for SGI 400 VTL as Permanent Storage	11
Figure 1-8	Sample Configuration for SGI 400 VTL as a Fast-Mount Cache	11
Figure 2-1	Replace the Filesystem Name	27
Figure 3-1	Adding the VTL Server	33
Figure 3-2	Connecting to the VTL Server	35
Figure 3-3	Determine the Total Size of the RAID Set	38
Figure 3-4	Create a New Virtual Library	41
Figure 3-5	Select a Shelf	42
Figure 3-6	Specify the Library Vendor and Name	44
Figure 3-7	Enter Virtual Drive Information	45
Figure 3-8	Enter the Barcode Information	47
Figure 3-9	Finish Creating the Library	49
Figure 3-10	Select Physical Devices	51
Figure 3-11	Enter the Number of Tapes to Be Configured	52
Figure 3-12	Verify and Create Tapes	53
Figure 3-13	Delete Selected Migration Tapes to Free Space Needed for Dumps (4.5 GUI)	59
Figure 3-14	Confirm the Deletion of a Tape on SGI-C00	60

Figure 3-15	Selecting Tapes to Delete (7.5 GUI)	61
Figure 3-16	Delete Tapes in the 7.5 VTL GUI	62
Figure 3-17	Select the Initiator WWPN	65
Figure 3-18	Relationship of the "Initiator WWPN" Field to the DMF Server Port	66
Figure 3-19	Verify the DMF Server as the Client	67
Figure 3-20	Select the Library to Assign to the SAN	68
Figure 3-21	Relationship of the "Name" Field to the SGI 400 VTL	69
Figure 3-22	Select a Fibre Channel Target	70
Figure 3-23	Relationship of the Target WWPN Field to the SGI 400 VTL Port	71
Figure 3-24	Verify that the Initiator and Target WWPNs Correspond	72
Figure 3-25	Relationships Among the Library, the DMF Server Initiator WWPN, and the VTL Target WWPN	73
Figure 4-1	Change Tape Properties	90
Figure 4-2	Menu to Change Tape Properties in VTL 7.5 GUI	91
Figure 4-3	Set the Maximum Capacity for Dump Tape Size	92
Figure B-1	Enabling Compression	110

About This Guide

This guide tells you how to configure an SGI[®] 400 virtual tape library (VTL) system for use with the SGI InfiniteStorage[™] Data Migration Facility (DMF) and the OpenVault[™] mounting service using IBM[®] Ultrium-TD4 virtual tapes.

Related Publications

The following publications contain additional information:

- *COPAN 400T/TX System Administration Guide*
- *OpenVault Administrator Guide for SGI InfiniteStorage*
- *DMF 6 Administrator Guide for SGI InfiniteStorage*
- *SGI InfiniteStorage Software Platform* release note (the `README.txt` file, see below to obtain)

Obtaining Publications

You can obtain SGI documentation as follows:

- See the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, man pages, and other information.
- You can view man pages by typing `man title` at a command line.
- The `/docs` directory on the ISSP DVD or in the Supportfolio[™] download directory contains the following:
 - The ISSP release note: `/docs/README.txt`
 - DMF release notes: `/docs/README_DMF.txt`
 - A complete list of the packages and their location on the media:
`/docs/RPMS.txt`
 - The packages and their respective licenses: `/docs/PACKAGE_LICENSES.txt`

- The release notes and manuals are provided in the `noarch/sgi-isspdocs` RPM and will be installed on the system into the following location:

`/usr/share/doc/packages/sgi-issp-ISSPVERSION/TITLE`

Conventions

The following conventions are used throughout this document:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)
[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.
<code>manpage(x)</code>	Man page section identifiers appear in parentheses after man page names.
GUI	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this publication, contact SGI. Be sure to include the title and document number of the publication with your comments. (Online, the document number is located in the front matter of the publication. In printed publications, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:

techpubs@sgi.com

- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:

SGI
Technical Publications
46600 Landing Parkway
Fremont, CA 94538

SGI values your comments and will respond to them promptly.

Introduction

This guide tells you how to configure an SGI[®] 400 virtual tape library (VTL) system for use with the SGI InfiniteStorage Data Migration Facility (DMF) and the OpenVault[™] mounting service using IBM[®] Ultrium-TD4 virtual tapes. It supports the both 4.5 and 7.5 versions of the VTL software.

This chapter discusses the following:

- "DMF in a SGI 400 VTL Environment" on page 1
- "Preliminary Requirements" on page 11
- "Overview of the Steps in this Guide" on page 17

Note: This guide documents the supported procedure for configuring DMF with SGI 400 VTL. If you deviate from this procedure, DMF may not function.

DMF in a SGI 400 VTL Environment

This section discusses the following:

- "System Overview" on page 1
- "DMF Overview" on page 4
- "Backups and DMF" on page 10
- "Naming Convention" on page 10
- "Configuring OpenVault" on page 10
- "Summary of the Completed Configuration" on page 10

System Overview

An SGI 400 VTL cabinet has up to eight shelves, shelf 0 (bottom) through shelf 7 (top). A site can have multiple cabinets with differing numbers of shelves. Each cabinet has up to two VTL application servers (referred to in this guide as VTL0 and VTL1) that control the shelves. VTL0 controls shelves 0–3, VTL1 controls shelves 4–7.

Figure 1-1 shows a conceptual drawing of DMF using SGI 400 VTL with eight shelves as permanent storage.

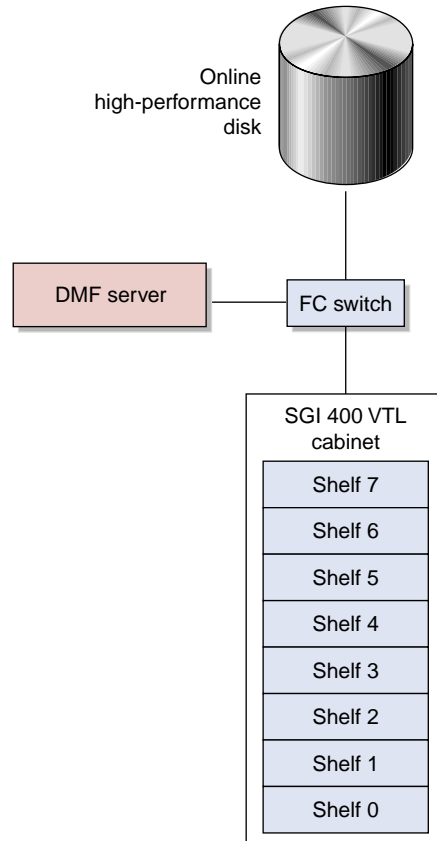


Figure 1-1 DMF with SGI 400 VTL as Permanent Storage

If you use SGI 400 VTL in a fast-mount cache configuration, you would also need a permanent migration target, such as physical tape. Figure 1-2 shows a conceptual drawing of DMF using SGI 400 VTL with four shelves as fast-mount cache.

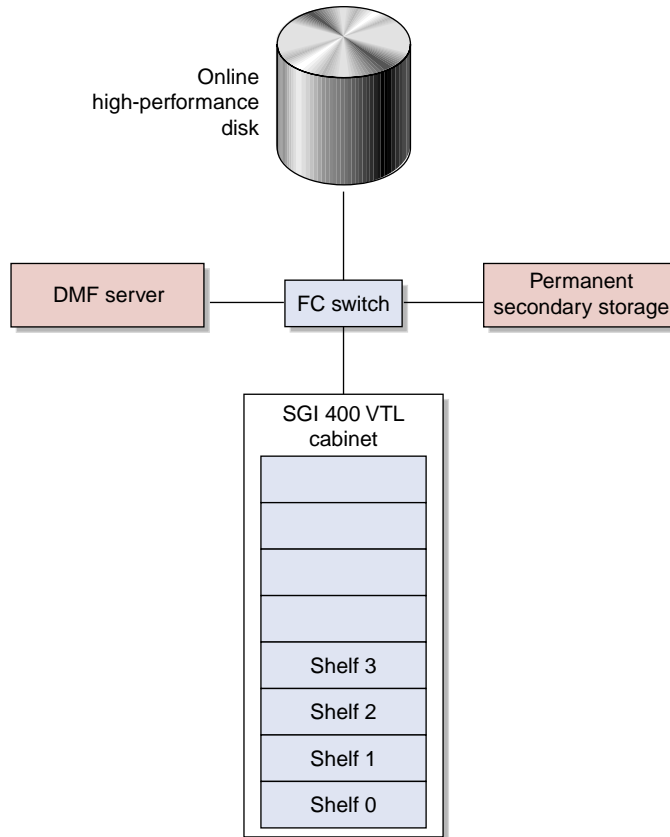


Figure 1-2 DMF with SGI 400 VTL as a Fast-Mount Cache

For users with higher throughput requirements, the *Parallel Data Mover Option* allows additional data movers on dedicated nodes to operate in parallel with the integrated data mover functionality on the DMF server, increasing data throughput and enhancing resiliency. The dedicated function of a *parallel data mover node* is to move data from the filesystem to secondary storage or from secondary storage back into the primary filesystem. Offloading the majority of I/O from the DMF server improves I/O throughput performance.

DMF Overview

Note: For complete details about DMF, see the DMF administrator guide.

When DMF migrates or recalls files, it issues a mount request to the OpenVault mounting service. OpenVault mounts the correct SGI 400 VTL archive filesystem and DMF performs I/O to the appropriate volume on that filesystem.

DMF continuously monitors the managed user filesystem according to the policies established in the DMF configuration file. Only the most timely data resides on the higher performance primary filesystem; less timely data is automatically migrated to the archive tiers. However, data always appears to be online to end users and applications, regardless of its actual location.

DMF moves file data to secondary storage on the SGI 400 VTL disk, but it leaves critical metadata (such as index nodes, or *inodes*, and directories) in the primary filesystem. A user retrieves a file simply by accessing it normally through NFS; DMF automatically recalls the file's data from the secondary storage, caching it on the primary filesystem. From a user's perspective, all content is visible all of the time.

You can use SGI 400 VTL as either of the following:

- "SGI 400 VTL as Permanent Storage" on page 4
- "SGI 400 VTL as Fast-Mount Cache" on page 7

SGI 400 VTL as Permanent Storage

Figure 1-3 and Figure 1-4 show the concepts of migrating and recalling file data when using SGI 400 VTL as permanent storage.

Note: For simplicity, these diagrams do not address a second copy. Data will be recalled from a second copy only if necessary.

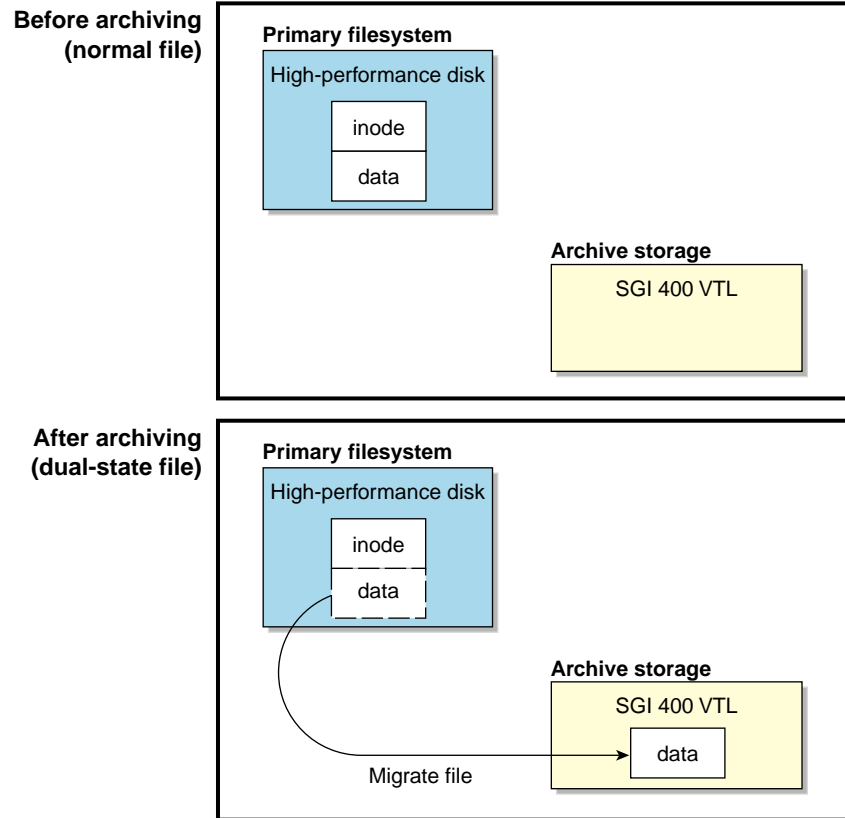


Figure 1-3 SGI 400 VTL as Permanent Storage: Migrating File Data

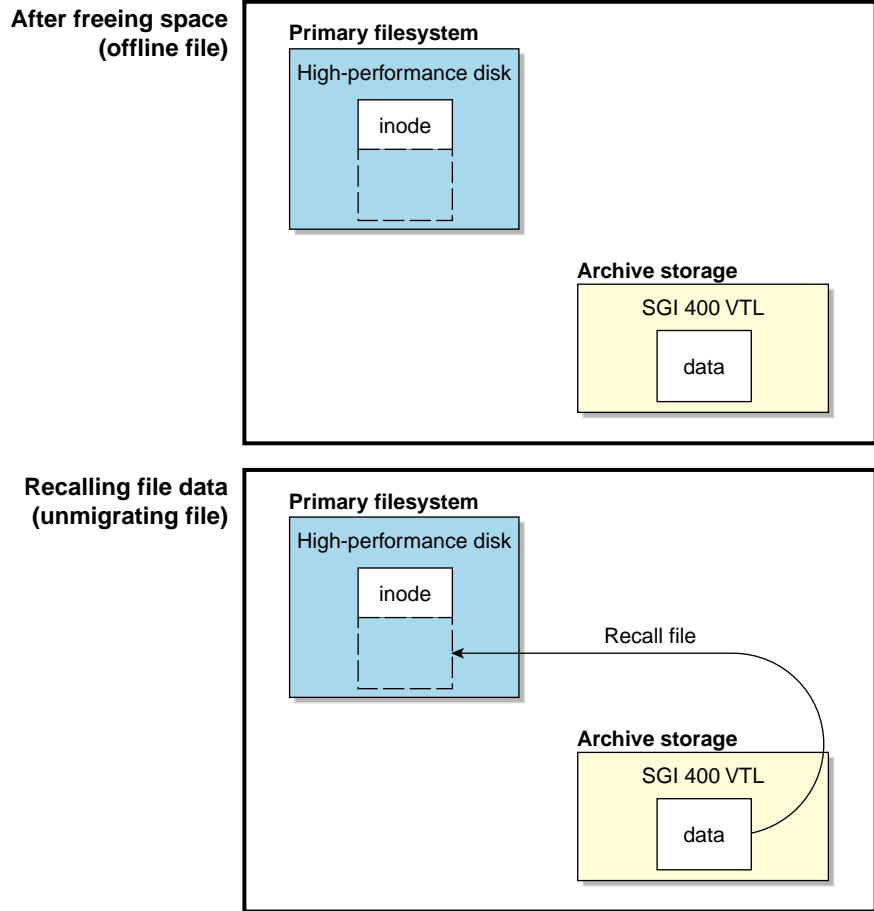


Figure 1-4 SGI 400 VTL as Permanent Storage: Freeing and Recalling File Data

The figures describe the archiving concept, showing that data is removed from the primary filesystem after migrating (represented by the dashed lines) and moved to the archive on VTL, but the inode remains in place in the primary filesystem. During the period when the file data has been copied to the secondary storage on VTL but has not yet been deleted from the primary filesystem, the file is considered to be *dual-state*. After the file data has been deleted from the primary filesystem, the file is considered to be *offline*. SGI recommends that you migrate two copies of a file to prevent file data loss in the event that a migrated copy is lost.

Because the inodes and directories remain online, users and applications never need to know where the file data actually resides; archived files remain cataloged in their original directories and are accessed as if they were still online. In fact, when drilling into directories or listing their contents, a user cannot determine whether a file is online or offline; determining the data's actual residence requires special commands or command options. The only difference users might notice is a slight delay in access time. Therefore, DMF allows you to oversubscribe your online disk in a manner that is transparent to users.

SGI 400 VTL as Fast-Mount Cache

The *fast-mount cache* configuration migrates data simultaneously to SGI 400 VTL as a temporary cache and to permanent storage on another migration target (such as physical tape). Volumes on the cache can be freed immediately when the fullness threshold is reached. SGI always recommends that you migrate at least two copies to permanent storage targets in order to prevent file data loss in the event that a migrated copy is lost. When using a fast-mount cache, SGI therefore recommends that you migrate at least three copies (one to the temporary cache and two to permanent storage targets).

Figure 1-5 and Figure 1-6 show the concepts of migrating and recalling file data when using SGI 400 VTL as a fast-mount cache in conjunction with permanent storage , in this case on physical tape.

Note: For simplicity, these diagrams do not address a second permanent storage copy. Data will be recalled from a second permanent storage copy only if necessary.

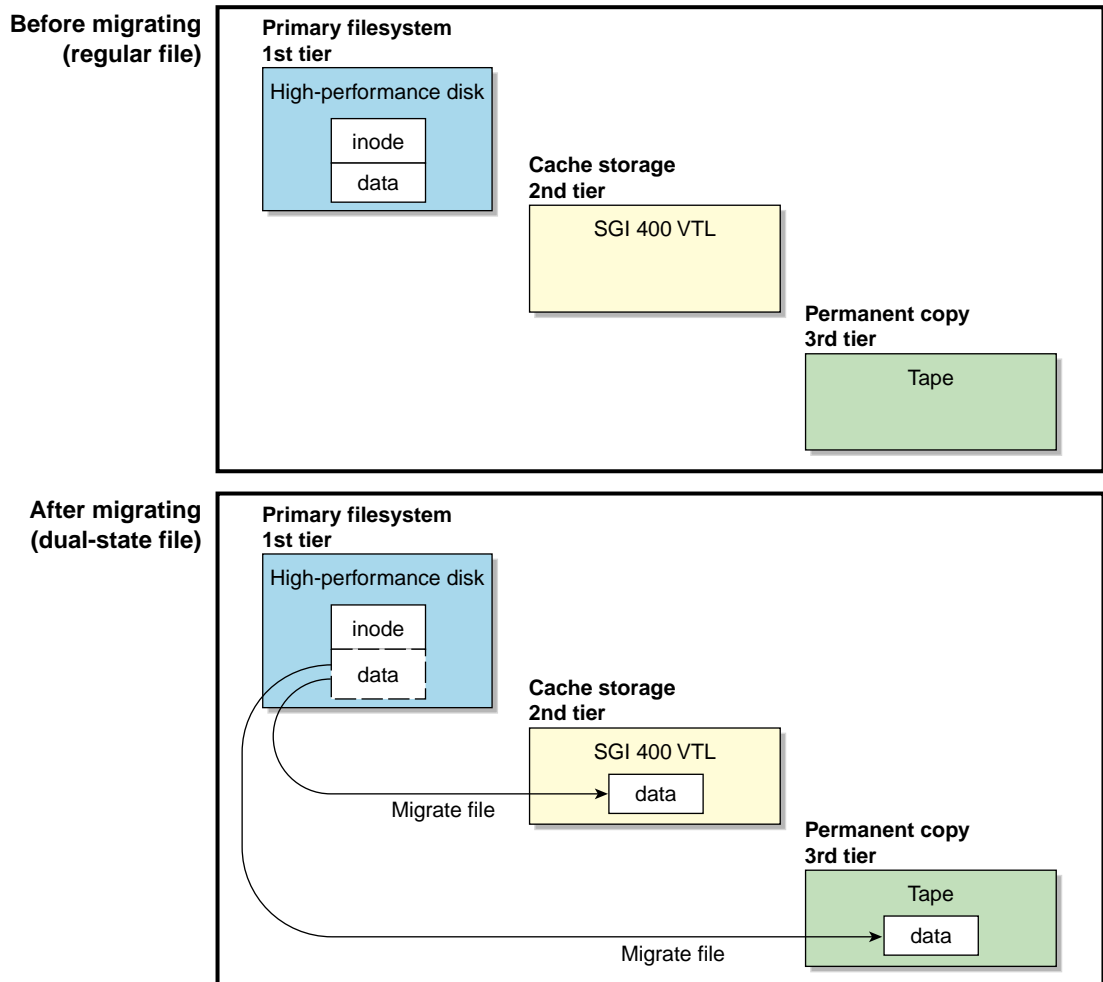


Figure 1-5 SGI 400 VTL as a Fast-Mount Cache: Migrating File Data

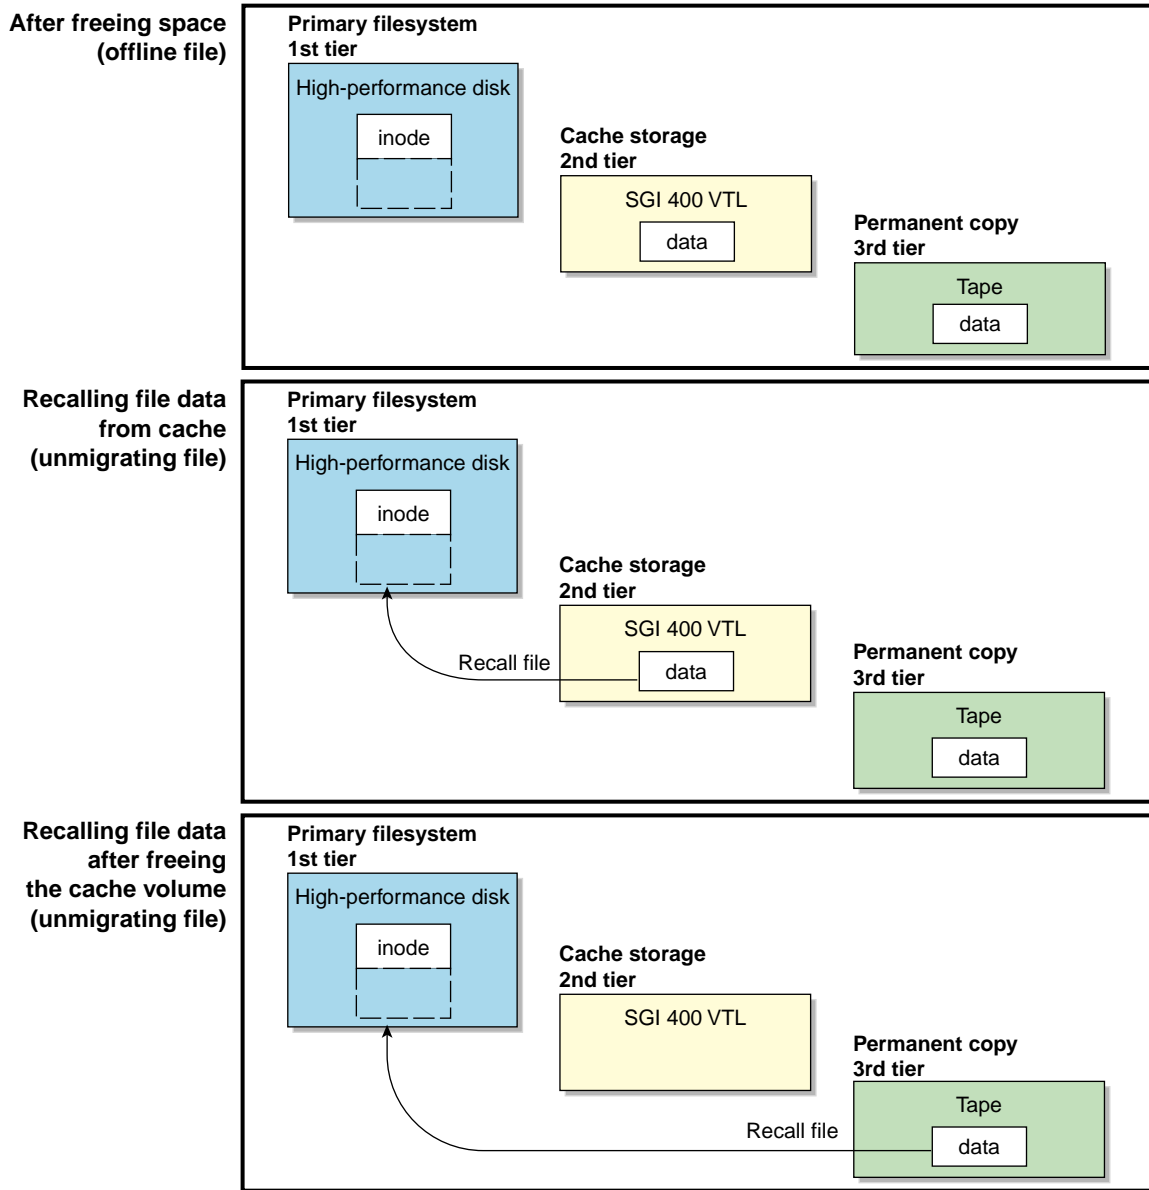


Figure 1-6 SGI 400 VTL as a Fast-Mount Cache: Freeing and Recalling File Data

Backups and DMF

DMF moves only the data associated with files, not the file inodes or directories. Therefore, you must still perform regular backups to protect files that have not been migrated, as well as inodes and directory structures. You must store the backups on physical tape or in a filesystem that is not managed by DMF.

Naming Convention

DMF uses a *volume serial number (VSN)* to uniquely identify a specific volume at your site. The volumes may be distributed throughout the SGI 400 VTL cabinet. All VSNs that are managed by a single OpenVault instance must be unique.

SGI recommends that you use an SGI 400 VTL naming convention where SGI-C00 represents C for *Cabinet*, 0 (zero) for cabinet 0, and 0 for shelf 0. SGI also recommends that you use a similar name but varying case to maintain a correspondence between the SGI 400 VTL view of the library (such as SGI-C00) and OpenVault's view of the library (such as c00). The names used in this guide are recommended and are also used in the sample DMF configuration file for SGI 400 VTL.

Configuring OpenVault

You will use the `ov_admin(8)` command to add the DMF application to OpenVault. See "Configure OpenVault for the Library" on page 74.

Summary of the Completed Configuration

As an example, the procedures in this guide will lead you through the process of configuring a single cabinet with multiple shelves. You must adjust the procedures and the DMF configuration file to fit your site-specific situation.

After you complete the procedures in this guide, the cabinet will have up to eight independent OpenVault libraries, c00 through c07, and a set of up to sixteen drives per shelf, such as c00d00 through c00d15 for shelf 0. Each library will be used with a corresponding volume group, such as vg_c00 through vg_c07. The volume groups will be managed as two DMF migrate groups (mg0 and mg1 for SGI 400 VTL as permanent storage or fmc1 and fmc2 for fast-mount cache).

Figure 1-7 and Figure 1-8 describe the parts of the configuration that relate to SGI 400 VTL.

SGI 400 VTL		OpenVault		DMF		
Cabinet	Virtual library	Library name	Drive names	Drive group	Volume group	Migrate group
shelf 7	SGI-C07	C07	C07d1-C07d7	dg_c07	vg7_c07	} mg1
shelf 6	SGI-C06	C06	C06d1-C06d7	dg_c06	vg6_c06	
shelf 5	SGI-C05	C05	C05d1-C05d7	dg_c05	vg5_c05	
shelf 4	SGI-C04	C04	C04d1-C04d7	dg_c04	vg4_c04	
server VTL1						
server VTL0						
shelf 3	SGI-C03	C03	C03d1-C03d7	dg_c03	vg3_c03	} mg0
shelf 2	SGI-C02	C02	C02d1-C02d7	dg_c02	vg2_c02	
shelf 1	SGI-C01	C01	C01d1-C01d7	dg_c01	vg1_c01	
shelf 0	SGI-C00	C00	C00d1-C00d7	dg_c00	vg0_c00	

Figure 1-7 Sample Configuration for SGI 400 VTL as Permanent Storage

SGI 400 VTL		OpenVault		DMF			
Cabinet	Shelf name	Library name	Drive names	Drive group	Volume group	Migrate group	Fast-mount cache
shelf 3	C03	C03	C03d00-C03d15	dg_c03	vg_c03	}	} copan_fmc2
shelf 2	C02	C02	C02d00-C02d15	dg_c02	vg_c02		
shelf 1	C01	C01	C01d00-C01d15	dg_c01	vg_c01	} mg_fmc	} copan_fmc1
shelf 0	C00	C00	C00d00-C00d15	dg_c00	vg_c00		

Figure 1-8 Sample Configuration for SGI 400 VTL as a Fast-Mount Cache

Preliminary Requirements

Before you configure the virtual library and DMF, do the following:

- "Set Up the DMF Server and SGI 400 VTL System" on page 12

- "Ensure that OpenVault Is Not Active" on page 12
- "Set Up the OpenVault Server and `dmf` Application" on page 13
- "Copy the Support Files to the VTL Application Servers" on page 13
- "Create the Directories that DMF Requires" on page 13
- "Determine the Backup Requirements For Your Site" on page 15

Set Up the DMF Server and SGI 400 VTL System

You must ensure that the DMF server and the SGI 400 VTL system are on the network and zoned appropriately. SGI 400 VTL disk drives must be visible only from the active DMF server, the passive DMF server (if applicable), and the parallel data mover nodes; the drives must not be visible to any other nodes. You must also install the required software for SGI 400 VTL and DMF.

For more information, see: the hardware documentation for your systems, your SGI 400 VTL system administration documentation, the *SGI InfiniteStorage Software Platform* release note (see the Preface for the location), and *DMF 6 Administrator Guide for SGI InfiniteStorage*.

Ensure that OpenVault Is Not Active

You must ensure that OpenVault is not actively being used by other applications. For example, verify that no processes are displayed by the `ov_procs(8)` command, as in the following example:

```
dmfserver# ov_procs
PID      PPID     COMMAND
dmfserver#
```

Note: For simplicity and efficiency, the procedures in this guide stop OpenVault at various points. If you have other applications that are currently running OpenVault, those applications will be affected.

Set Up the OpenVault Server and dmF Application

If this is a new DMF installation, you must use the `ov_admin(8)` command to do the following:

1. Initially configure the OpenVault server
2. Add a privileged and an unprivileged instance of the dmF application

For more information, see the chapter about mounting service configuration tasks in *DMF 6 Administrator Guide for SGI InfiniteStorage*.

Copy the Support Files to the VTL Application Servers

You must copy all of the files in the `/usr/lib/dmf/support/copan` directory on the DMF server to a new `/usr/local/dmf` directory on each of the VTL application servers (VTL0 and VTL1).¹

For example, for application server VTL0 and DMF server mynode:

```
[root@VTL0 vtl]# pwd
/usr/local/dmf
[root@VTL0 vtl]# mkdir dmF
[root@VTL0 vtl]# cd dmF
[root@VTL0 dmF]# scp mynode.mycompany.com:/usr/lib/dmf/support/copan/* .
Password:
dmreadycopantapes          100% 5251      1.3KB/s   00:00
dmreadycopantapesV1       100% 1900      1.3KB/s   00:00
iscli_copan_cmd           100% 3367      2.0KB/s   00:00
xslt3_copan                100% 1992      2.0KB/s   00:00
[root@VTL0 dmF]# ls -l
-rwxr-xr-x 1 root sys 5251 2012-02-24 13:41 dmreadycopantapes
-rwxr-xr-x 1 root sys 1900 2012-02-24 13:41 dmreadycopantapesV1
-rwxr-xr-x 1 root sys 3367 2012-02-24 13:41 iscli_copan_cmd
-rwxr-xr-x 1 root sys 1992 2012-02-24 13:41 xslt3_copan
```

Create the Directories that DMF Requires

Use the information about configuring DMF administrative directories appropriately in the “Best Practices” chapter of the DMF administrator’s guide to create the

¹ For historic reasons, the SGI 400 VTL is sometimes referred to as *COPAN* in literals and the graphical user interfaces.

required filesystems and directories of the appropriate size on a general-purpose RAID storage system.

In a production system, SGI in most cases recommends that you restrict these directories to DMF use and make them the mountpoint of a filesystem, in order to limit the loss of data in the case of a filesystem failure.

You specify the location of these directories by using parameters in the DMF configuration file. The following lists show the directory names used by the sample configuration file (using these names will minimize the number of changes you must make):

- **Required** to be dedicated to DMF use and to be a filesystem mountpoint:
 - *(If used)* `/dmf/move` for the scratch filesystem that is used to move files between volume groups

Note: You must specify a scratch filesystem if you intend to use the `dmmove(8)` command. The scratch filesystem should be dedicated to the `dmmove` function. You must use the `dmi,mtpt=/dmf/move` option.

- **Recommended** to be dedicated to DMF use and to be a filesystem mountpoint:
 - `/dmf/cache` for storing chunks of data while merging them from sparse volumes
 - `/dmf/home` for directories in which the DMF daemon database, library server (LS) database, and related files reside
 - `/dmf/journals` for directories in which the journal files for the daemon database and LS database will be written
 - `/dmf/spool` for directories in which DMF log files are kept
 - `/dmf/tmp` for directories in which DMF puts temporary files for its own internal use
- **Additional directories:**
 - `/dmf/journals/database_copies` and `/var/lib/dmf/database_copies` for DMF database copies
 - `/dmf/dumps` for DMF dump files created by DMF backup tasks

You can choose to use different locations by modifying the `DATABASE_COPIES` configuration parameter in step 6c of Chapter 2, "Configure DMF" on page 21.

You can use the `df(1)` command to verify that the filesystems required for the DMF administrative directories and the DMF-managed filesystems are mounted.

Determine the Backup Requirements For Your Site

This section discusses the following:

- "Site-Specific Factors to Consider for Backups" on page 15
- "Default Backup Policy for DMF and SGI 400 VTL" on page 15
- "Values You Must Determine For Backups" on page 16

Site-Specific Factors to Consider for Backups

The backup requirements depend upon a number of very site-specific factors, including the following:

- The amount of data that is not migrated at the time a dump takes place
- The size of the DMF databases, which is a function of the number of migrated files and the number of DMF copies of each file
- The number of full and/or partial backups to retain, as determined by the frequency of backups and the dump retention period

For more details, see the information about configuring DMF administrative filesystems and directories appropriately in the "Best Practices" chapter of the *DMF 6 Administrator Guide for SGI InfiniteStorage*, particularly the information about `HOME_DIR` size.

Default Backup Policy for DMF and SGI 400 VTL

The sample DMF configuration file for SGI 400 VTL does the following:

- Performs a full backup once a week (Sundays at 00:00) and a partial backup the remaining days (Monday—Saturday at 00:00), creating dumps of the following filesystems:
 - `/dmf/home`

- /dmf/journals
- All DMF-managed filesystems
- Causes all of the file data in the DMF-managed filesystems to be migrated before the dumps take place (except for files that do not meet the site migration policy).
- Removes the bit-file identifiers from the DMF databases for permanently deleted files.
- Retains the dumps for four weeks (28 days). The backup each day requires at least one dump tape. The dumps for each day's backup will be put on a single tape if possible. Tapes are recycled after the retention period is completed.

If you want to change these backup policies, you must modify the DMF configuration file and the procedures in this guide accordingly. For more information, see the *DMF 6 Administrator Guide for SGI InfiniteStorage*.

Values You Must Determine For Backups

You must determine the following:

- "Space Required for the Daily Backup"
- "Number of Dump Tapes Required"

Space Required for the Daily Backup

The amount of space that will be required by your site for the dumps created by each day's backup depends upon the information discussed in "Site-Specific Factors to Consider for Backups" on page 15. This amount is the *Dump_Space_Needed_Per_Day* value, which you will use later in step 1a and step 2 of "Free Space in Library SGI-C00 for Dump Tapes" on page 55.

In general, a smaller size is better for dump tapes in order to use space efficiently, so long as all of the dumps for each day's backup can fit onto one tape. (However, if the dumps for day's backup require more than one tape, it will increase the complexity and time requirements of restores.)

Number of Dump Tapes Required

The number of dump tapes that will be used depends upon the retention period and the information in "Site-Specific Factors to Consider for Backups" on page 15.

Tapes are recycled only after the retention period is completed, therefore you must have more dump tapes than are required to fulfill the retention period (at least one extra tape). Assuming that dumps are done daily, the minimum number of tapes required is:

$$\text{Retention_Period_In_Days} + 1 = \#_Dump_Tapes$$

For example, using the retention period of 4 weeks (28 days) in the DMF sample configuration file for SGI 400 VTL:

$$28 + 1 = 29 \text{ tapes}$$

So long as the dumps for each day's backup can fit onto one dump tape, this means that at a minimum 29 dump tapes are required for dumps, assuming that backups are performed each day. You will use the `#_Dump_Tapes` value in step 1b of "Free Space in Library SGI-C00 for Dump Tapes" on page 55.

Note: You should monitor the dump report daily to verify that there are sufficient tapes available for future dumps.

Overview of the Steps in this Guide

This guide leads you through the following steps, which you will perform as the `root` user:



Caution: You must successfully complete each step before moving on to the next step. If you encounter an error, fix the problem before continuing.

1. Configure DMF:
 - a. Install the DMF license.
 - b. Set up the environment for DMF.
 - c. Create the DMF configuration file by using DMF Manager.
2. Create and configure each virtual library in turn, beginning with shelf 0, which will be configured as library `SGI-C00`:
 - a. Invoke the VTL Console graphical user interface and attach to the appropriate VTL server.

- b. Determine the size of migration tapes for the library.
- c. Create the library.

Note: If your site's data would benefit from software compression, you can enable SGI 400 VTL compression mode (and therefore DMF compression) for each application server (VTL0 and VTL1). See the optional Appendix B, "Enabling Compression " on page 109. However, compression carries an overhead CPU penalty and is not appropriate in all cases.

- d. Configure virtual migration tapes for the library:
 - I. Create migration tapes.
 - II. Relabel the barcode (known as the VSN in DMF)for the migration tapes.

Note: The DMF library server uses the barcode when selecting tapes. A script relabels the barcode so that it indicates the specific RAID set in which the tape is stored, allowing DMF to optimize I/O to and from the shelf.

- e. Free space in the library on shelf 0 (SGI-C00) for dump tapes. You will do this only for shelf 0.

Note: All of the dump tapes must be on a single shelf. The sample DMF configuration file defines this to be shelf 0. If you want to use a different shelf, you must modify the DMF configuration file drive group and the procedures in this guide accordingly.

- f. Export library access to the DMF server.
- g. Configure OpenVault for the library:
 - I. Add a library control program (LCP) for the library.
 - II. Add the set of drive control programs (DCPs) per library (one for each drive).
 - III. Import the migration tapes for the library.
 - IV. Test that OpenVault can mount a migration tape from the library.

After completing these steps for library SGI-C00, repeat them for shelf 1 and library SGI-C01, and so on until you have completed the last library (such as SGI-C07).

3. Reuse the space freed from migration tapes in step 2e above to create virtual dump tapes:
 - a. Create the dump tapes.
 - b. Relabel the barcode for the dump tapes.
 - c. Import the dump tapes into OpenVault.
 - d. Test that OpenVault can mount a dump tape.
4. Test DMF:
 - a. Validate the DMF configuration by using DMF Manager or the `dmcheck(8)` command.
 - b. Verify the DMF-managed filesystems by using the `dmaudit(8)` command.
 - c. Test the DMF configuration.



Caution: After your system is configured and running, you should not stop the disks in a SGI 400 VTL shelf if those disks are currently being used by an OpenVault LCP. First view the output from `dmstat(8)` or `ov_stat(8)` to ensure that DMF is not using the disks in that shelf, then stop the OpenVault LCP associated with that shelf, and finally stop the disks. For more information, see the “Best Practices” chapter of the *DMF 6 Administrator Guide for SGI InfiniteStorage*.

Configure DMF

This chapter describes the minimum steps required to get DMF up and running with a SGI 400 VTL system:

- "Install the DMF License" on page 22
- "Set Up the Environment for DMF" on page 22
- "Create the DMF Configuration File" on page 23

For more information, see the following:

- *DMF 6 Administrator Guide for SGI InfiniteStorage*
- *DMF 5 Filesystem Audit Guide for SGI InfiniteStorage*
- DMF man pages:

- `dmput(1)`
- `dm1s(1)`
- `dmf.conf(5)`
- `dmdidle(8)`

Install the DMF License

To install the DMF license, copy the DMF license key into the `/etc/lk/keys.dat` file on the DMF server. For more information, see the chapter about licensing in the DMF administrator's guide.

Set Up the Environment for DMF

To set up the environment for DMF, do the following:

1. Verify that the required DMF administrative filesystems and DMF-managed filesystems are mounted. For example:

```
dmfserver# df -h
...
/dev/lxvm/lun5s3      47G   65M   47G   1% /dmf/move
/dev/lxvm/lun5s5      47G   66M   47G   1% /dmi_fs
/dev/lxvm/lun5s0      47G   70M   47G   1% /dmf/tmp
/dev/lxvm/lun5s2      47G   79M   47G   1% /dmf/home
/dev/lxvm/lun5s1      47G  107M   47G   1% /dmf/spool
/dev/lxvm/lun5s9      47G   69M   47G   1% /dmf/journals
/dev/lxvm/lun5s7      47G   65M   47G   1% /dmf/cache
...
```

For more information, see "Create the Directories that DMF Requires" on page 13.

2. Create an empty `/dmf/home/dump_tasks/tapes` file. For example:

```
dmfserver# touch /dmf/home/dump_tasks/tapes
```

Create the DMF Configuration File

This section discusses the following:

- "Summary of Configuration Parameters" on page 23
- "DMF Configuration Procedure" on page 24

Summary of Configuration Parameters

Table 2-1 and Table 2-2 highlight the configuration file parameters that are of particular interest when using SGI 400 VTL shelves, grouped by the object in the sample file that contains the parameter. For details, see *DMF 6 Administrator Guide for SGI InfiniteStorage*.

Table 2-1 Parameters of Particular Interest for SGI 400 VTL as Permanent Storage

Sample Object Name	Object Type	Parameter
DMF (shield icon)	base	ADMIN_EMAIL SERVER_NAME
vtl_ls	libraryserver	COPAN_VSNS DRIVE_GROUPS
dg_c00-dg_c07	drivegroup	MOUNT_SERVICE_GROUP
vg_c00-vg_c07	volume group	MAX_PUT_CHILDREN RESERVED_VOLUMES ZONE_SIZE
mg0,mg1	migrategroup	GROUP_MEMBERS MULTIPLIER
dump_tasks	taskgroup	DUMP_FILE_SYSTEMS

Table 2-2 Parameters of Particular Interest for SGI 400 VTL as Fast-Mount Cache

Sample Object Name	Object Type	Parameter
DMF (shield icon)	base	ADMIN_EMAIL SERVER_NAME
<i>nodename</i>	node	INTERFACE (<i>Parallel Data Mover Option only</i>)
copan_ls	libraryserver	COPAN_VSNS DRIVE_GROUPS
vg_c00 through vg_c03	volumegroup	MAX_PUT_CHILDREN MERGE_THRESHOLD RESERVED_VOLUMES ZONE_SIZE
mg_fmc	migrategroup	GROUP_MEMBERS
fmc_task1 and fmc_task2	taskgroup	FILE_RETENTION_DAYS FMC_MOVEFS FMC_NAME FREE_VOLUME_MINIMUM FREE_VOLUME_TARGET
copan_fmc1 and copan_fmc2	fastmountcache	CACHE_MEMBERS

DMF Configuration Procedure

To create the DMF configuration using the DMF Manager graphical user interface, do the following:

1. Start the DMF Manager service:

```
dmfserver# service dmfman start
```

2. Invoke DMF Manager by pointing your browser to the following location:

```
https://your_hostname:1179
```

3. Accept the security certificate.

For more information, see the chapter about using DMF Manager in the DMF administrator's guide. You can use the **Help** button in DMF Manager to access the DMF administrator's guide.

4. Log in as the `admin` user. Click the **Log In** button in the upper-right corner of the window.

Note: The first time you configure DMF, you are automatically given `admin` access. After you save the configuration, you must thereafter log in as the `admin` user to make changes. For more information, see the chapter about DMF Manager in the DMF administrator's guide.

5. Populate DMF Manager with the sample DMF configuration file for SGI 400 VTL use. To do this, right-click in the **Overview** panel and select one of the following:
1

Pre-Configured

> **COPAN VTL LS Sample**

Pre-Configured

> **Fast-Mount Cache Sample**

COPAN VTL LS Sample creates the following:

- One library server: `vtl_ls`
- Eight drive groups: `dg_c00` through `dg_c07`
- Eight volume groups: `vg_c00` through `vg_c07`
- Two migrate groups: `mg0` and `mg1`
- The DMF administrative filesystems, as described in "Create the Directories that DMF Requires" on page 13

Fast-Mount Cache Sample creates the following:

- One library server: `copan_ls`
- Two fast-mount caches: `copan_fmc1` and `copan_fmc2`
- Four COPAN drive groups: `dg_c00` through `dg_c03`

¹ For historic reasons, the SGI 400 VTL is sometimes referred to as "COPAN" in literals and the graphical user interfaces.

- Four COPAN volume groups: `vg_c00` through `vg_c03`
- One tape library server: `tape_ls`
- Two tape drive groups: `dg1` and `dg2`
- Four tape volume groups: `vg1a`, `vg1b`, `vg1c`, and `vg2`
- Two migrate groups: `mg0`, `mg_fmc`
- One resource watcher: `rw`
- The DMF administrative directories, as described in "Create the Directories that DMF Requires" on page 13

Depending upon your site-specific configuration, or if you choose a different naming convention ("Naming Convention" on page 10), you must modify the DMF configuration accordingly.

6. Make the following changes to the sample configuration, based on the configuration at your site:

Note: This procedure describes the minimum changes you must make. If your configuration differs from the sample configuration, which is likely, you may need to make more modifications.

- a. For each filesystem that DMF will manage, copy and paste the `/dmi_fs` object, renaming it with the filesystem name. Right-click the `/dmi_fs` object and select:

Configure...
> **Copy**

In the copy, delete the `/dmi_fs` name and replace it with the name of your filesystem. Figure 2-1 shows an example using the name `/myfs`.

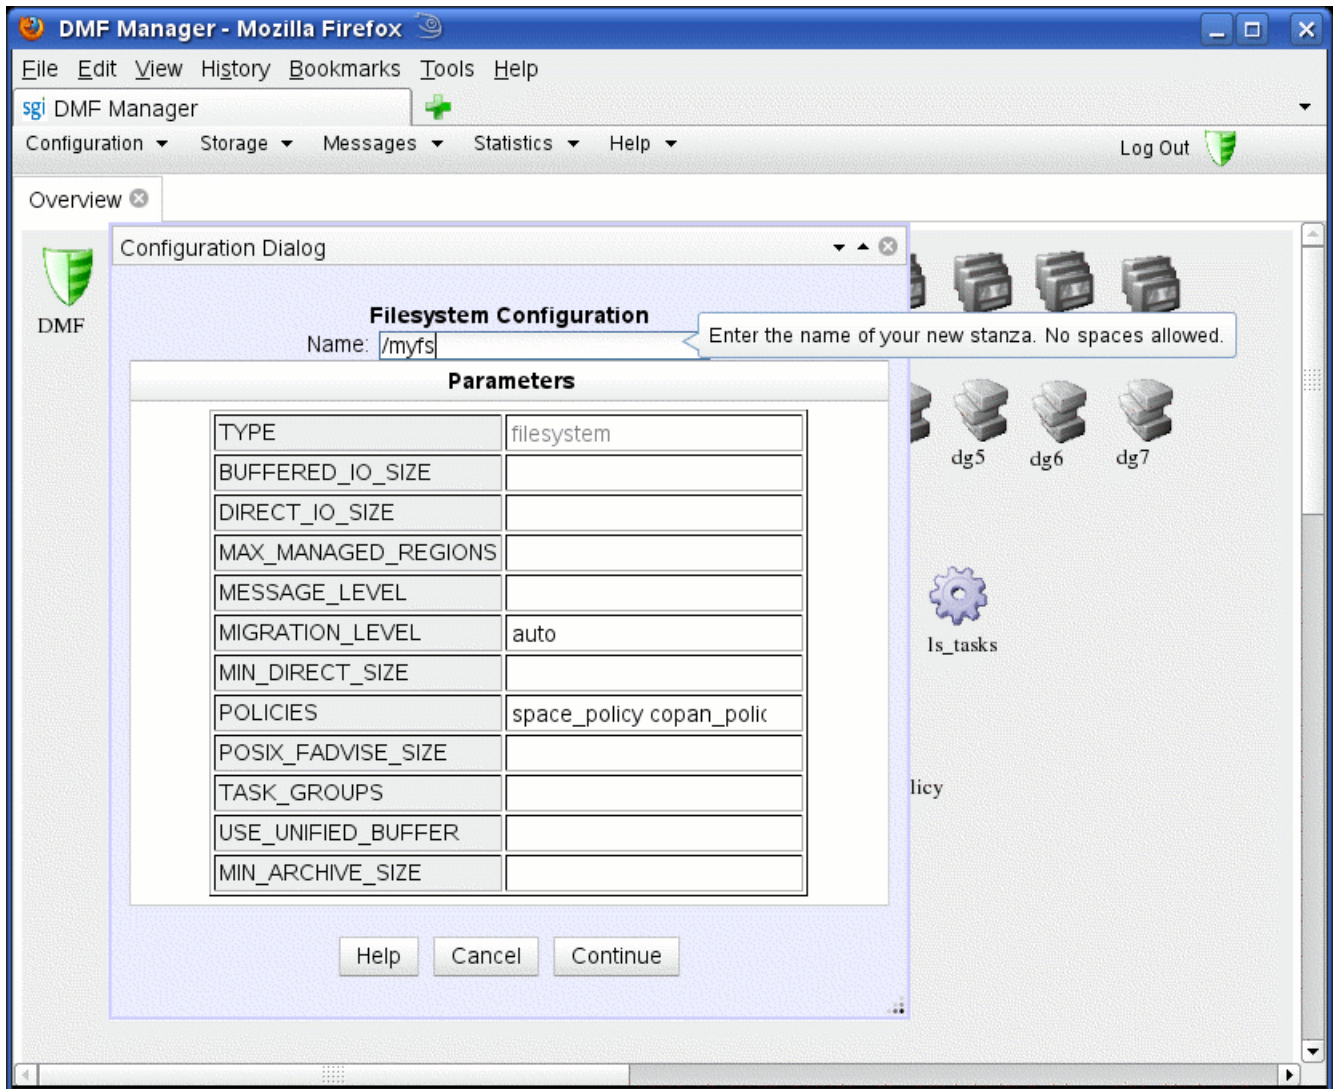


Figure 2-1 Replace the Filesystem Name

You may want to enter a value for the `DIRECT_IO_SIZE` parameter in order to improve I/O performance. Then click **Continue**.

Note: For help on a parameter, right-click the parameter name.

After have created an object for each filesystem, delete the sample `/dmi_fs` object.

- b. Modify the `DUMP_FILE_SYSTEMS` parameter so that it contains all of the filesystems that you want backed up. You should include `/dmf/home`, `/dmf/journals`, and all of the user filesystems that are managed by DMF. Right-click the **dump_tasks** icon and select:

Configure ...
> **Modify**

- c. Modify the `DATABASE_COPIES` parameter so that it specifies the directories in which you want to store copies of the DMF databases (see "Create the Directories that DMF Requires" on page 13). Right-click the **daemon_tasks** icon and select:

Configure ...
> **Modify**

- d. Change the `ADMIN_EMAIL` parameter to an appropriate address for your site. Right-click the DMF shield icon and select:

Configure ...
> **Modify Base**

Delete the default address and enter in the address for your site. Click **Continue**.

- e. Modify the number of drive groups and volume groups according to the number of COPAN shelves in your configuration. Using the procedures in this guide, there should be one drive group and one volume group per shelf, and the migrate groups should contain a balanced number of volume groups. For example, if you have six shelves (0-5) and you are using SGI 400 VTL as permanent storage, you would do the following:

- Delete `dg_c06`, `dg_c07`, `vg_c06`, and `vg_c07`

- Move `vg_c03` to `mg1` (so that `mg0` and `mg1` each have three volume groups)
- f. If you are using the Parallel Data Mover Option:
- Right-click the DMF shield icon and select:
 - Configure ...**
 - > **Modify Base**
- Set `SERVER_NAME` to the system hostname of your DMF server. Click **Continue**.
- For each node in your DMF configuration, right-click anywhere within the **Overview** panel and select:

- Add Object...**
- > **Node...**

Set the name of the object to the node's system hostname. If you are using alternative networks for DMF communication, set `INTERFACE` appropriately.

7. For fast-mount cache, modify the parameters for the permanent storage target (such as a tape library) as appropriate for your site.
8. Validate your changes. Right-click anywhere within the **Overview** panel and select:

- Configure**
- > **Validate Configuration**

You will see errors about OpenVault and a DMF database, because they are not yet configured:

```
ERROR: No OpenVault drives configured.
ERROR: OpenVault server is not up or client is misconfigured.
ERROR: Database (/dmf/home/vtl_ls/tpvrdm.dat) does not exist.
```

The above errors are expected at this point, and you can disregard them. However, if there are any other errors or warnings, fix them and then rerun the validation. Proceed to step 9 when there are no errors and preferably no warnings, other than the expected OpenVault and database errors.

9. Save the configuration. Right-click anywhere within the **Overview** panel and select:

Configure
 > **Apply Configuration**

Note: If you do not use the suggested naming conventions later in this procedure (such as `c00` when you get to step 3g of "Add the Set of DCPs for the Library" on page 80), you must modify the `MOUNT_SERVICE_GROUP` parameter to tell the DMF drive group the name of the OpenVault drive group. If you follow the suggestions in the procedure, no change is required.

Create and Configure Each Virtual Library

Do the following for each virtual library in turn:

- "Invoke the Console and Attach to the VTL Server" on page 31
- "Determine the Migration Tape Size for the Library" on page 36
- "Create the Library" on page 40
- "Configure Virtual Migration Tapes for the Library" on page 50
- "Free Space in Library SGI-C00 for Dump Tapes" on page 55
- "Export Library Access to the DMF Server" on page 63
- "Configure OpenVault for the Library" on page 74

Note: The example output for some commands shown here has been edited for length and readability.

The example screen shots are taken using the 4.5 version of the VTL GUI. Important differences with the 7.5 version of the VTL GUI are noted.

Invoke the Console and Attach to the VTL Server

This section discusses the following:

- "Adding the VTL Server at First Use" on page 31
- "Subsequently Connecting to the VTL Server " on page 34

Adding the VTL Server at First Use

The very first time you access the VTL Console, you must add the VTL Server. Do the following:

1. Start the VTL Console according to the directions in your SGI 400 VTL system administrator's guide.

2. Add the VTL server and log in to the VTL Console:
 - a. Right-click the **COPAN *nmn* Servers** entry in the left-hand tree and select **Add**, as shown in Figure 3-1.
 - b. Enter the following in the dialog box:
 - The name of the VTL server (such as `VTL0`) or its IP address
 - The user name `root`
 - The password `CPN200t`
 - c. Click **OK**.

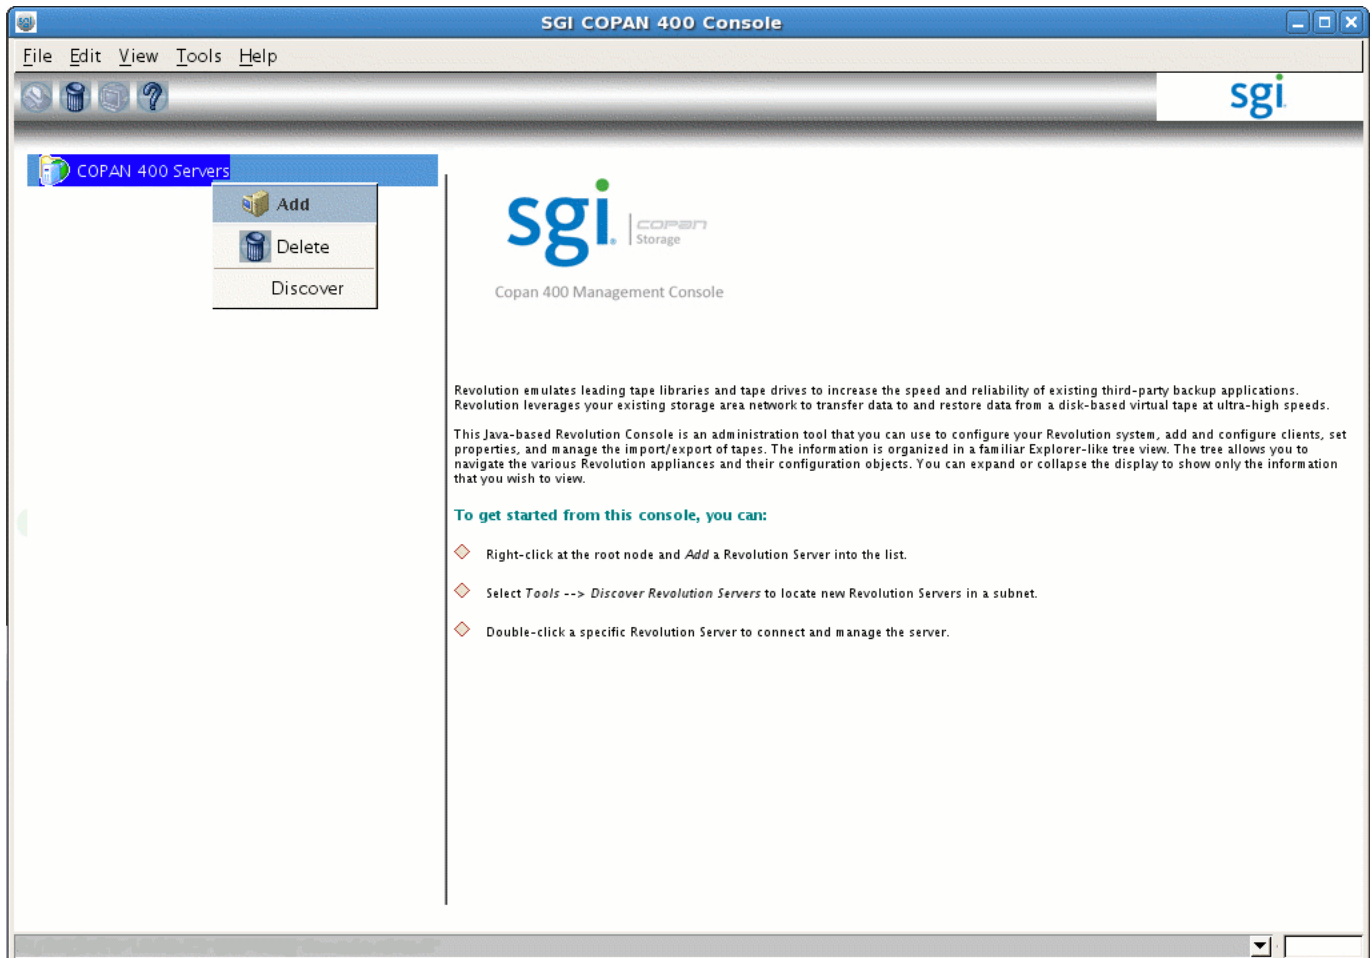


Figure 3-1 Adding the VTL Server

You can then proceed to "Determine the Migration Tape Size for the Library" on page 36.

Subsequently Connecting to the VTL Server

After the initial log in, you can connect to the VTL Console via the VTL server. Do the following:

1. Start the VTL Console according to the directions in your SGI 400 VTL system administrator's guide.
2. Right-click the VTL server name, such as **VTL0**, in the left-hand tree and select **Connect**, as shown in in Figure 3-2.
3. Enter the user name and password.
4. Click **OK**.

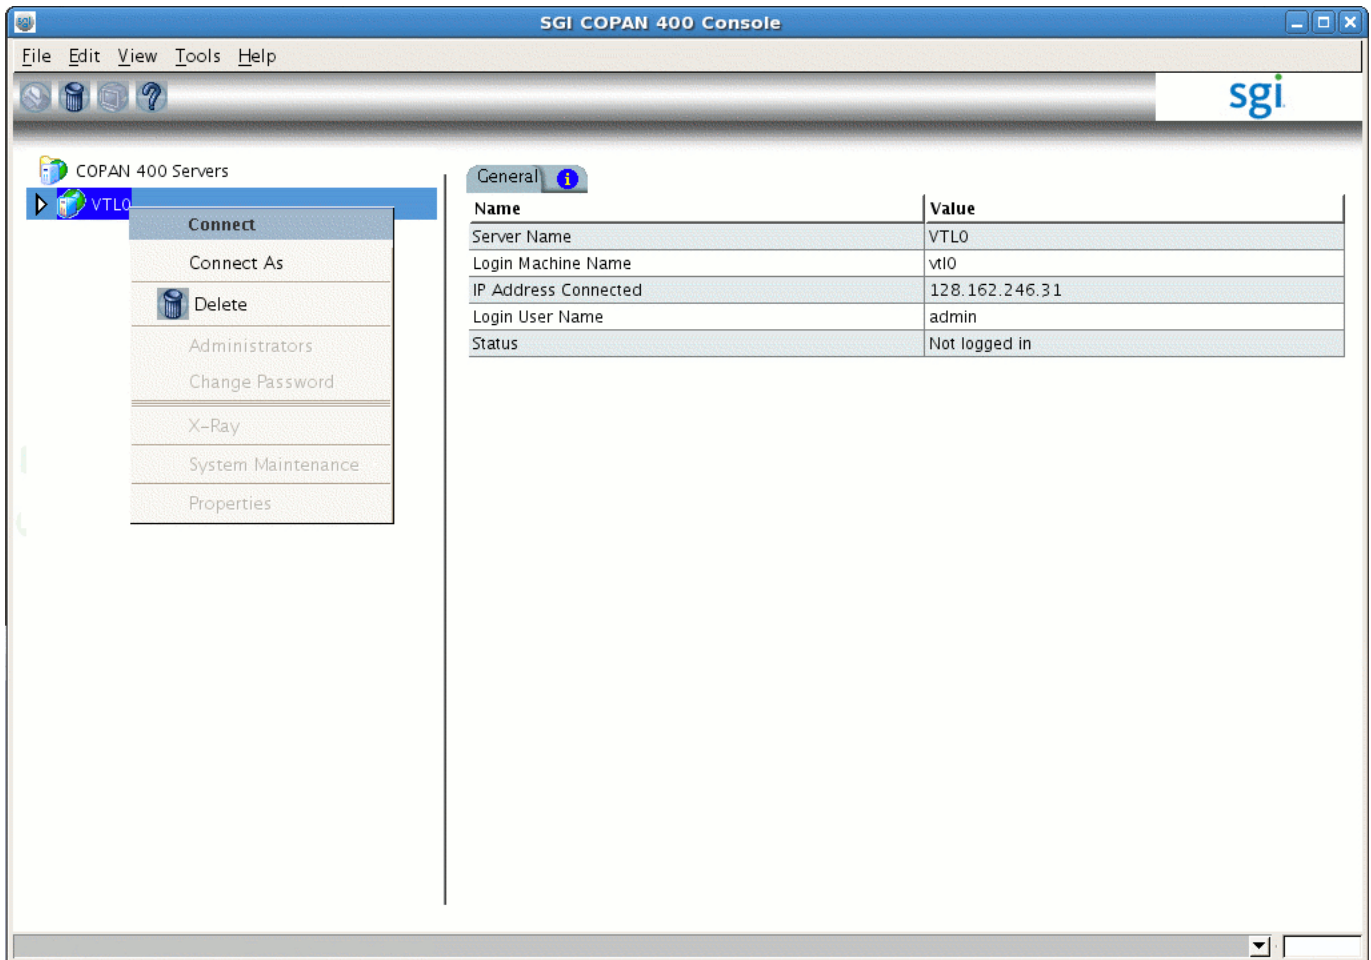


Figure 3-2 Connecting to the VTL Server

You can then proceed to "Determine the Migration Tape Size for the Library" on page 36.

Determine the Migration Tape Size for the Library

SGI 400 VTL shelves use *redundant array of idle disks (RAID)* technology. There are up to eight shelves, and each shelf has a RAID controller and 26 groups of disks. Each group of disks is referred to as a *RAID set*. Each RAID set is configured to hold a fixed number of virtual tapes, all of which are the same size. You must specify the size of the virtual tapes.

Note: It is possible for different shelves in the same COPAN cabinet to contain RAID sets with different capacities, so each shelf could have a different optimal virtual tape size.

You should select the number of migration tapes and the tape size that best accommodate the way you want to use DMF and that will best fit into the usable space on the RAID set. (A certain percentage of space on the RAID is required for overhead. A greater number of tapes requires a larger amount of overhead.)

By default, the Console creates virtual tapes that are the same size as the corresponding physical tapes that are being emulated; because a virtual tape cannot span RAID sets, this can end up wasting otherwise usable RAID space. For example, a RAID set with 1,433 GB of usable space could accommodate only one 768-GB tape, wasting RAID space. Therefore, you want to select a different tape size if you are not obligated to match the physical size. You should choose a number of virtual tapes and a virtual tape size that wastes as little space as possible.

Note: If you want a huge number of tapes (more than the default slot number of 678), including both migration tapes and dump tapes, see Appendix A, "Using an Extremely Large Number of Tapes" on page 107.

1. Use the left-hand tree to access the Fibre Channel devices for the shelf you want to configure, such as the following for shelf 0:

VTL0

- > **Physical Resources**
- > **Storage Devices**
- > **Fibre Channel Devices**
- > **Shelf 0**

2. Record the number that appears in the right panel under the column header **Size (MB)** (such as the number 1,430,414 shown in Figure 3-3). This is the total size of the RAID set.

3: Create and Configure Each Virtual Library

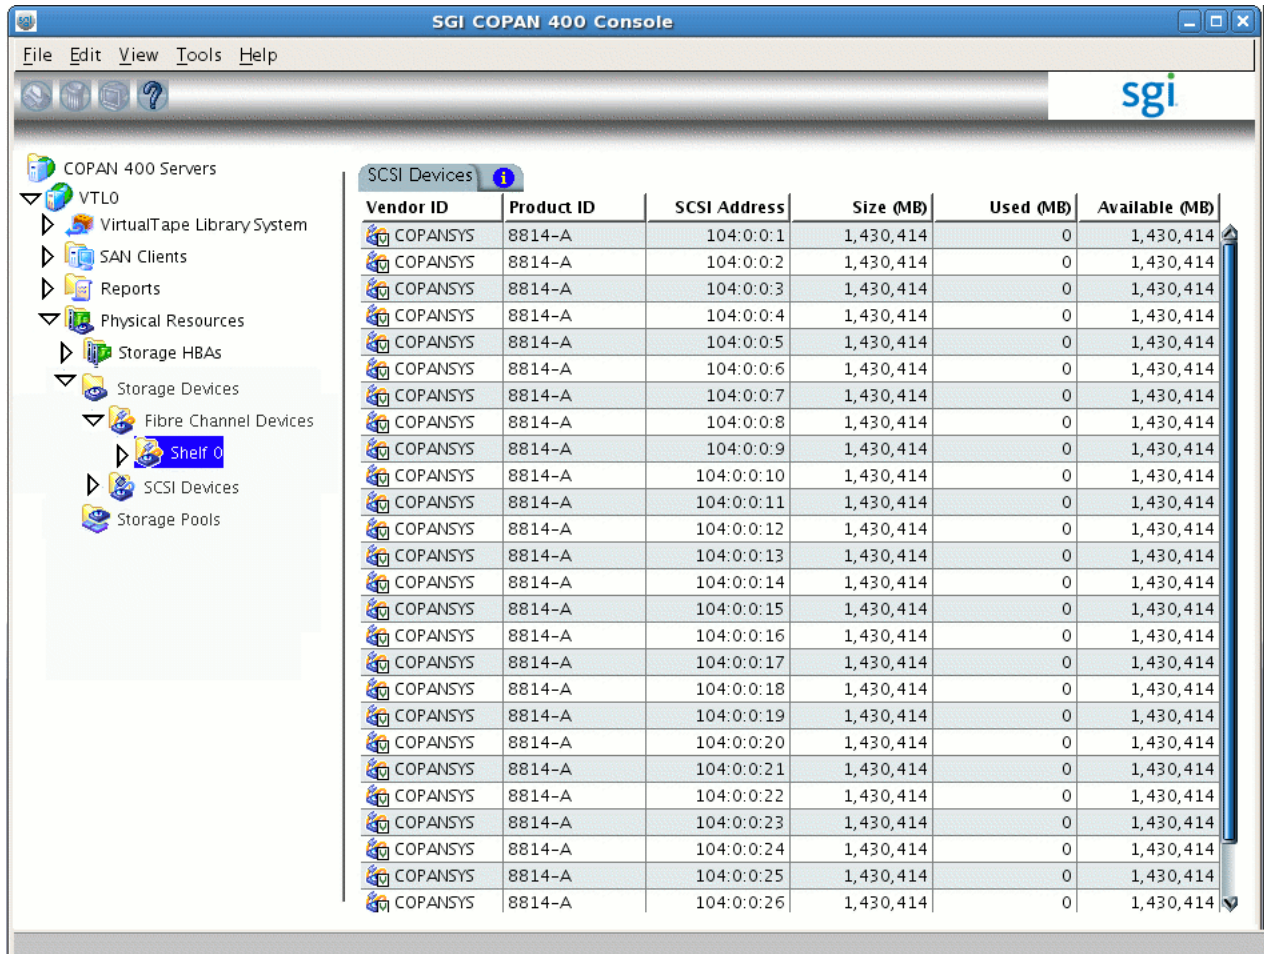


Figure 3-3 Determine the Total Size of the RAID Set

Table 3-1 and Table 3-2 show suggested tape sizes for common sizes of RAID sets.

Table 3-1 Suggested Tape Sizes for Common RAID-Set Sizes (4.5 VTL)

RAID Unit Size (MB)	Number of Tapes Per RAID Unit	Virtual Tape Size (GB)
715,229	1	698
715,229	2	349
1,430,414	2	698
1,430,414	3	465
1,430,414	4	349

Table 3-2 Suggested Tape Sizes for Common RAID-Set Sizes (7.5 VTL)

RAID Unit Size (MB)	Number of Tapes Per RAID Unit	Virtual Tape Size (GB)
8,582,968	1	8,376
8,582,968	2	4,188
8,582,968	3	2,792
8,582,968	4	2,094

If you have a RAID-set size listed in one of the above tables, record the value for use later (in step 7b of "Create the Library" on page 40). Otherwise, go on to step 3 to determine the tape size for the library.

3. Multiply the total size of the RAID set by .976 to get the approximate amount of usable space. You can then divide that result by the desired number of tapes to see what each corresponding virtual tape size is; choose the size that will give you the number of tapes you prefer that wastes the least amount of space.

In general, a larger number of smaller-sized virtual tapes will result in fewer partially-full tapes in the DMF database (and therefore more room for new data). This can potentially provide faster recalls of migrated data because there are more tapes available for reading (you cannot simultaneously write to and read from the same tape). However, you do not want to use tapes that are unreasonably small,

as that might cause excessive mounts and unmounts. Typically, you want to use a migration tape that is at least 500 GB.

Use the following formula as a guideline:

$$(Total_RAID_Set_Size * .976) / \#_Migration_Tapes = Migration_Tape_Size$$

For example, if your total RAID size is 1,430,414 MB and you want two tapes, the resulting virtual tape size would be approximately 698,042.032 MB (or 698 GB).

$$(1430414 \text{ MB} * .976) / 2 = 698042.032 \text{ MB (or 698 GB)}$$

Record this information for use later (in step 7b of "Create the Library" on page 40).

Create the Library

Do the following:

1. Open the left-hand tree and place the cursor over **Virtual Tape Libraries**, such as the following for VTLO:

```
VTLO  
  > Virtual Tape Library System  
    > Virtual Tape Libraries
```

Right-click and select **New** as shown in Figure 3-4.

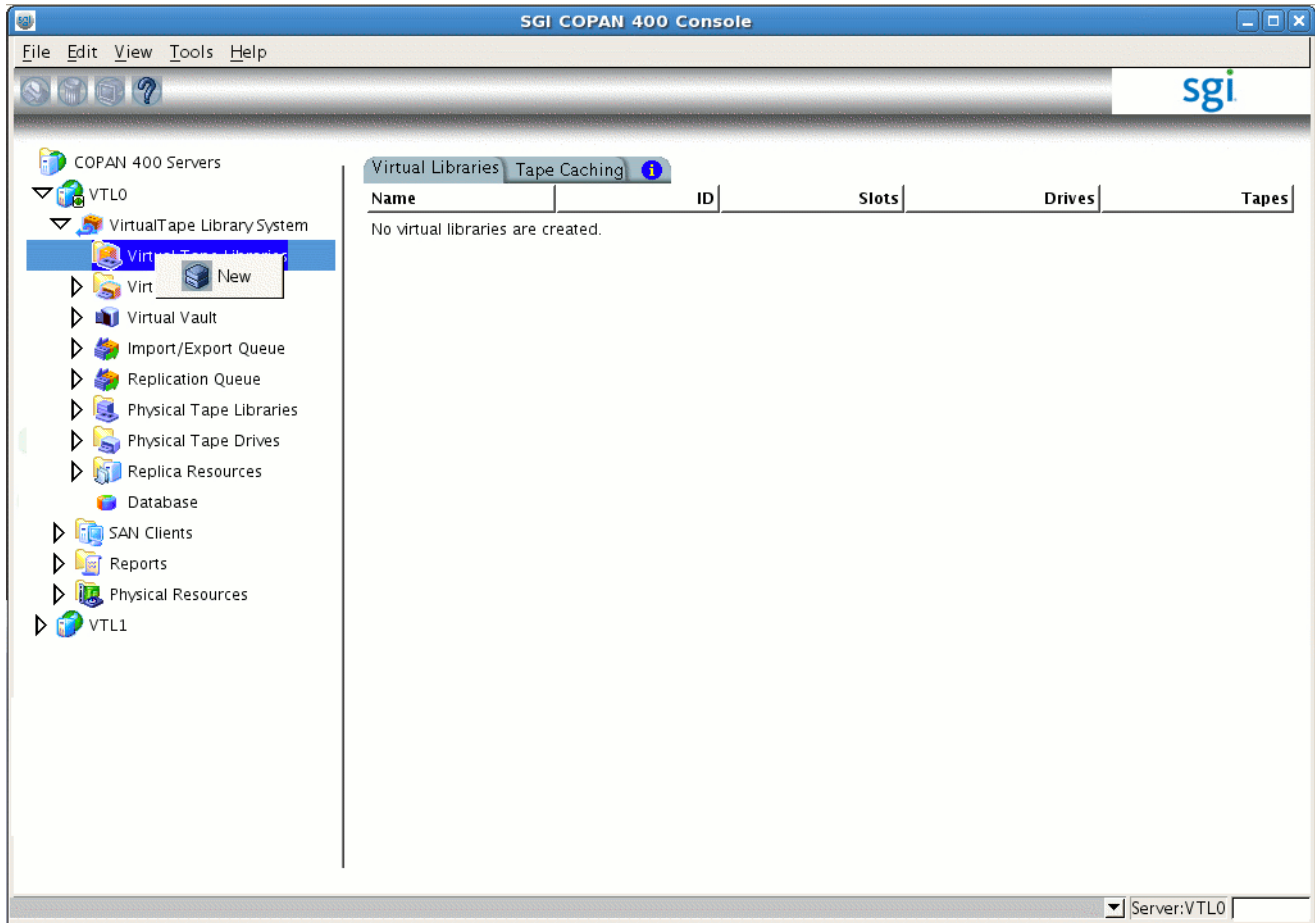


Figure 3-4 Create a New Virtual Library

2. Select a shelf by clicking the check box. Begin with shelf number 0, as shown in Figure 3-5.

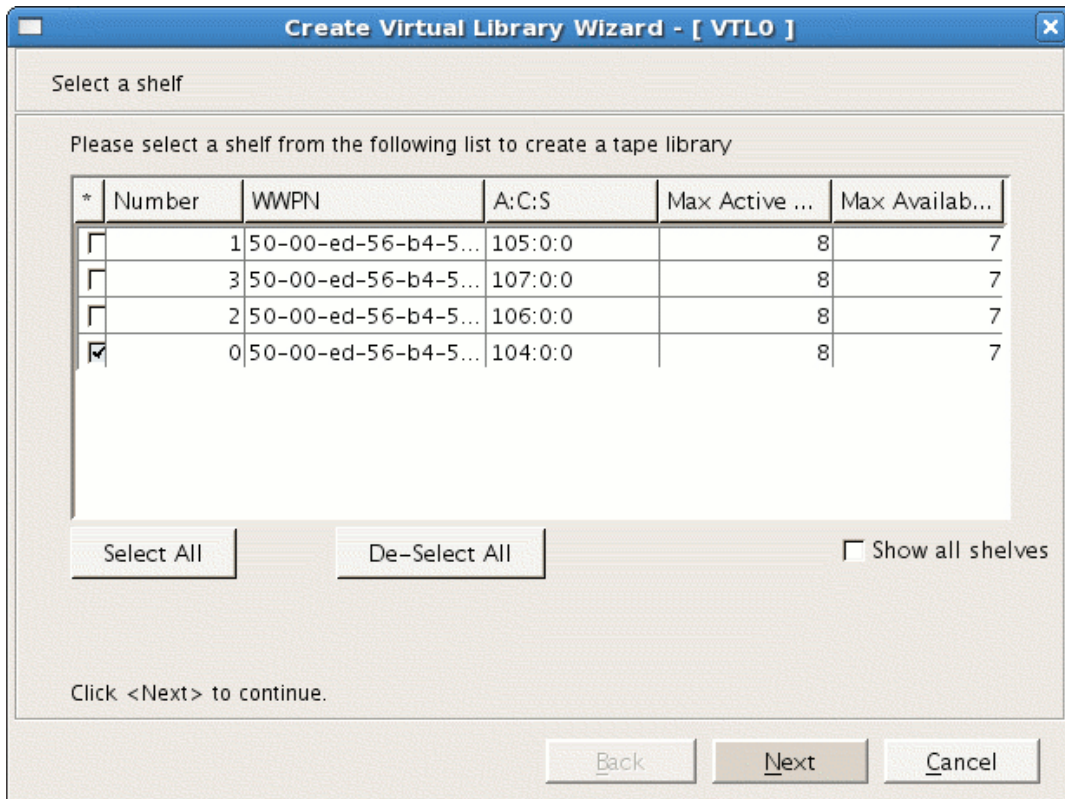


Figure 3-5 Select a Shelf

Note: In the 7.5 VTL GUI, the values for maximum active LUNs and maximum available drives are both 15.

Take note of the **Max Available** drives (for use later in step 4c) and click **Next**.

3. Select the vendor ID type and then name the virtual library **in that order**, as shown in Figure 3-6.

Note: If you name the library first and then select the vendor ID type, the library name will be replaced by the default name. For simplicity during operation, SGI recommends that you **do not use the default name** but instead manually enter a name that corresponds to the shelf number.

- a. Select **SGI** under **Vendor ID**.
- b. Change the library name in the **Virtual Library Name** field so that it corresponds to the SGI 400 VTL shelf number, such as **SGI-C00** for shelf 0, **SGI-C01** for shelf 1, and so on.
- c. Click **Next**.

3: Create and Configure Each Virtual Library

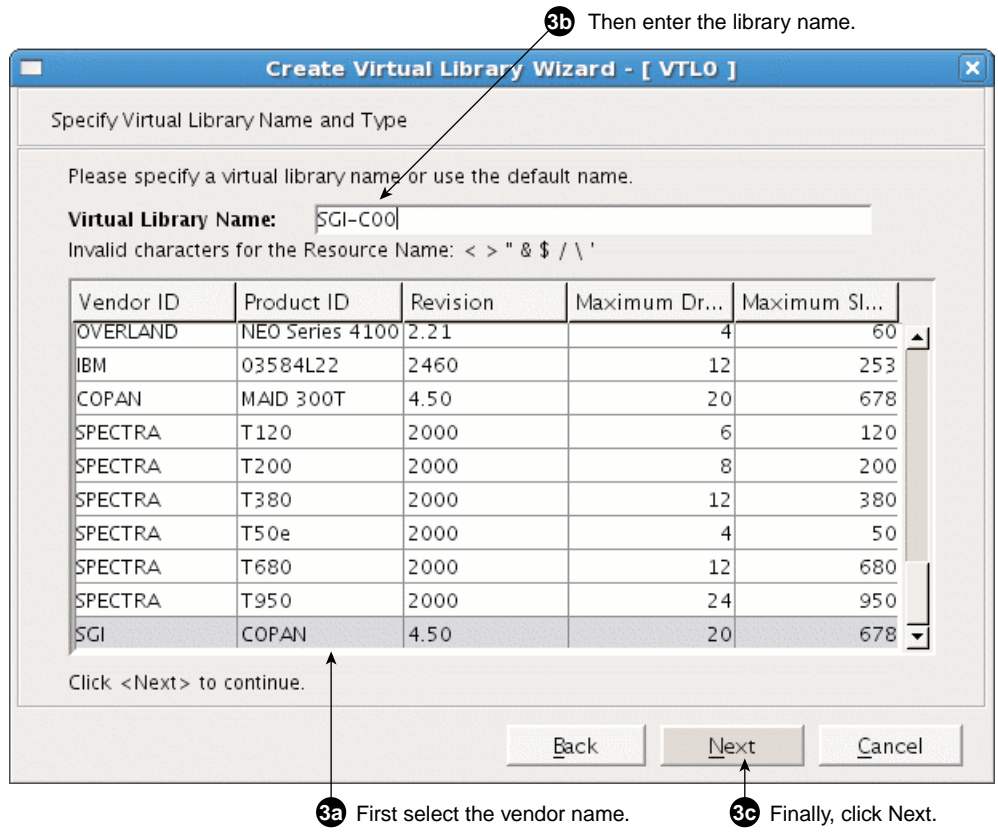


Figure 3-6 Specify the Library Vendor and Name

4. Enter the virtual drive information:
 - a. Select the **IBM ULTRIUM-TD4** product ID row.
 - b. Use the resulting default name in the **Virtual Drive Prefix** field.
 - c. Specify the **Max Available** value noted above in step 2 (rather than use the default) for **Total Virtual Drives**.

5. Do not make any changes on the next set of screens:
 - a. Do not check the **Enable Automated Tape Caching** box; just click **Next**.
 - b. Do not check the **Auto Archive / Replication** box or make any changes on the virtual library information screen; just click **Next**.
 - c. (7.5 VTL GUI only) Do not check the **Enable Tape Duplication** box.
6. Enter the remaining virtual library information:
 - a. Enter the barcode start and end numbers for the shelf, using the following required naming scheme:
 - Start: `CCS9AA`
 - End: `CCS9ZZ`

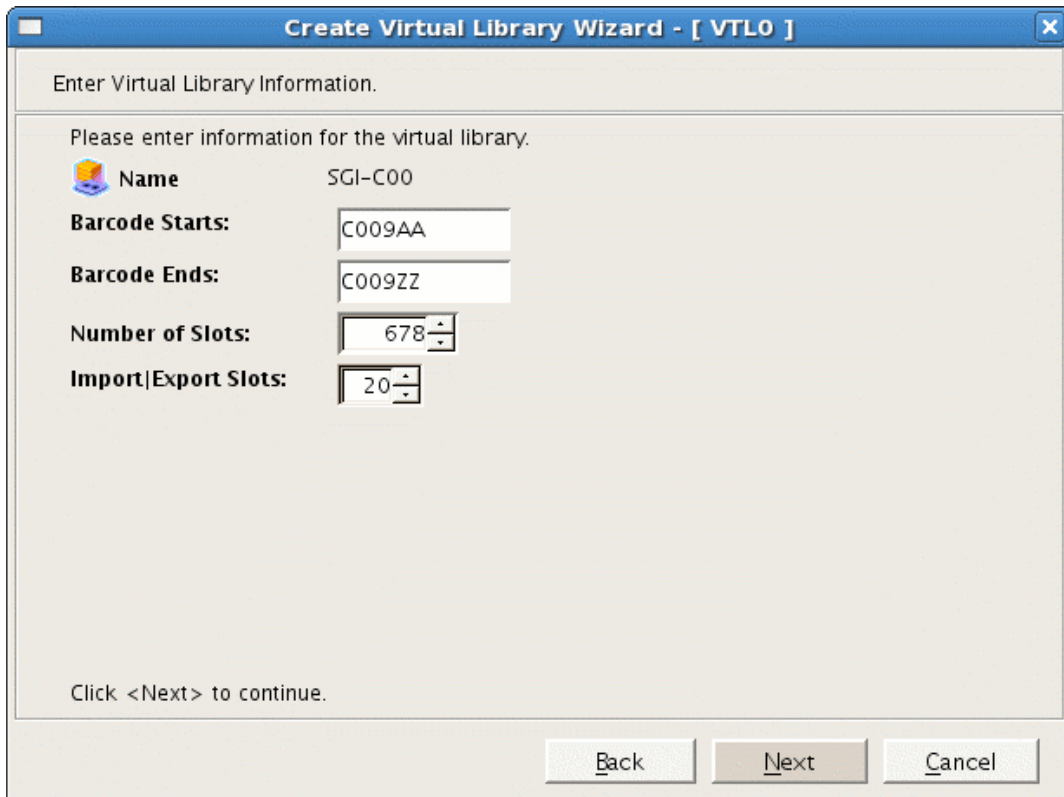
where:

- *CC* uniquely identifies a cabinet at your site. This guide uses `C0` for (“cabinet 0”). You may use a different naming convention if you prefer, using the characters `[0--9A-Z][0--9A-Z]`.
- *S* represents the shelf position in the cabinet. There are up to eight shelves in an SGI 400 VTL cabinet, shelf 0 (bottom) through shelf 7 (top).
- `9AA` and `9ZZ` are required values.

Note: The naming convention allows a script to relabel the barcode (also known as the *volume serial number, VSN*) so that it indicates the specific RAID set in which the tape is stored. This in turn allows DMF to optimize I/O to and from the shelf.


For example, for the virtual library on cabinet 0 and shelf 0 (therefore `SGI-C00`), you would enter the following barcode information, as shown in Figure 3-8:

```
C009AA
C009ZZ
```



Enter Virtual Library Information.

Please enter information for the virtual library.

 **Name** SGI-C00

Barcode Starts: C009AA

Barcode Ends: C009ZZ

Number of Slots: 678

Import|Export Slots: 20

Click <Next> to continue.

Back Next Cancel

Figure 3-8 Enter the Barcode Information

For the virtual library on the next shelf (cabinet 0 and shelf 1, therefore SGI-C01), you would enter the following barcode information:

C019AA
C019ZZ

Note: Verify that you have used the correct shelf number.

- b. Use the defaults for **Number of Slots** and **Import | Export Slots**.

Note: If you must increase the number of slots, see Appendix A, "Using an Extremely Large Number of Tapes" on page 107.

- c. Click **Next**.
7. Enter the virtual tape properties:
 - a. *(4.5 VTL GUI only)* Do not select **Tape Capacity on Demand**.
 - b. Enter the **Maximum Capacity** tape value from "Determine the Migration Tape Size for the Library" on page 36.
 - c. Click **Next**.
8. Verify that the library information displayed matches what you intended to configure and click **Finish** to complete creation of the library, as shown in Figure 3-9.

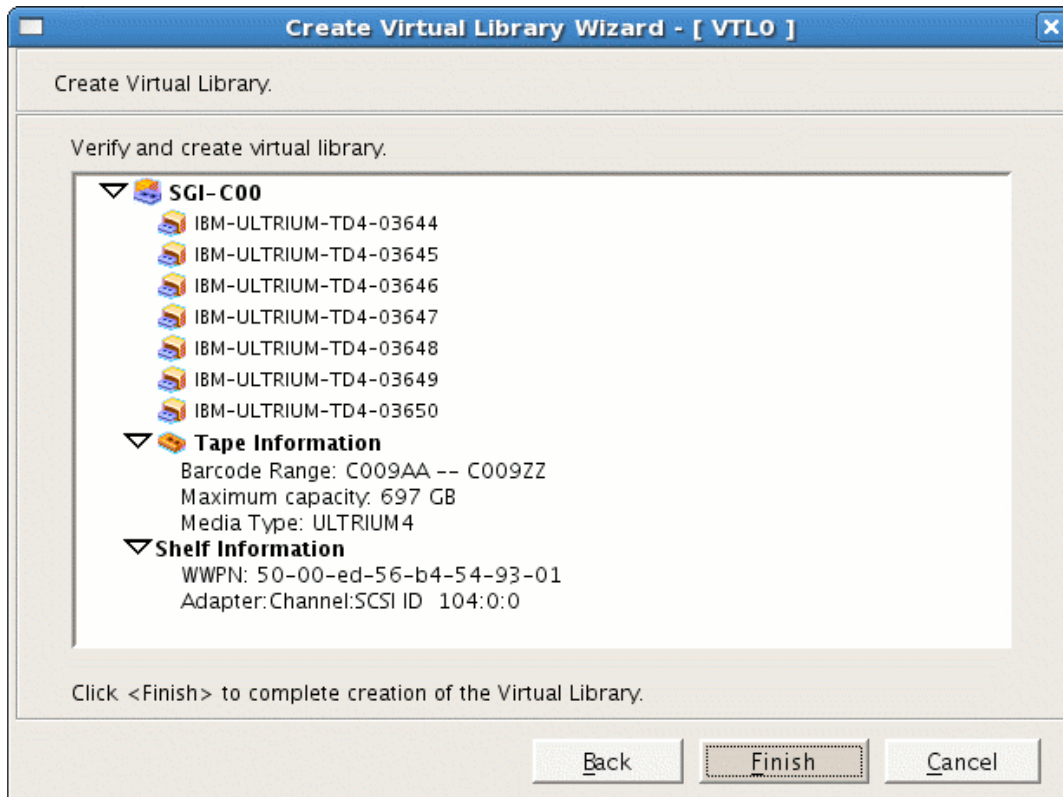


Figure 3-9 Finish Creating the Library

9. Click **Yes** to create the virtual tapes and go to the next section, "Configure Virtual Migration Tapes for the Library" on page 50.

Note: If your site's data would benefit from software compression, you can enable SGI 400 VTL compression mode (and therefore DMF compression) for each application server (VTL0 and VTL1). See Appendix B, "Enabling Compression " on page 109. However, if you are using the 4.5 VTL software, compression carries an overhead CPU penalty and is not appropriate in some cases.

Configure Virtual Migration Tapes for the Library

This section discusses the following steps:

- "Create Migration Tapes" on page 50
- "Relabel the Migration Tapes" on page 54

Create Migration Tapes

Note: This can be a time-consuming process.

Do the following to create migration tapes for the library:

1. Select the default **Express** method for creating tapes and click **Next**.
2. Click **Next** to accept all of the physical devices (selected by default), as shown in Figure 3-10.

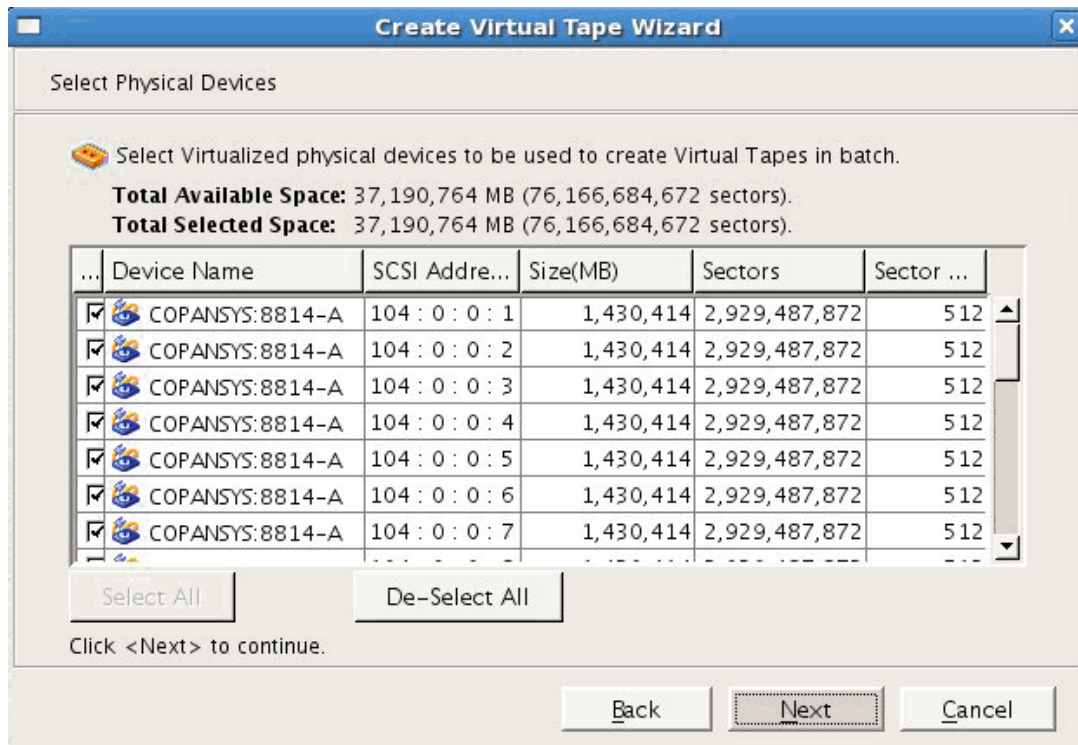


Figure 3-10 Select Physical Devices

- Enter in the number of tapes to be configured and click **Next**, as shown in Figure 3-11.

Note: This step is creating the actual logical tape volumes on the RAID, which will require that disks be spun up; the more tapes you create, the more time will be required. With the 7.5 VTL GUI, you can create all of the tapes in one batch at a convenient time. However, if you are using the 4.5 VTL GUI, you cannot enter a number greater than 50. For example, if there are 52 tapes to be configured, you must do them in two groups (such as 2 and 50) by repeating some of these steps.

3: Create and Configure Each Virtual Library

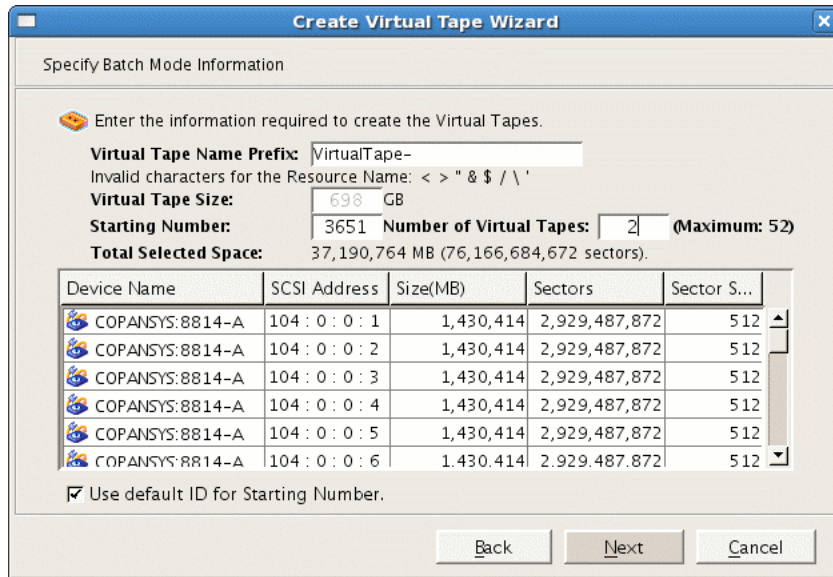


Figure 3-11 Enter the Number of Tapes to Be Configured

4. Do not change the **Barcode Starts** information (originally input in step 6a of "Create the Library" on page 40) because the Console will continue from the correct point. Do not click **Use the following barcode range for this batch**.

Click **Next** to select the defaults for the barcode range option.

5. Verify the tapes to be created and select **Finish** to create the first group of tapes for this library, as shown in Figure 3-12.

Note: The more tapes created, the longer this step will take. For example, depending upon the disk and power configuration, creating 50 tapes can take more than 15 minutes (during which time there will be no updates displayed).

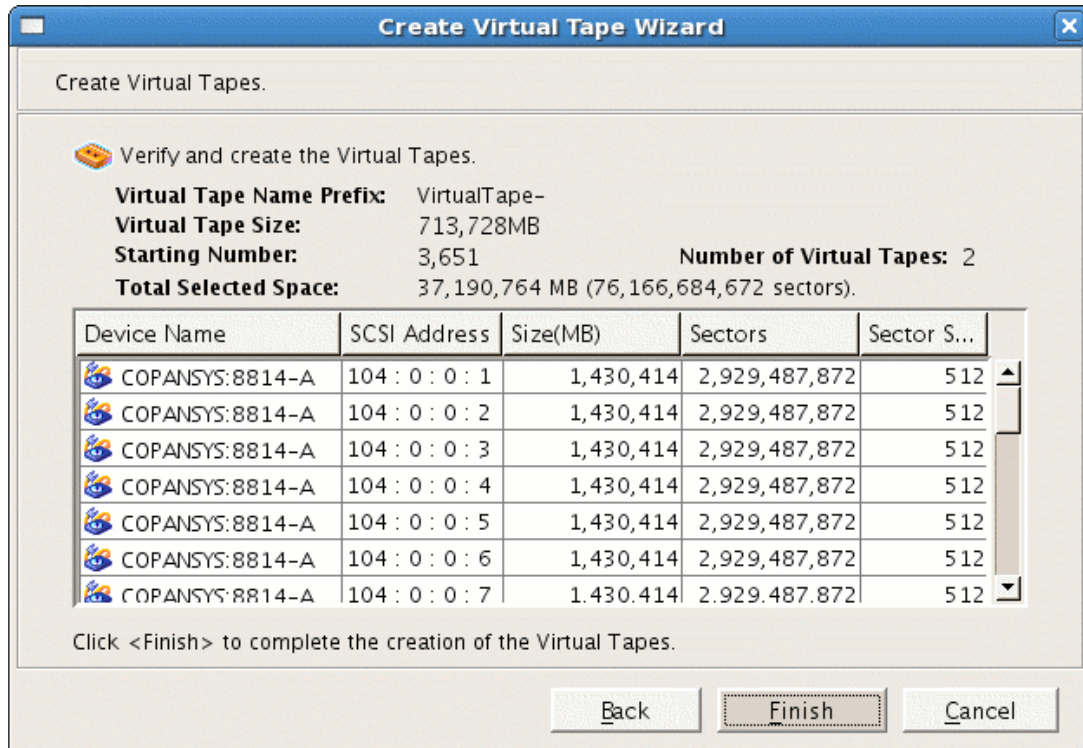


Figure 3-12 Verify and Create Tapes

6. Click **OK** to complete the step.
7. Open the left-hand tree and place the cursor over **Tapes**, such as the following for library SGI-C00:

```

VTL0
  > Virtual Tape Library System
    > Virtual Tape Libraries
      > SGI-C00
        > Tapes
    
```

Right-click and select **New Tape(s)**.

8. Repeat steps 1 through 7 above to create the remaining tapes for this library. (For example, if there are 52 total tapes possible and you entered 2 previously for step 3, enter 50 this time.) The more tapes created, the longer the **Finish** step will take.

Relabel the Migration Tapes

The DMF library server uses the barcode when selecting tapes. You must run a script to relabel each virtual tape's barcode so that it indicates the specific RAID set in which the tape is stored, in order to optimize I/O to and from the shelf. (For example, C009AA might be transformed into C00AAA, where the fourth character identifies the RAID set to DMF.)

Do the following:

1. Ensure that there is no other activity on the Console or on the VTL application server that controls the library you have just created (such as VTL0 for library SGI-C00). This is a requirement for the duration of this procedure.
2. Run the `dmreadycopantapes` script on the VTL application server from the `/usr/local/dmf` directory.

Note: This step requires that you previously downloaded the required files from the DMF server to each VTL application server, as specified in "Copy the Support Files to the VTL Application Servers" on page 13.

The script will output an `executed successfully` line for each barcode that is changed. Depending upon the number of tapes and drives, this can be a time-consuming process. For example:

- 4.5 VTL software:

```
[root@VTL0 vtl]# cd /usr/local/dmf
[root@VTL0 dmf]# ./dmreadycopantapes

Command: settapeproperty executed successfully.
...
Command: settapeproperty executed successfully.
[root@VTL0 dmf]#
```

- 7.5 VTL software:

```
[root@VTL0 vtl]# cd /usr/local/dmf
[root@VTL0 dmf]# ./dmreadycopantapes

/usr/local/sgi/vtl/bin/iscli settapeproperty --server-name=127.0.0.1 --vdev=10000121 --barcode=C00AAZ --force
...
/usr/local/sgi/vtl/bin/iscli settapeproperty --server-name=127.0.0.1 --vdev=10000172 --barcode=C00ZCE --force
[root@VTL0 dmf]#
```

For more information, add the `-v` option.



Caution: The script must complete successfully in order for DMF to function properly. If there are errors, fix them before continuing. If necessary, contact SGI Support.

Free Space in Library SGI-C00 for Dump Tapes

Note: You will only perform this step once, after creating the first library on shelf 0. All of the dump tapes must be on a single shelf. The sample DMF configuration file defines this to be shelf 0. If you want to use a different shelf, you must modify the DMF configuration file `drivegroup` object and these procedures accordingly.

In this step, you will free disk space from only one library (SGI-C00) so that the space can later be reclaimed for the purpose of dump tapes (which will occur in Chapter 4, "Configure Virtual Dump Tapes" on page 89.) The dump tapes must be on one library only. To avoid problems in case of potential hardware failure, the dump tapes should be spread across the RAID sets on the shelf. To make the best use of disk space, the dump tapes should be sized to fit the amount of data needed to accommodate all of the dumps for a day's full backup.

Note: If the dumps for a day's backup require more than one tape, it will increase the complexity and time requirements of restores.

Do the following:

1. Determine how many migration tapes must be deleted in order to free up the required amount of dump space per day that you established in "Determine the Backup Requirements For Your Site" on page 15. Do the following:
 - a. Divide the migration tape size (determined in "Determine the Migration Tape Size for the Library" on page 36) by the dump space needed per day (estimated in "Space Required for the Daily Backup" on page 16) to determine the number of dump tapes that you can create from the space occupied by one migration tape:

$$\text{Migration_Tape_Size} / \text{Dump_Space_Needed_Per_Day} = \#_Dump_Tapes_Per_Migration_Tape$$

For example, suppose you determined that your migration tape size is 698 GB and you estimated that you require 100 GB for each day's dump set. You can therefore make six dump tapes out of each migration tape:

$$698 / 100 = 6.98 \text{ (rounds down to 6 dump tapes per migration tape)}$$

Note: This means that there are 98 GB of unused space that are left on the RAID set. In step 2 below, you will make use of this space in the final tape size.

- b. Divide the total number of dump tapes required (established in "Number of Dump Tapes Required" on page 16) by the number of dump tapes per migration tape to determine the number of migration tapes to be deleted:

$$\#_Dump_Tapes / \#_Dump_Tapes_Per_Migration_Tape = \#_Migration_Tapes_To_Be_Deleted$$

For example, using the default DMF configuration file for SGI 400 VTL, you will retain the dumps for 28 days (requiring a minimum of 29 tapes). Because you can make 6 dump tapes out of each migration tape, you therefore need to delete 5 migration tapes:

$$29 / 6 = 4.83 \text{ (rounds up to 5 migration tapes to be deleted)}$$

Because you must delete 5 migration tapes and you can get 6 dump tapes per migration tape, you will end up with a total of 30 dump tapes, which is 1 greater than the absolute minimum of 29; having an extra dump tape is a good precaution.

2. Determine the size of the dump tape so that it includes the space that was left over after the calculation you made in step 1a above. To do this, divide the amount of space left over in the calculation of the number of dump tapes in step 1a above (in this case, 98 GB) by the number of dump tapes created per migration tape and add it to the original dump space needed per day. This will determine the final dump tape size:

$$\text{Total_Left_Over} / \#_Dump_Tapes_Per_Migration_Tape = \text{Left_Over_Per_Dump_Tape}$$

$$\text{Left_Over_Per_Dump_Tape} + \text{Dump_Space_Needed_Per_Day} = \text{Dump_Tape_Size}$$

For example:

$$98 / 6 = 16.333333 \text{ (rounds down to 16)}$$

$$16 + 100 = 116$$

Take note of the *Dump_Tape_Size* value for use later as the **Maximum Capacity** field in step 2 of "Create the Dump Tapes" on page 89.

Note: It is possible to later reclaim more space for dump tapes, as long as there is still space on the shelf. However, it is more complicated. See Appendix C, "Reclaiming Space for Dumps" on page 111.

3. Delete the required number of migration tapes individually, selecting tapes across the RAID set by choosing tapes with barcodes that have different letters of the alphabet for the fourth character (the fourth character specifies the RAID set). Doing so will reduce the risk due to possible hardware failures.

For example, if you wanted to delete 5 migration tapes, you might choose to delete the following (highlighting the fourth character):

```
C00AAA
C00DAD
C00GAG
C00JAJ
C00MAM
  ^
```

To delete a given tape, do the following:

- "Delete Tapes with the 4.5 VTL GUI" on page 58
- "Delete Tapes with the 7.5 VTL GUI" on page 60

Delete Tapes with the 4.5 VTL GUI

To delete tapes with the 4.5 VTL GUI, do the following:

1. Open the left-hand tree and place the cursor over the library name, such as the following for library `SGI-C00`:

```
VTL0  
  > Virtual Tape Library System  
    > Virtual Tape Libraries  
      > SGI-C00
```

2. Right-click and select **Delete Tape(s)**.
- 3.

Click the boxes for the tapes you want to delete and click **Delete**, as show in Figure 3-13.

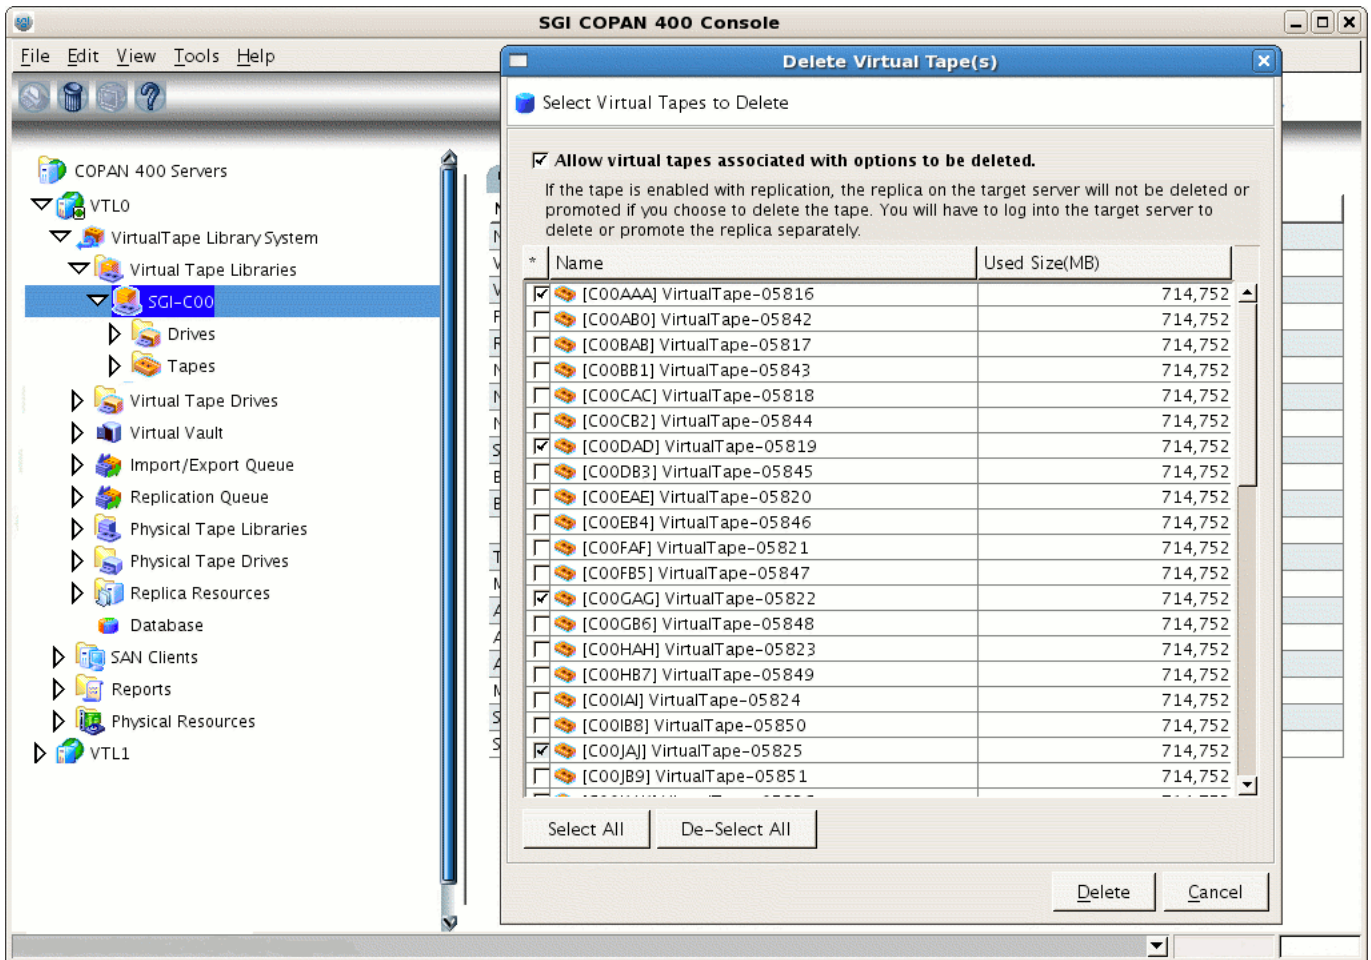


Figure 3-13 Delete Selected Migration Tapes to Free Space Needed for Dumps (4.5 GUI)

4. In the **WARNING** dialog, type **YES** to confirm that you want to delete the tapes and click **OK**, as shown in Figure 3-14. (There is no data on the tapes at this point, so it is safe to delete them.)

3: Create and Configure Each Virtual Library

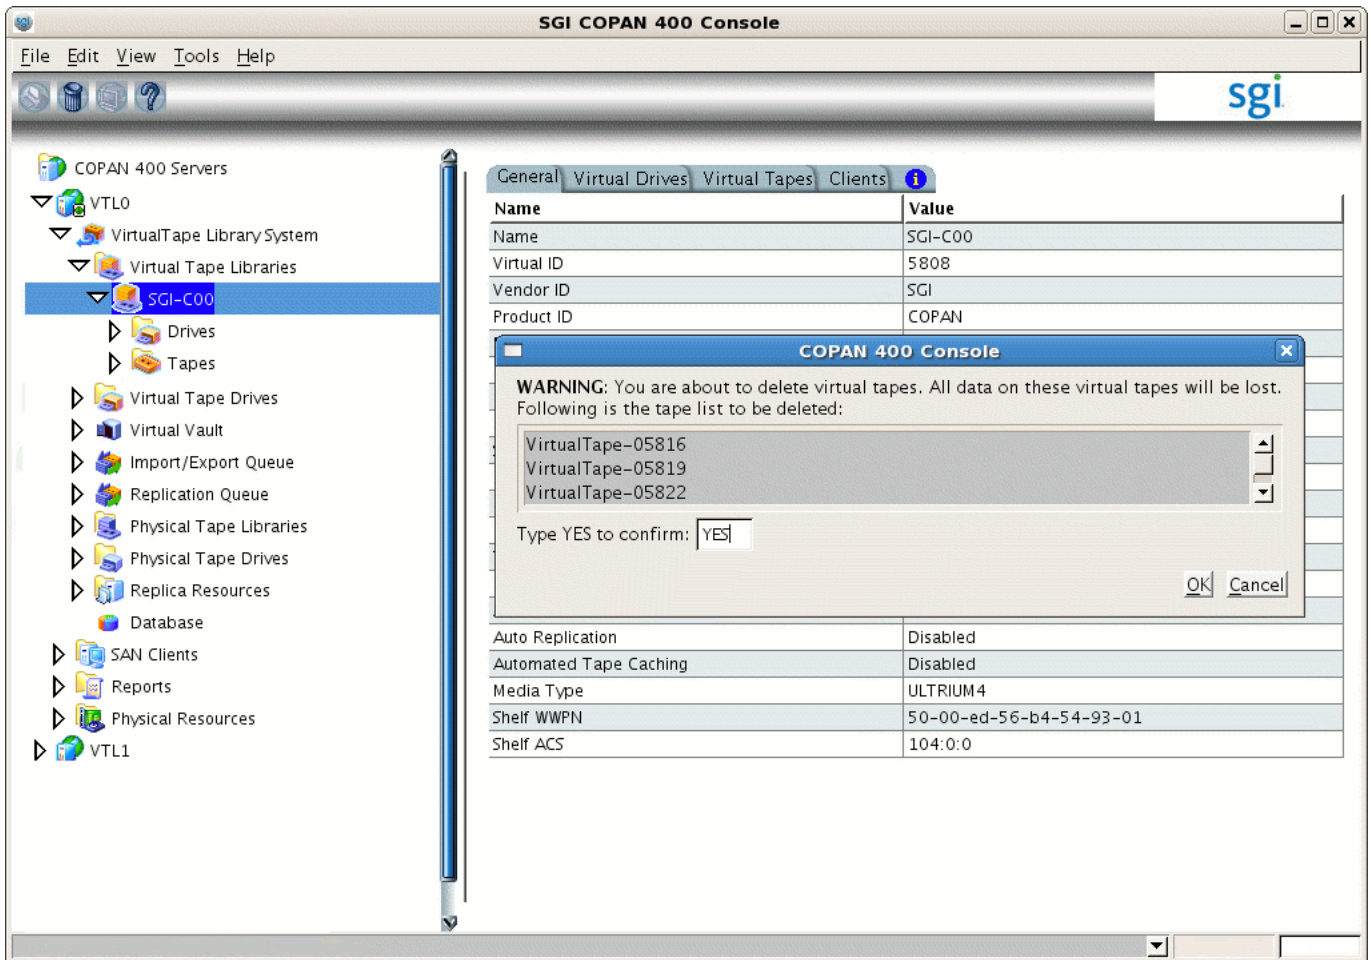


Figure 3-14 Confirm the Deletion of a Tape on SGI-C00

Delete Tapes with the 7.5 VTL GUI

To delete tapes with the 7.5 VTL GUI, do the following:

1. Open the left-hand tree tapes display for the library, such as the following for library SGI-C00:

VTLO
 > **Virtual Tape Library System**
 > **Virtual Tape Libraries**
 > **SGI-C00**
 > **Tapes**

- In the right-hand panel, click the names of the tapes that you want to delete. Use **Ctrl+click** to select multiple tapes or **Ctrl+Shift+click** to select multiple tapes in sequence. For example, Figure 3-15 shows two tapes that are highlighted.

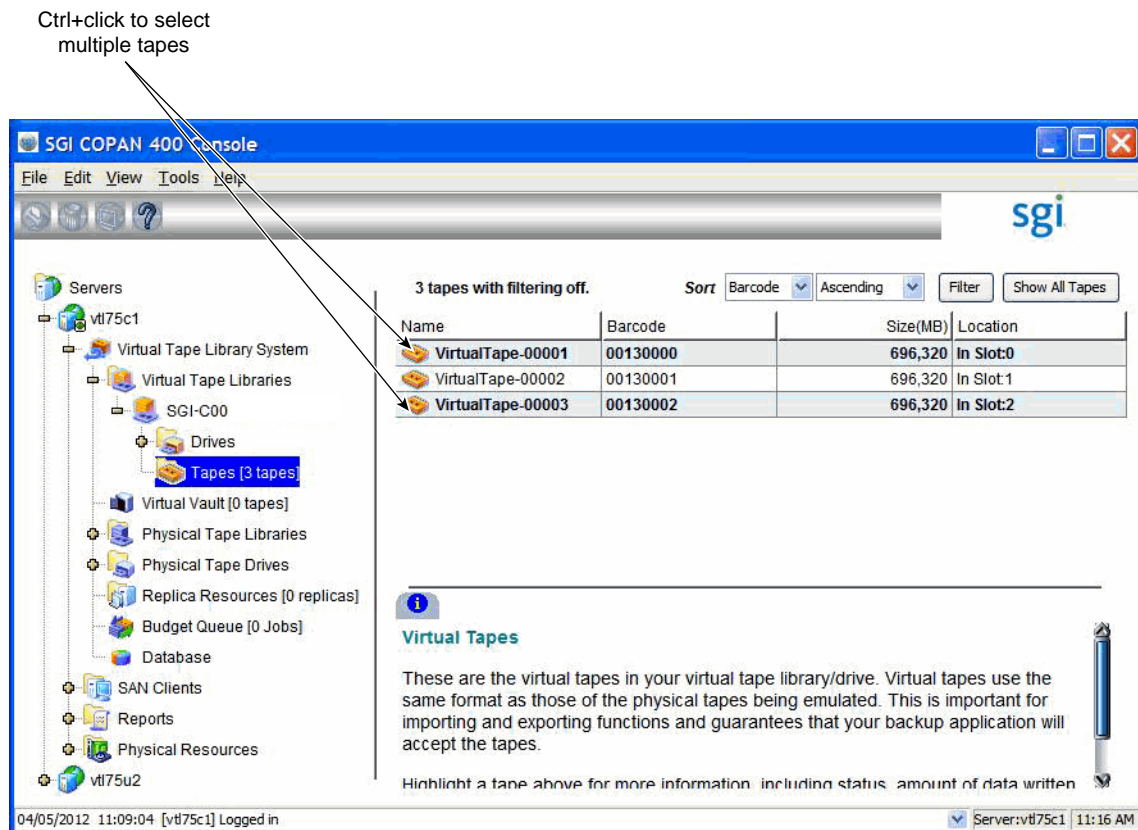


Figure 3-15 Selecting Tapes to Delete (7.5 GUI)

- Right-click and select **Delete Tapes** in the pop-up window.

4. Verify that the correct tapes are selected. For example, Figure 3-16 shows that three tapes are selected for deletion.

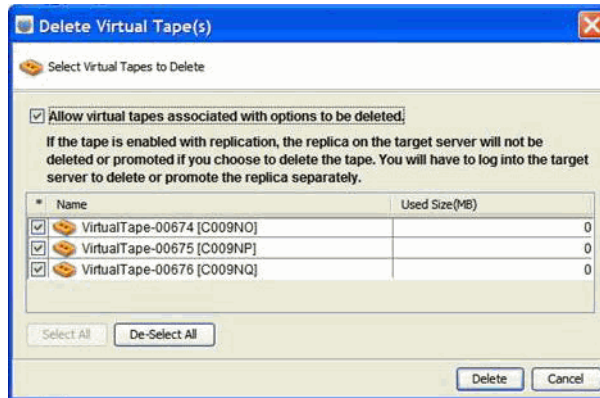


Figure 3-16 Delete Tapes in the 7.5 VTL GUI

5. Click **Delete**.

Export Library Access to the DMF Server

This section exports SGI 400 VTL access to a particular DMF server port via a particular SGI 400 VTL server port using world wide port names (WWPNs). To export library access, you must do the following:

- Identify the ports you want to connect on the DMF server (which uses the *initiator WWPN*) and on the SGI 400 VTL server (which uses the *target WWPN*). There are four ports available for each VTL server (`vtt10` and `vtt11`); the number of ports available on the DMF server will vary by site.
- Configure a target WWPN on the SGI 400 VTL shelf (such as shelf 0) to connect the pair of desired ports. This target WWPN is internal to the SGI 400 VTL cabinet and is similar to a switch's zone.

Note: For optimum bandwidth, SGI recommends that you configure each shelf target WWPN to a unique pair of ports. Ideally, each target WWPN on the VTL server should connect to a different initiator WWPN on the DMF server. However, if this is impossible due to insufficient ports on the DMF server, you should configure each shelf individually through a unique target WWPN (rather than configuring multiple shelves on a single target WWPN) in order to use unique initiator/target pairs.

The SGI 400 VTL cabinet has a built-in zoning capability that lets you control which host ports in the SAN can see which virtual libraries. By default, the libraries are not visible anywhere; you must use the steps in this section to configure each library to accept Fibre Channel commands from only the particular port on the particular host that will be running OpenVault, thereby preventing all other initiator WWPNs from causing problems. Essentially, you are creating a Fibre Channel zone where the only members in the zone are the OpenVault host bus adapter (HBA) and the virtual library.

Note: To accomplish this, you must add your own SAN clients. You should not use the **EveryoneFC** client because it allows every HBA in every host on the Fibre Channel SAN to see the virtual library, which could interfere with legitimate users of the library (for example, by inappropriately rewinding tapes while the legitimate user is writing them).

Do the following:

1. Log in to the DMF server and display the WWPN of the DMF server port that you want to connect to the SGI 400 VTL shelf for this library. To do this, view the contents of the appropriate `port_name` file:

```
dmfserver# cat /sys/class/fc_host/port/port_name
```

In this case, suppose you connected DMF server `burn` on port `host5` to shelf 0 (library `SGI-C00`). To determine the WWPN of `host5`, view the contents of its `port_name` file. For example:

```
burn# cat /sys/class/fc_host/host5/port_name  
0x210100e08b3c337d
```

You will use the WWPN (in this case, `210100e08b3c337d`) later in step 5.

2. In the left-hand panel of the Console, right-click **SAN Clients** and select **Add**.
3. Enter the name of the DMF server in the **Client Name** field, click **Find**, then click **Next**.

Note: If the client WWPN is in a zone, the Console will automatically let you select initiators only from that zone and will skip this step.

4. Select **Fibre Channel** and click **Next**.
5. Select the initiator WWPN that corresponds to the WWPN of the DMF server port that you displayed in step 1 above.

Note: If there is only one initiator WWPN in the client, the Console will automatically select it for you and this dialog will not be displayed.

For example, step 1 displayed `210100e08b3c337d`, so you would check that initiator WWPN line in the Console and click **Next**, as shown in Figure 3-17. Figure 3-18 shows the relationship of the **Initiator WWPN** field to the DMF server port.

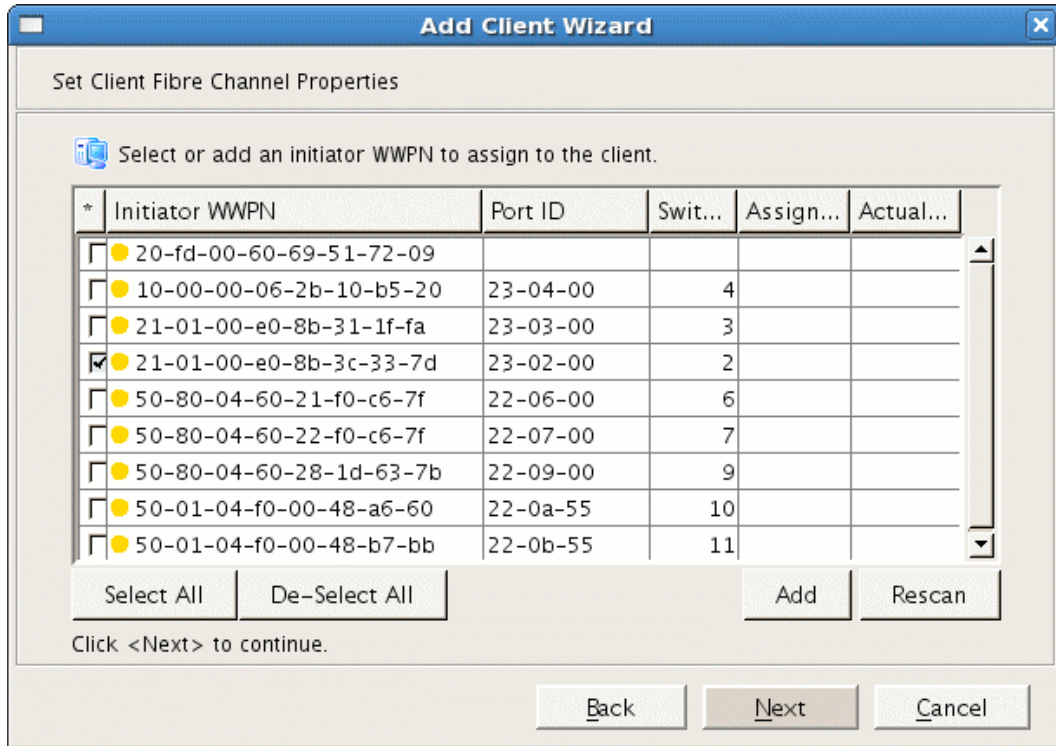


Figure 3-17 Select the Initiator WWPN

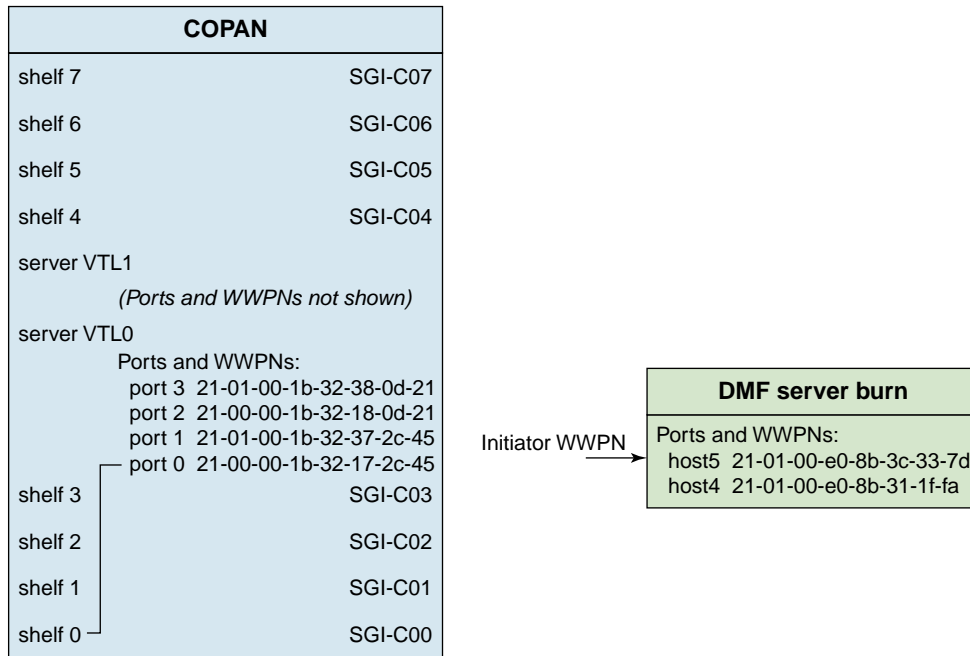


Figure 3-18 Relationship of the "Initiator WWPN" Field to the DMF Server Port

6. Do not check the boxes for **Enable Volume Set Addressing** or **Enable Celerra Support**; just click **Next**.
7. Verify that the client name and initiator WWPN information are as you intended and click **Finish**. Figure 3-19 shows an example.

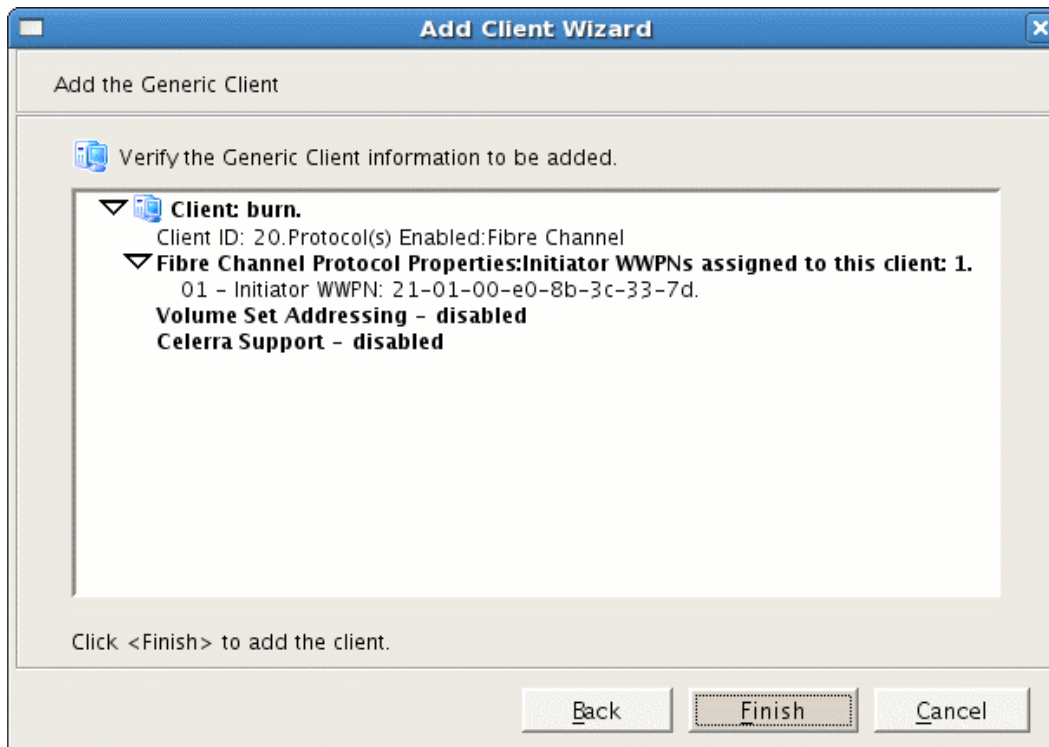


Figure 3-19 Verify the DMF Server as the Client

Note: The 7.5 VTL GUI also displays the following:

IBM System i Support — disabled

8. In the left-hand panel of the Console, place the cursor over the name of the DMF server under **SAN Clients**. For example, for DMF server **burn**:

```

VTLO
  > SAN Clients
    > burn

```

Right-click and select **Assign**.

9. Select the library name, such as SGI-C00 for library 0 on shelf 0. (The green dot indicates that the resource is available and has not been assigned to a client.) Click **Next**. Figure 3-20 shows an example for SGI-C00. Figure 3-21 shows the relationship of the **Name** field to the SGI 400 VTL.

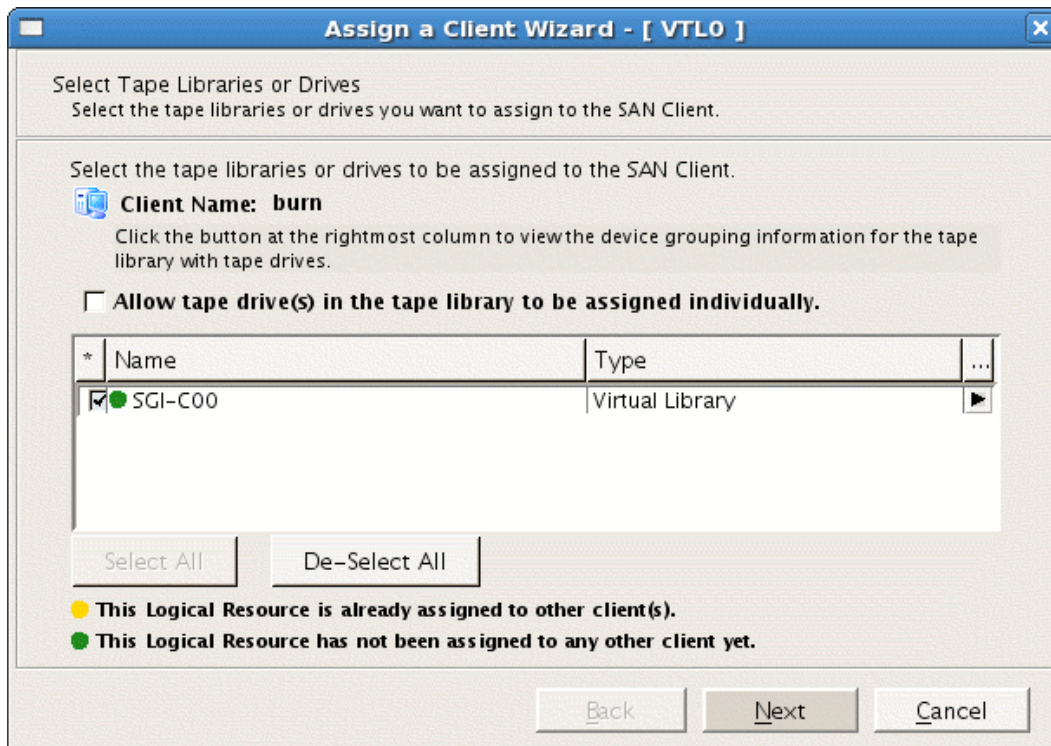


Figure 3-20 Select the Library to Assign to the SAN

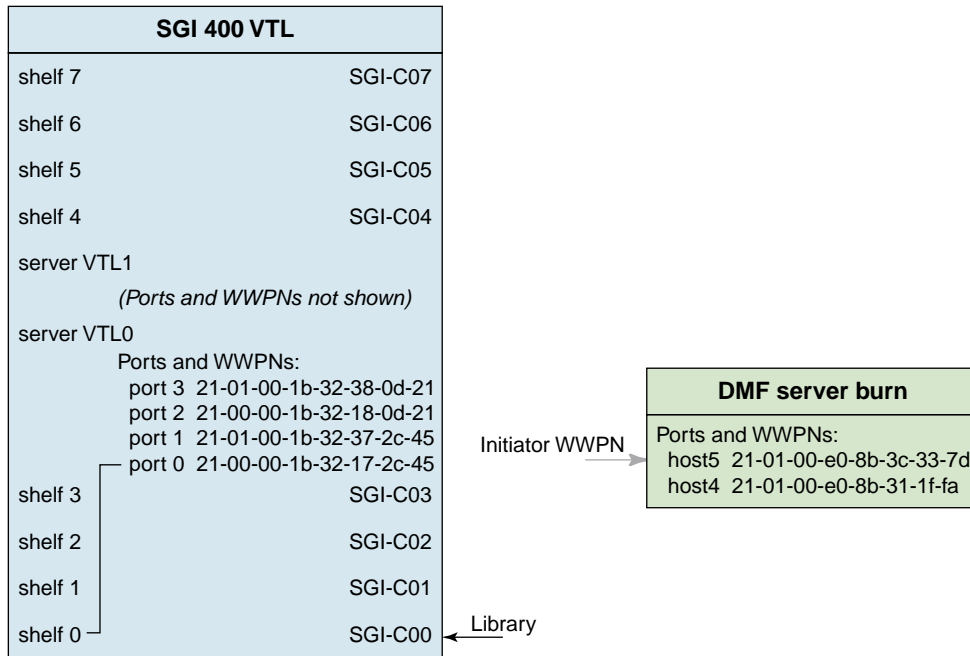


Figure 3-21 Relationship of the "Name" Field to the SGI 400 VTL

10. Select a target WWPN that will be used to connect the library's shelf and the DMF server, as shown in Figure 3-22. You can choose any of the target WWPNs, but you should use a different WWPN for each shelf. Figure 3-23 shows the relationship of the **Target WWPN** field to the SGI 400 VTL port. After you complete this step, the selected shelf will communicate with the DMF server through this port.

Note: Take note of the target WWPN you select here. You will use it later to identify the shelf when configuring OpenVault in step 2 of "Add an LCP for the Library" on page 75.

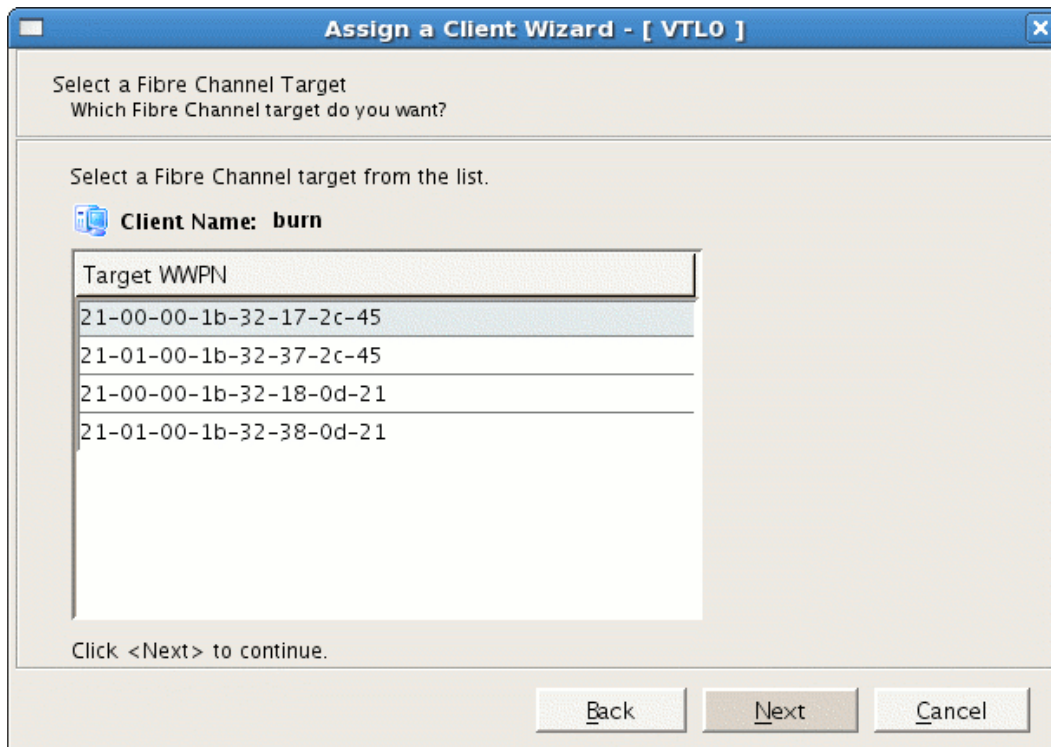


Figure 3-22 Select a Fibre Channel Target

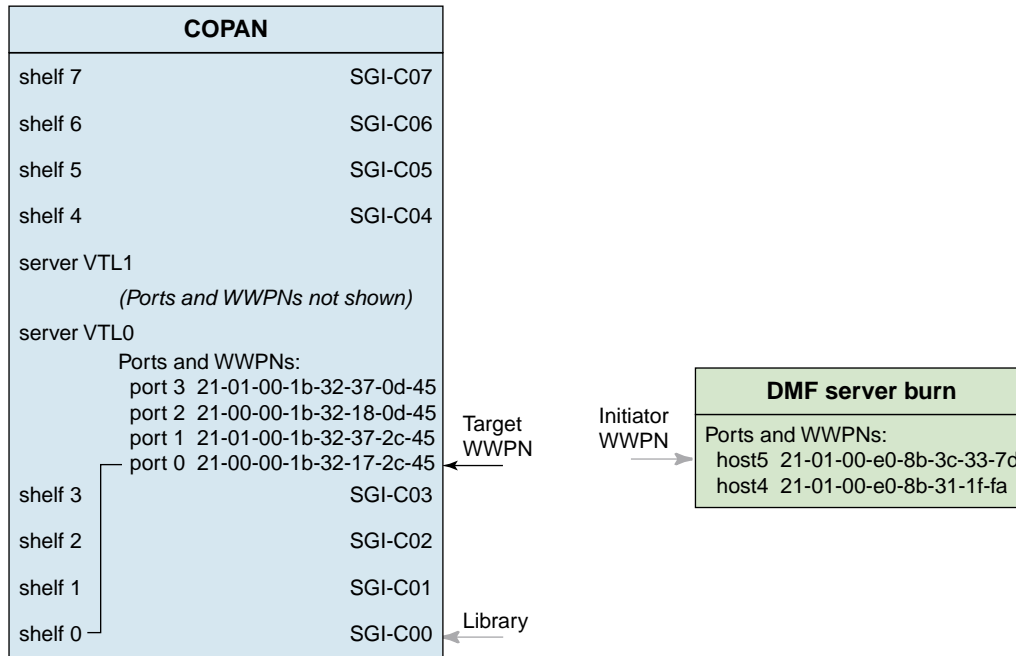


Figure 3-23 Relationship of the **Target WWPN** Field to the SGI 400 VTL Port

- Verify that the initiator WWPN from the DMF server corresponds to the target WWPN for the correct shelf. Click **Finish** to complete the configuration. Figure 3-24 shows an example. Figure 3-25 shows the relationships among the library, the DMF server initiator WWPN, and the SGI 400 VTL target WWPN.

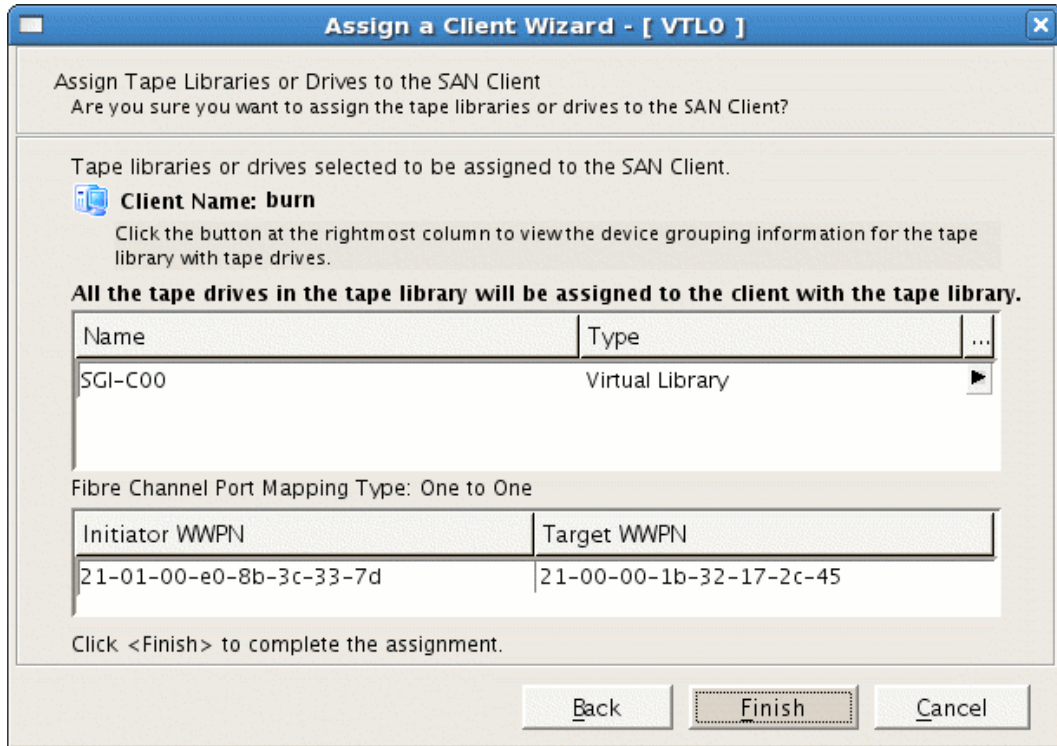


Figure 3-24 Verify that the Initiator and Target WWPNs Correspond

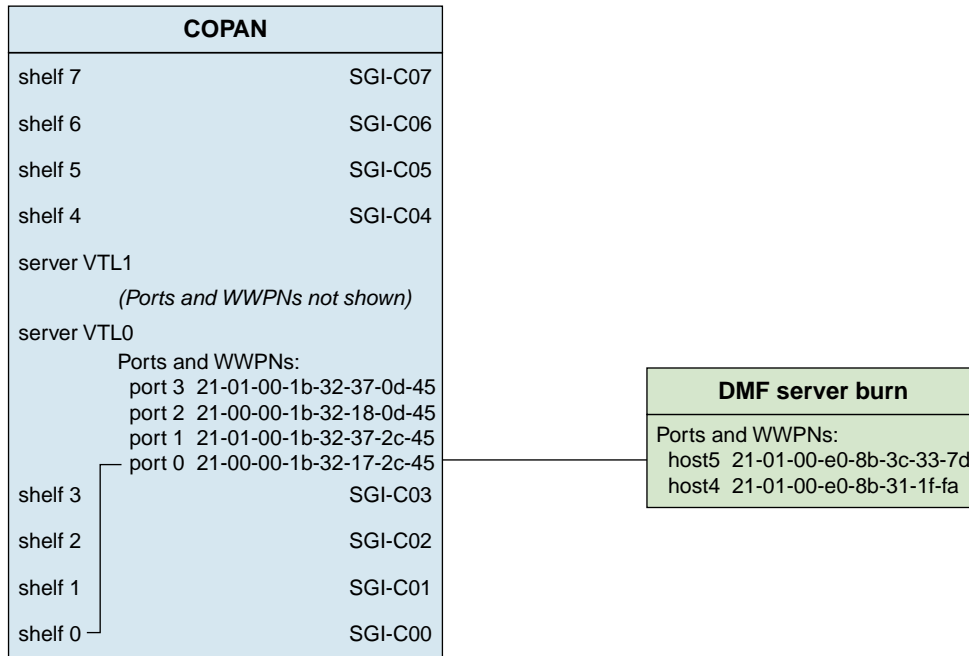


Figure 3-25 Relationships Among the Library, the DMF Server Initiator WWPN, and the VTL Target WWPN

12. Log in to the DMF server and force the DMF server to scan this port for the new connection:

```
dmfserver# echo "- - -" > /sys/class/scsi_host/port/scan
```

For example, for port host5 on DMF server burn:

```
burn# echo "- - -" > /sys/class/scsi_host/host5/scan
```

13. On the DMF server, use the `ts-lsscsi(8)` command to verify that the library is visible on the port:

```
dmfserver# ts-lsscsi
```

For example, the following output shows that the SGI 400 VTL cabinet and the eight LUNs are visible, and highlights the world wide node name (WWN) and WWPN information that is displayed:

```
burn# ts-lsscsi
[0:0:0:0] disk ATA HDS722580VLSA80 V320 /dev/sda
[1:0:0:0] disk ATA HDS722580VLSA80 V320 /dev/sdb
...
                    WWN                               WWPN
[5:0:0:0] mediumx SGI COPAN          4.50 -          |-----| |-----|
[5:0:0:1] tape IBM ULTRIUM-TD4 4C17 /dev/ts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun1
[5:0:0:2] tape IBM ULTRIUM-TD4 4C17 /dev/ts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun2
[5:0:0:3] tape IBM ULTRIUM-TD4 4C17 /dev/ts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun3
[5:0:0:4] tape IBM ULTRIUM-TD4 4C17 /dev/ts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun4
[5:0:0:5] tape IBM ULTRIUM-TD4 4C17 /dev/ts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun5
[5:0:0:6] tape IBM ULTRIUM-TD4 4C17 /dev/ts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun6
[5:0:0:7] tape IBM ULTRIUM-TD4 4C17 /dev/ts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun7
...

```

Configure OpenVault for the Library

To configure OpenVault for the library, do the following:

- "Add an LCP for the Library" on page 75
- "Add the Set of DCPs for the Library" on page 80
- "Import the Migration Tapes for the Library" on page 85
- "Test that OpenVault Can Mount a Migration Tape from the Library" on page 88

For more details, see *OpenVault Administrator Guide for SGI InfiniteStorage*.

Add an LCP for the Library

You must add one library control program (LCP) for each virtual library. Do the following:

1. Stop OpenVault by using the `ov_stop(8)` command:

```
dmfserver# ov_stop
```

2. Discover the device name that maps to the SGI 400 VTL shelf by examining the output from the following `ov_scandev(8)` command:

```
dmfserver# ov_scandev -p vendor,product,serial -L
```

The WWPN in the output is the WWPN of the shelf assigned in step 10 of "Export Library Access to the DMF Server" on page 63. For example, for the DMF server `burn`, the WWPN `2100001b32172c45` (highlighted here) is shelf 0. This is the device name that you must select later in step 3d. Also note the serial number as well for confirmation in step 3e:

```
burn# ov_scandev -p vendor,product,serial -L
/dev/tape/by-path/pci-0002:00:02.1-fc-0x2100001b32172c45:0x0000000000000000 SGI COPAN 8VM1008S00
      |                |                                     |         |
      -----                                     -----
                        WWPN                               SERIAL NUMBER
```

As you add libraries, they will be reflected in the `ov_scandev` output. For example, the following shows the output after you have added another library (the line ending in `8VM1008S01` highlighted in this example):

```
burn# ov_scandev -p vendor,product,serial -L
/dev/tape/by-path/pci-0002:00:02.1-fc-0x2101001b32372c45:0x0000000000000000 SGI COPAN 8VM1008S01
/dev/tape/by-path/pci-0002:00:02.1-fc-0x2100001b32172c45:0x0000000000000000 SGI COPAN 8VM1008S00
```

3: Create and Configure Each Virtual Library

3. Use the `ov_admin` command to add the LCP, using the defaults except as noted in the following example (comments below):

```
burn# ov_admin
```

```
### STEP a
```

```
What is the name of the OpenVault server? [burn]
Waiting for OpenVault to initialize ...
The OpenVault server was successfully started.
```

```
OpenVault Configuration Menu for server "burn"
```

```
Configuration on Machines Running LCPs and DCPs
  1 - Manage LCPs for locally attached Libraries
  2 - Manage DCPs for locally attached Drives
```

```
Configuration on Admin-Enabled Machines
  11 - Manage Cartridge Groups
  12 - Manage Drive Groups
  13 - Import Media
```

```
Configuration on the OpenVault Server Machine
  21 - Manage Applications
  22 - Manage OpenVault Client Machines
```

```
q - Exit.
```

```
Which operation would you like to do: 1
```

```
### STEP b
```

```
OpenVault LCP Configuration Menu
```

```
  1 - Create a new SCSI LCP
  2 - Create a new network LCP
```

```
  r - Return to Main Menu.
  q - Exit.
```

```
Which operation would you like to do: 1
```

```
### STEP c
```


3: Create and Configure Each Virtual Library

```
For the drive at location "4th drive from bottom",
    Enter a drive name for the element address "503": C00d4
For the drive at location "5th drive from bottom",
    Enter a drive name for the element address "504": C00d5
For the drive at location "6th drive from bottom",
    Enter a drive name for the element address "505": C00d6
For the drive at location "7th drive from bottom",
    Enter a drive name for the element address "506": C00d7
```

LCP Configuration Parameters:

```
Library name           : C00
LCP name               : C00@burn
LCP type               : STK-Lseries
OpenVault Server host name: burn
Library control path   : /dev/tape/by-path/pci-0002:00:02.1-fc-0x2100001b32172c45:0x0000000000000000
Number of drives       : 7
Default cartridge shape : LTO
```

Drives in the Library

Drive Name	Drive Address
-----	-----
C00d1	500
C00d2	501
C00d3	502
C00d4	503
C00d5	504
C00d6	505
C00d7	506

Create the LCP now? [Yes]

OpenVault -- LCP Creation Menu

Pathname	Library	LCP
- /dev/tape/by-path/pci-0002:00:02.1-fc-0x2100001b32172c45:0x0000000000000000	C00	C00@burn

r - Return to Previous Menu.
q - Exit.

For which Library would you like to add a LCP: **q**

STEP h

Comments:

- a. Log in to the DMF server and invoke the OpenVault administration tool:

```
dmfserver# ov_admin
```
- b. Enter 1 to manage LCPs.
- c. Enter 1 to create a new SCSI LCP.
- d. For the SCSI library selection, enter the number associated with the device name that corresponds to the library (determined in step 2).
- e. Confirm that the serial number (such as 8VM1008S00 indicating shelf 0 and library SGI-C00) matches the WWPN of the device (such as 2000001b32172c45) that was displayed in step 2 above.
- f. Enter the device name `c0n` where `n` matches the OpenVault library number, such as C00 for library SGI-C00.

Note: SGI recommends that you use a similar name but varying case to maintain a correspondence between the SGI 400 VTL view of the library (such as the C00 in SGI-C00) and OpenVault's view of the library (such as c00).

This is the value that is used for the `MOUNT_SERVICE_GROUP` parameter in the DMF configuration file. If you use a different naming scheme, you must modify the DMF configuration file. You must use a matching value later in step 3g of "Add the Set of DCPs for the Library" on page 80.

- g. Enter the drive names in turn, such as from C00d1 through C00d7 for library SGI-C00. In an SGI 400 VTL environment, these names correspond to the logical unit numbers (LUNs) for this library, lun1 through lun7.
- h. Enter **q** to exit.

Add the Set of DCPs for the Library

To add the drive control programs (DCPs) for the library, do the following:

1. Stop OpenVault by using the `ov_stop` command:

```
dmfserver# ov_stop
```

2. Review the device name that maps to the SGI 400 VTL shelf by examining the output from the following `ov_scandev` command:

```
dmfserver# ov_scandev -p vendor,product,serial -L
```

For example, for the DMF server `burn` shelf 0, the WWPN portion of the device name (such as `2100001b32172c45` highlighted in this example) is the device name that you must verify later in step 3d.

```
burn# ov_scandev -p vendor,product,serial -L
/dev/tape/by-path/pci-0002:00:02.1-fc-0x2100001b32172c45:0x0000000000000000 SGI COPAN 8VM1008S00
```

```
          |           |
          -----
          WWPN
```

3. Add the DCP, using the defaults except as noted in the following example (comments below):

```
burn# ov_admin
```

STEP a

```
What is the name of the OpenVault server? [burn]
Waiting for OpenVault to initialize ...
The OpenVault server was successfully started.
```

```
OpenVault Configuration Menu for server "burn"
```

```
Configuration on Machines Running LCPs and DCPs
  1 - Manage LCPs for locally attached Libraries
  2 - Manage DCPs for locally attached Drives
```

```
Configuration on Admin-Enabled Machines
 11 - Manage Cartridge Groups
 12 - Manage Drive Groups
 13 - Import Media
```

```
Configuration on the OpenVault Server Machine
```


- 21 - Manage Applications
- 22 - Manage OpenVault Client Machines

q - Exit.

Which operation would you like to do: **2**

STEP b

OpenVault DCP Configuration Menu

1 - Create a new DCP

r - Return to Main Menu.

q - Exit.

Which operation would you like to do: **1**

STEP c

STEP d

VERIFY THAT THE

WWPN MATCHES THE

ov_scandev OUTPUT

OpenVault -- DCP Creation Menu

WWPN

Pathname	WWPN	Drive
1 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun1		
2 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun2		
3 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun3		
4 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun4		
5 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun5		
6 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun6		
7 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun7		

r - Return to Previous Menu.

q - Exit.

For which Drive would you like to add a DCP: **1**

STEP e

lun1 CORRESPONDS TO DRIVE 1

Device information:

3: Create and Configure Each Virtual Library

Vendor: IBM
Product: ULTRIUM-TD4
Serial Number: 1285424968
DCP Type: Ultrium4
DCP Control Path: /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun1

Enter a name for the device, or <enter> to use the default name [tapel] **C00d1** **### STEP f**
MATCH THE DRIVE NAME

The following Drive Groups currently exist:

drives

What Drive Group do you want to use (or create)? [drives] **C00** **### STEP g**
ENTER A NAME THAT
ASSOCIATES TO THE
LIBRARY NAME

What is the unload delay time for this Drive Group? [0]

Select an Application which should be granted access to this Drive Group

- 1 - ov_umsh
- 2 - dmf

- d - Done (no more selections).

Which item would you like to select: **2** **### STEP h**

Application "dmf" now has access to Drive Group "C00".

Select another Application which should be granted access to this Drive Group

- 1 - ov_umsh

- d - Done (no more selections).

Which item would you like to select: **d** **### STEP i**

Drive C00d1 was successfully created.

What name would you like to give to this DCP? [C00d1@burn]

OpenVault -- DCP Creation Menu

```

      Pathname
- /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun1  C00d1  C00d1@burn
1 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun2
2 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun3
3 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun4
4 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun5
5 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun6
6 - /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun7

r - Return to Previous Menu.
q - Exit.

```

For which Drive would you like to add a DCP: 1

STEP j

Device information:

```

Vendor:          IBM
Product:         ULTRIUM-TD4
Serial Number:   1285424969
DCP Type:        Ultrium4
DCP Control Path: /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun2

```

Enter a name for the device, or <enter> to use the default name [tapel] C00d2

**### REPEAT THE STEPS
 ### FOR EACH DRIVE
 ### (SUCH AS C00d2-C00d7)**

...

Drive C00d7 was successfully created.

What name would you like to give to this DCP? [C00d7@burn]

OpenVault -- DCP Creation Menu

```

      Pathname
      Drive  DCP

```

3: Create and Configure Each Virtual Library

```
- /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun1 C00d1 C00d1@burn
- /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun2 C00d2 C00d2@burn
- /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun3 C00d3 C00d3@burn
- /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun4 C00d4 C00d4@burn
- /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun5 C00d5 C00d5@burn
- /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun6 C00d6 C00d6@burn
- /dev/pts/pci0002:00:02.1/fc/2000001b32172c45-2100001b32172c45/lun7 C00d7 C00d7@burn
```

r - Return to Previous Menu.
q - Exit.

For which Drive would you like to add a DCP: q

STEP k

Comments:

- a. Invoke the OpenVault administration tool:

```
dmfserver# ov_admin
```

- b. Enter 2 to manage DCPs.
- c. Enter 1 to create a new DCP.
- d. Verify that the WWPN in the device name displayed in the LUNs matches the library for which you want to configure the DCPs, displayed in step 2 above.
- e. Enter the menu selection number for the pathname that corresponds to the drive for which you want to add a DCP. This pathname will include the WWPN noted in step 2 above and will end with a LUN number that corresponds to the drive number, such as lun1 for drive 1.

Note: This WWPN/LUN name ordering convention is only true for SGI 400 VTL systems. For physical tape libraries, the WWPN/LUN names will be unique, but not necessarily ordered.

- f. Enter a name for the device that associates it with the library (rather than using the default name), such as C00d1 for library SGI-C00.
- g. Enter a name for the drive group that associates it with the library (rather than using the default name), such as C00.

Note: This is the value that is used for the `MOUNT_SERVICE_GROUP` parameter in the DMF configuration file. This value must match the value used earlier in step 3f of "Add an LCP for the Library" on page 75.

- h. Enter 2 to give access to the `dmf` application.
 - i. Enter `d` to complete the drive group definition.
 - j. Repeat steps 3e through 3i for each successive drive group until you have completed the last drive (such as `C00d7`).
 - k. Enter `q` to exit.
4. Restart OpenVault by using the `ov_start(8)`:

```
burn# ov_start
```

Import the Migration Tapes for the Library

Do the following to import the migration tapes for use by OpenVault and DMF:

1. View the available library and drives by using the following `ov_stat(8)` command:

```
dmfserver# ov_stat -ld
```

For example:

```
burn# ov_stat -ld
```

Library Name	Broken	Disabled	State	LCP State
C00	false	false	ready	ready

Drive Name	Group	Access	Broken	Disabled	SoftState	HardState	DCP State	Occupied	Cartridge PCL
C00d1	C00	true	false	false	ready	unloaded	ready	false	
C00d2	C00	true	false	false	ready	unloaded	ready	false	
C00d3	C00	true	false	false	ready	unloaded	ready	false	
C00d4	C00	true	false	false	ready	unloaded	ready	false	
C00d5	C00	true	false	false	ready	unloaded	ready	false	
C00d6	C00	true	false	false	ready	unloaded	ready	false	
C00d7	C00	true	false	false	ready	unloaded	ready	false	

3: Create and Configure Each Virtual Library

2. List the slot information to see all of the virtual tapes that can be imported.

```
dmfserver# ov_stat -L OV_library_name -s
```

The Cart ID field should be blank, indicating that no tapes (whose barcodes are listed in the PCL, *physical cartridge label*, column) have yet been imported. For example:

```
burn# ov_stat -L C00 -s
```

Library Name	Broken	Disabled	State	LCP State
C00	false	false	ready	ready

```
Library: 'C00'
```

Library	Slot Name	Slot Type	Occupied	PCL	Cart ID
C00	slot 1000	LTO	true	C00AAA	
C00	slot 1001	LTO	true	C00BAB	
C00	slot 1002	LTO	true	C00CAC	
C00	slot 1674	LTO	true	C00YT0	
...					
C00	slot 1675	LTO	true	C00ZT1	
C00	slot 1676	LTO	true	C00AT2	
C00	slot 1677	LTO	true	C00BT3	

Library	Bay Name	Slot Type	Total Slots	Free Slots
C00	bay 1	DLT	0	0

3. Import the virtual tapes for the library by tape type into OpenVault and add them to DMF according to the DMF volume group (VG) name (such as `vg0` for OpenVault library C00 and SGI 400 VTL library SGI-C00) by entering the following `dmov_loadtapes(8)` command:

```
dmfserver# dmov_loadtapes -t Ultrium4-800 -l OV_library_name -s tapesize_in_bytes VG_name
```

Note: The tape type `Ultrium4-800` is the name recognized by the `dmov_loadtapes` command for the **IBM ULTRIUM-TD4** tape type. You must enter the tape type on the command line or set the `OV_DEFAULT_CART_TYPE` environment variable.

This command will invoke an editor session. If there are any tapes that you do not want added to the library server database, such as any that should be reserved for use by an application other than DMF, you should delete those barcodes from the list. Otherwise, you should exit the editor.

For example, for library SGI-C00, you would enter C00 (the corresponding OpenVault library name) and vg0 (the corresponding VG name for library SGI-C00), and specify the 698 GB volume size in bytes (698000000000):

```
burn# dmov_loadtapes -t Ultrium4-800 -l C00 -s 698000000000 vg0
...
(editor session is invoked)
Unused tapes in library C00.
Delete any you do not wish to use for allocation or volume group vg0.
C00AAA
C00BAB
C00CAC
...
(exit the editor session)
...
(wait a few moments for the process to complete)
...
Loading tapes into allocation/volume group vg0.
Creating Ultrium4-800 cartridges in cartgroup carts for vg0.
```

The tapes for the library are now imported and ready for use. A tape that has been imported into OpenVault can be referred to as a *tape volume*.

4. Verify that the tapes now have OpenVault cartridge IDs assigned to them by examining the output from the following command:

```
dmfserver# ov_stat -L OV_library_name -s
```

For example, the following output shows that the tape with the barcode C00AAA now has an OpenVault cartridge ID of ooA69kzOtNYADeH5:

```
burn# ov_stat -L C00 -s
Library Name  Broken  Disabled  State    LCP State
C00           false   false     ready    ready
Library: 'C00'

Library  Slot Name  Slot Type  Occupied PCL      Cart ID
C00     slot 1000  LTO        true     C00AAA   ooA69kzOtNYADeH5
C00     slot 1001  LTO        true     C00BAB   ooA69kzOtNcAATXk
C00     slot 1002  LTO        true     C00CAC   ooA69kzOtNcAAwqN
...
C00     slot 1675  LTO        true     C00ZT1   ooA69kzOtSkAAQGF
C00     slot 1676  LTO        true     C00AT2   ooA69kzOtSkAAtZU
```

3: Create and Configure Each Virtual Library

```
C00      slot 1677  LTO      true      C00BT3  ooA69kzOtSkABKsZ
```

Library	Bay Name	Slot Type	Total Slots	Free Slots
C00	bay 1	DLT	0	0

Test that OpenVault Can Mount a Migration Tape from the Library

To test that OpenVault can mount a tape from the library for the `dmf` application, enter the following `ov_mount(8)` command:

```
dmfserver# ov_mount -A dmf -v VSN
```

For example, to test that OpenVault can mount the tape with the VSN (barcode) `C00BAB` and then exit:

```
burn# ov_mount -A dmf -v C00BAB
Mounted C00BAB on /var/opt/openvault/clients/handles//tEyedOqkcssi
burn:/sys/class/fc_host/host5 # exit
exit
```

Repeat the steps in this chapter for each shelf.

Configure Virtual Dump Tapes

In this chapter, you will reuse the space in SGI 400 VTL shelf 0 (library `SGI-C00`) reclaimed in "Free Space in Library SGI-C00 for Dump Tapes" on page 55 to create a set of dump tapes for the DMF system. There will be one set of dump tapes for the entire system, all located on shelf 0.

Note: All of the dump tapes must be on a single shelf. The sample DMF configuration file defines this to be shelf 0. If you want to use a different shelf, you must modify the DMF configuration file `drivegroup` object and the procedures in this guide accordingly.

This chapter discusses the following:

- "Create the Dump Tapes" on page 89
- "Relabel the Dump Tapes" on page 93
- "Import the Dump Tapes into OpenVault" on page 94
- "Test that OpenVault Can Mount a Dump Tape" on page 97

Create the Dump Tapes

Note: This can be a time-consuming process.

Do the following to reuse the space reclaimed in "Free Space in Library SGI-C00 for Dump Tapes" on page 55:

1. Start the VTL Console as directed in "Subsequently Connecting to the VTL Server" on page 34.

4: Configure Virtual Dump Tapes

2. Open the left-hand tree and place the cursor over the SGI-C00 library name:

VTLO
> **Virtual Tape Library System**
> **Virtual Tape Libraries**
> **SGI-C00**

Right-click and select **Properties**, as shown in Figure 4-1.

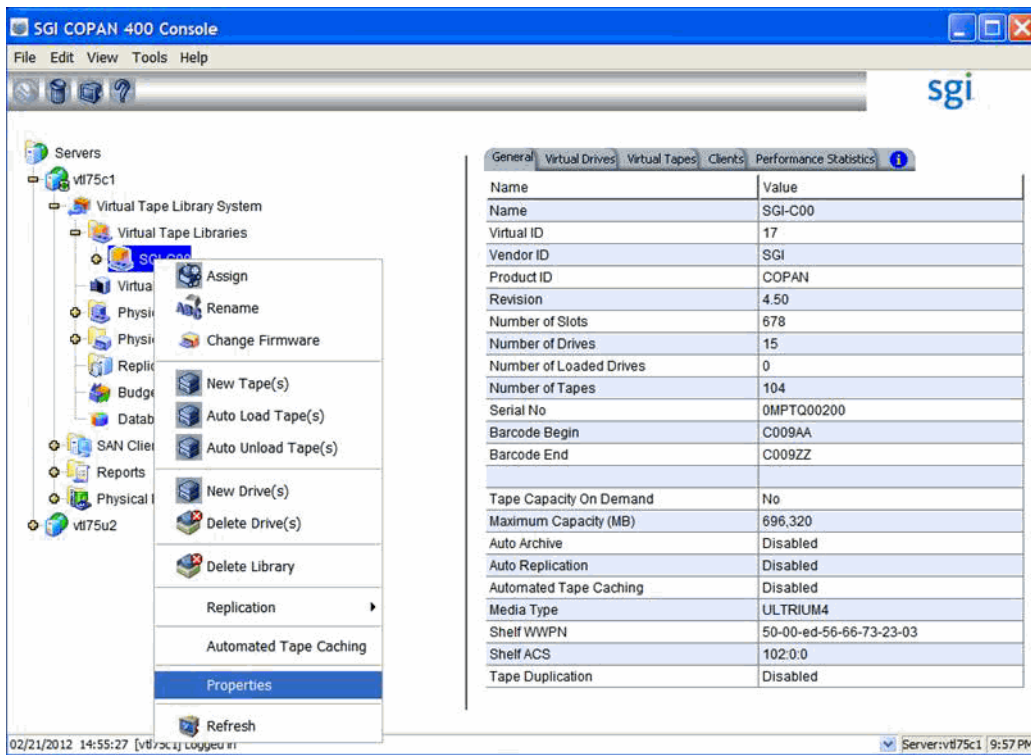


Figure 4-1 Change Tape Properties

Note: In the 7.5 version of the GUI, the menu is slightly different:



Figure 4-2 Menu to Change Tape Properties in VTL 7.5 GUI

4: Configure Virtual Dump Tapes

Then set the **Maximum Capacity** field in the **Tape Properties** tab to the *Dump_Tape_Size* value you determined in step 2 of "Free Space in Library SGI-C00 for Dump Tapes" on page 55 and click **OK**. For example, Figure 4-3 shows setting the size to be 116 GB for dump tapes.

Note: This example uses 116 GB as an illustration. You should use a value appropriate for your site.

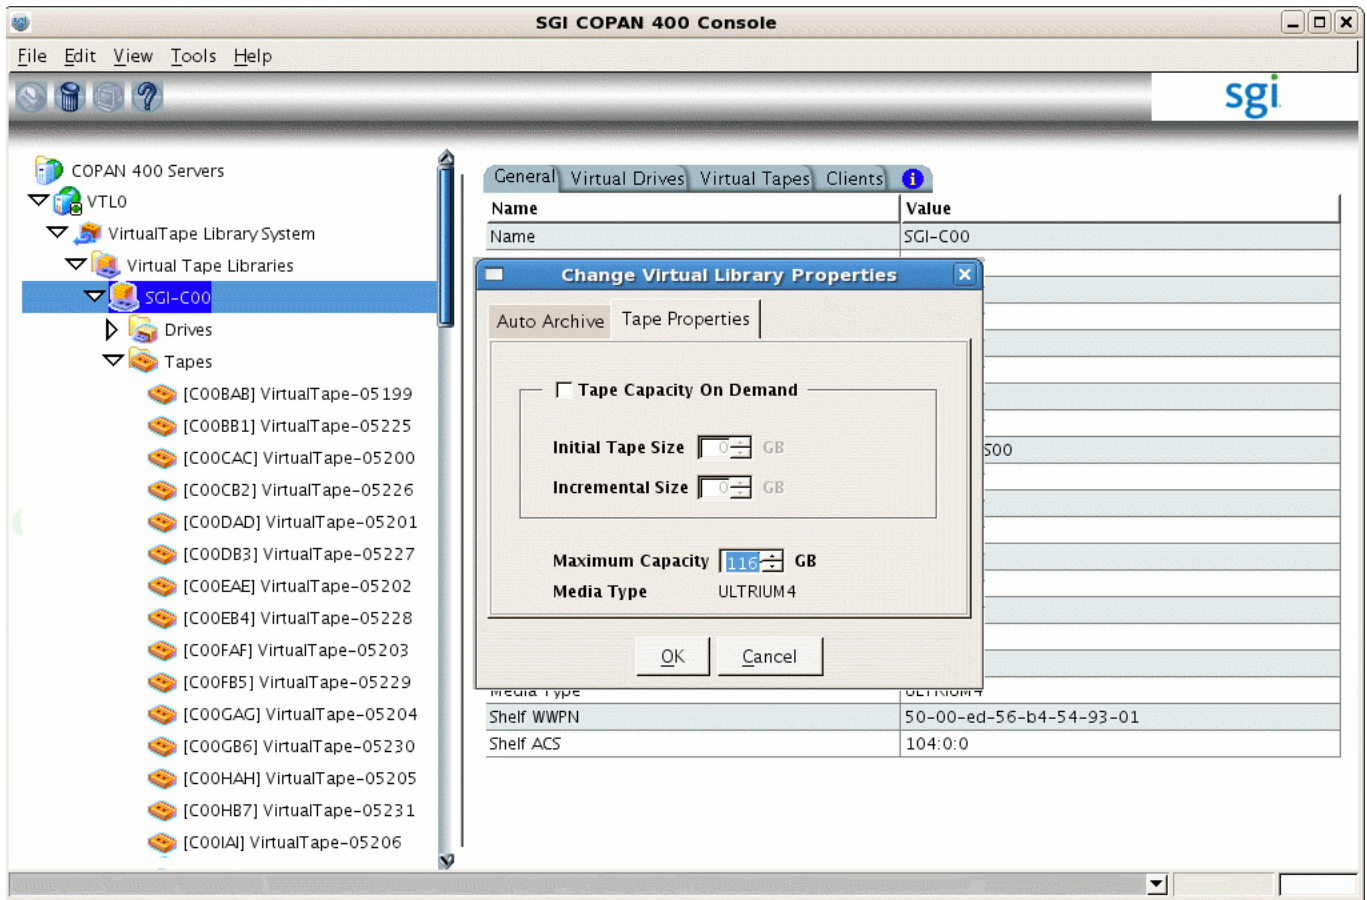


Figure 4-3 Set the Maximum Capacity for Dump Tape Size

3. Place the cursor over the library name, for example:

```
VTL0  
  > Virtual Tape Library System  
    > Virtual Tape Libraries  
      > SGI-C00
```

Right-click and select **New Tape(s)**.

4. Repeat the actions in steps 1 through 8 of "Configure Virtual Migration Tapes for the Library" on page 50 to create the required number of dump tapes.

Relabel the Dump Tapes

The DMF library server uses the barcode (or *volume serial number, VSN*) when selecting tapes. To make the virtual tapes created in the previous section ready for use as dump tapes, you must run the `dmreadycopantapes` script to relabel each virtual tape's barcode so that it indicates the specific RAID set in which the tape is stored, allowing DMF to optimize I/O to and from the shelf. (For example, `C009AA` might be transformed into `D00AAA`, where the first character `D` indicates that it is a dump tape, and the fourth character identifies the RAID set to DMF.)

Do the following:

1. Ensure that there is no other activity on the VTL Console or on the VTL application server that controls first library (such as `VTL0` for library `SGI-C00`). This is a requirement for the duration of this procedure.
2. Run the `dmreadycopantapes` script on the VTL application server from the `/usr/local/dmf` directory and provide the library portion of the barcode (the first three characters, such as `C00` in `C00AAA`).

The script will output an `executed successfully` line for each barcode that is changed. Depending upon the number of tapes and drives, this can be a time-consuming process.

For example, for dump tapes created on library SGI-C00, which have barcodes that begin with C00:

```
[root@VTLO vtl]# cd /usr/local/dmf
[root@VTLO dmf]# ./dmreadycopantapes -d -l C00

Command: settapeproperty executed successfully.
...
Command: settapeproperty executed successfully.
[root@VTLO dmf]#
```

For more information, add the `-v` option.



Caution: The script must complete successfully in order for DMF to function properly. If there are errors, fix them before continuing. If necessary, contact SGI Support.

Import the Dump Tapes into OpenVault

Do the following to import the dump tapes for use by OpenVault and DMF:

1. Stop and restart OpenVault in order to refresh the cartridge inventory:

```
dmfserver# ov_stop
dmfserver# ov_start
```

For example:

```
burn# ov_stop
OpenVault clients stopped
Successfully initiated shutdown on OpenVault server burn with the message:
Shutdown by root
burn# ov_start
OpenVault server started
OpenVault clients started
```

After OpenVault starts, wait a moment for other initialization to fully take place.

2. View the available library and drives by using the following `ov_stat(8)` command:

```
burn# ov_stat -ld
```

Library Name	Broken	Disabled	State	LCP State
C00	false	false	ready	ready

Drive Name	Group	Access	Broken	Disabled	SoftState	HardState	DCP State	Occupied	Cartridge PCL
C00d1	C00	true	false	false	ready	unloaded	ready	false	
C00d2	C00	true	false	false	ready	unloaded	ready	false	
C00d3	C00	true	false	false	ready	unloaded	ready	false	
C00d4	C00	true	false	false	ready	unloaded	ready	false	
C00d5	C00	true	false	false	ready	unloaded	ready	false	
C00d6	C00	true	false	false	ready	unloaded	ready	false	
C00d7	C00	true	false	false	ready	unloaded	ready	false	

3. List the slot information to see all of the dump tapes that can be imported:

```
dmfserver# ov_stat -s -L OV_library_name
```

For example, for OpenVault library C00 (SGI 400 VTL library SGI-C00):

```
burn# ov_stat -s -LC00
```

Library Name	Broken	Disabled	State	CP State
C00	false	false	ready	ready

Library: 'C00'

Library	Slot Name	Slot Type	Occupied	PCL	Cart ID
C00	slot 1000	LTO	true	D00ABW	
C00	slot 1001	LTO	true	C00BAB	ooA69k0rR6IACBC+
C00	slot 1002	LTO	true	C00CAC	ooA69k0rR6IACeS1
C00	slot 1003	LTO	true	D00DBX	
C00	slot 1004	LTO	true	C00EAE	ooA69k0rR6IAC71x
C00	slot 1005	LTO	true	C00FAF	ooA69k0rR6IADY40
C00	slot 1006	LTO	true	D00GBY	
C00	slot 1007	LTO	true	C00HAH	ooA69k0rR6MAACDa
C00	slot 1008	LTO	true	C00IAI	ooA69k0rR6MAAfWQ
C00	slot 1009	LTO	true	D00JBZ	

...

Library	Bay Name	Slot Type	Total Slots	Free Slots
C00	bay 1	DLT	0	0

The above output shows that the dump tape whose PCL column contains the VSN (barcode) D00ABW has not yet been imported into OpenVault, because it has no cartridge ID listed in the Cart ID column. (The VSNs for the migration tapes, such as C00BAB, do have cartridge IDs.)

4. Import the dump tapes for the library by tape type into OpenVault and add them to DMF according to the DMF configuration file `TASK_GROUPS` name by entering the following:

```
dmfserver# dmov_loadtapes -t Ultrium4-800 -l OV_library_name TASK_GROUPS_name
```

This command will invoke an editor session. If there are any tapes that you do not want added to the dump tapes list, such as any that should be reserved for use by an application other than DMF, you should delete those VSNs from the list. Otherwise, you should exit the editor.

For example, for SGI 400 VTL library SGI-C00, you would enter C00 (the corresponding OpenVault library name) and `dump_tasks` (the corresponding `TASK_GROUPS` name from the sample DMF configuration file for SGI 400 VTL use):

```
burn# dmov_loadtapes -t Ultrium4-800 -l C00 dump_tasks
...
(editor session is invoked)
Unused tapes in library C00.
Delete any you do not wish to use for task group dump_tasks.
D00ABW
D00DBX
D00GBY
...
(exit the editor session)
...
(wait a few moments for the process to complete)
...
Loading tapes into /dmf/home/dump_tasks/tapes for task group dump_tasks.
Creating Ultrium4-800 cartridges in cartgroup carts for dump_tasks.
```

The dump tapes for the library are now imported and ready for use.

5. Verify that the dump tapes now have OpenVault cartridge IDs assigned to them by examining the output from the following command:

```
dmfserver# ov_stat -s -L OV_library_name
```


For example, the following output shows that the dump tape with the VSN D00ABW in the PCL column now has an OpenVault cartridge ID of ooA69k0sSawABQWg:

```
burn# ov_stat -s -L C00
Library Name  Broken    Disabled  State    LCP State
C00           false    false     ready    ready
Library: 'C00'

Library  Slot Name  Slot Type  Occupied  PCL      Cart ID
C00     slot 1000  LTO        true      D00ABW   ooA69k0sSawABQWg
C00     slot 1001  LTO        true      C00BAB   ooA69k0rR6IACBC+
C00     slot 1002  LTO        true      C00CAC   ooA69k0rR6IACeS1
C00     slot 1003  LTO        true      C00BAB   ooA69k0rR6IACBC+
C00     slot 1002  LTO        true      D00DBX   ooA69k0sSawABtiq
C00     slot 1004  LTO        true      C00EAE   ooA69k0rR6IAC71x
C00     slot 1005  LTO        true      C00FAF   ooA69k0rR6IADY40
C00     slot 1006  LTO        true      D00GBY   ooA69k0sSawACK2A
C00     slot 1007  LTO        true      C00HAH   ooA69k0rR6MAACDa
C00     slot 1008  LTO        true      C00IAI   ooA69k0rR6MAAfWQ
C00     slot 1009  LTO        true      D00JBZ   ooA69k0sSawACoJB

...
Library  Bay Name  Slot Type  Total Slots  Free Slots
C00     bay 1     DLT        0             0
```

Test that OpenVault Can Mount a Dump Tape

To test that OpenVault can mount a dump tape from the library for the `dmf` application, enter the following:

```
dmfserver# ov_mount -A dmf -v VSN
```

For example, to test that OpenVault can mount the tape with the VSN (barcode) D00ABW and then exit:

```
burn# ov_mount -A dmf -v D00ABW
Mounted D00ABW on /var/opt/openvault/clients/handles//3yhZogXOqM4P
burn:/sys/class/fc_host/host5 # exit
exit
```


Test the DMF System

This chapter describes how to test the DMF system:

- "Validate the DMF Configuration" on page 99
- "Enable Audits of the DMF-Managed Filesystems" on page 100
- "Test the DMF Configuration" on page 104

Validate the DMF Configuration

To verify the DMF configuration after the tapes have been created and imported into OpenVault, repeat the validation check in DMF Manager (as in step 8 of "Create the DMF Configuration File" on page 23) by right-clicking anywhere within the **Overview** panel and selecting:

Configure
 > Validate Configuration

You could instead run the `dmcheck(8)` command on the DMF server, which provides similar functionality. There should not be any errors or warnings. For example:

```
burn# dmcheck

Checking DMF installation.
Linux burn 2.6.32.23-0.3-default #1 SMP 2010-10-07 14:57:45 +0200 ia64 ia64 ia64 GNU/Linux -
dmfserver
SuSE-release: SUSE Linux Enterprise Server 11 (ia64)
SuSE-release: VERSION = 11
SuSE-release: PATCHLEVEL = 1
sgi-issp-release: SGI InfiniteStorage Software Platform, version 2.2, Build
702rp41.sles11-1101052113
sgi-foundation-release: SGI Foundation Software 2.2, Build 702rp41.sles11-1101052113
lsb-release:
LSB_VERSION="core-2.0-noarch:core-3.2-noarch:core-4.0-noarch:core-2.0-ia64:core-3.2-ia64:core-4.0-ia64"
DMF version 5.4.0-0 rpm dmf-5.4.0-0 installed.

Checking DMF config file
Scanning for non-comment lines outside define/endif pairs
```

5: Test the DMF System

```
Scanning for DMF parameters without values
Checking all objects for invalid names
Checking base
Checking daemon
Checking policy space_policy
Checking policy copan_policy
Checking filesystem /dmi_fs
Checking Library Server vtl_ls
Checking Drive Group dg_c00
Checking Volume Group vg_c00
Checking Migrate Group mg0
Checking selection rules in policy space_policy.
Checking selection rules in policy copan_policy.
Checking Task Group daemon_tasks
Checking Task Group dump_tasks
Checking Task Group ls_tasks
Checking for unreferenced objects
Cross-checking LSs and task groups for duplicate VSNs

Checking other daemons.
  Checking OpenVault
  Checking chkconfig

No Errors found.
No warnings found.
```

If there are any errors or warnings, fix them and then rerun the validation. For more information, see the *DMF 6 Administrator Guide for SGI InfiniteStorage*.

Enable Audits of the DMF-Managed Filesystems

The DMF configuration file for SGI 400 VTL performs a nightly audit at 23:00 to check for database consistencies. To enable these audits of the DMF-managed filesystems, you must initialize the `dmaudit(8)` working environment. For more information about `dmaudit`, see *DMF 5 Filesystem Audit Guide for SGI InfiniteStorage*.

Do the following:

1. Start the DMF service:

```
dmfserver# service dmfs start
```

2. Run the `dmaudit(8)` command to verify the filesystems:

```
dmfserver# dmaudit
```

In particular, set up the file system scan list and run a simple snapshot.

3. Enter `edit` to edit a fresh copy of the filesystem scan list.
4. Delete the `/dcmmssp` and `/dmf/move` filesystems, leaving just your DMF-managed filesystem (such as `/myfs`).
5. Enter `accept`.
6. Enter `snapshot`.
7. Enter `quit`.

For example, on DMF server `burn` (the characters `###` in the right margin highlight comments related to the steps above):

```
burn# dmaudit
```

```
This program must create a working directory named 'working_dir' to hold
several very large work files.  The files may be needed in subsequent
executions of this program, so the directory should be placed in a file system
that is not cleared frequently.
```

```
Please enter the full path name of an existing directory in which subdirectory
'working_dir' can be created (<CR> to quit): /dmf/home
```

```
You must select the file systems to be scanned when this program searches for
migrated files.  By default the list contains all file systems currently
mounted.
```

```
IF THERE ARE ANY FILE SYSTEMS CONTAINING MIGRATED FILES THAT DO NOT APPEAR IN
THIS LIST, EXIT THIS PROGRAM AND MOUNT THEM BEFORE CONTINUING!!!
```

```
Failure to do so could mean the loss of the data in those files, because this
program sometimes removes database entries if it cannot find matching migrated
files for them.  File systems that you are sure do not contain migrated files
may be removed from this list to speed execution.  Remember that using restore
to load files dumped from a DMF-configured file system can result in migrated
files in other file systems.  If there is any question whether a file system
contains migrated files, leave it in the list.
```

5: Test the DMF System

Select:

<view> View the current file system scan list
<edit> Edit a fresh copy of the file system scan list
<accept> Accept the current file system scan list
<quit> Quit

Please enter your selection: **edit**

STEP 3

/dcmmssp
/dmf/move
/myfs

Delete /dcmmssp and /dmf/move

STEP 4

/dmi_fs

Select:

<view> View the current file system scan list
<edit> Edit a fresh copy of the file system scan list
<accept> Accept the current file system scan list
<quit> Quit

Please enter your selection: **accept**

STEP 5

MAIN MENU

Select:

<snapshot> Take a snapshot and report status of file systems and databases
<config> Examine or modify configuration information
<quit> Quit

Please enter your selection: **snapshot**

STEP 6

DAEMON DATABASE ERROR REPORT

No errors were discovered comparing the file systems against the daemon database.

DATA MIGRATION CONFIGURATION

```
Data migration home directory:      /dmf/home
Data migration binaries directory:  /usr/sbin

Server name:                        daemon
Server home directory:              /dmf/home/daemon
Server spool directory:             /dmf/spool/daemon

Data migration daemon process ID:   19275
```

No MSPs are configured

```
Name      Type
-----
vtl_ls    dmatls
vg_c00    volume group  in Library Server vtl_ls
vg_c01    volume group  in Library Server vtl_ls
vg_c02    volume group  in Library Server vtl_ls
vg_c03    volume group  in Library Server vtl_ls
vg_c04    volume group  in Library Server vtl_ls
vg_c05    volume group  in Library Server vtl_ls
vg_c06    volume group  in Library Server vtl_ls
vg_c07    volume group  in Library Server vtl_ls
mg0       migrate group Group Members: vg_c00
mg1       migrate group Group Members: vg_c01
```

LIST OF FILE SYSTEMS SCANNED

```
/myfs    xfs
```

MAIN MENU

```
Select:
  <inspect>  Inspect and correct file system and database errors
  <report>   Reprint status report for the current snapshot
  <verifymsp> Check the dmatls tape library server databases against the
```

```

                daemon databases
<snapshot>     Take a snapshot and report status of file systems and
                databases
<free>         Release all file space used by the current snapshot
<config>       Examine or modify configuration information
<quit>         Quit
```

Please enter your selection: **quit**

STEP 7

Test the DMF Configuration

To test the completed DMF configuration, do the following from the DMF server:

1. Change to a DMF-managed filesystem directory and create a test file of at least 1 MB. For example, to create the file `testfile` in the filesystem `/dmi_fs` on the DMF server burn:

```
burn# cd /dmi_fs
burn# /usr/lib/dmf/support/dmtwf -l 1m testfile
```

2. List the DMF status of the test file by using the `dmls(1)` command, which should show that the file is a regular (REG) file. For example, from the DMF server burn:

```
burn# dmls -l
total 980
-rw-r--r--  1 root      root      1000000 2010-11-05 12:36 (REG) testfile
```

3. Migrate the file and release its data blocks after the offline copies are made by using the `dmp(1)` command:

```
burn# dmp -r testfile
```

4. Idle the DMF daemon so that the `dmp` request will take place immediately by using the `dmdidle(8)` command:

```
burn# dmdidle
```

5. List the progressing status of the file by executing the `dmls` command multiple times:

```
burn# dmls -l
```


For example, the following repeated execution of `dmls` shows that `testfile` moves from migrating (MIG) to offline (OFF):

```
burn# dmls -l
total 980
-rw-r--r--  1 root      root      1000000 2010-11-05 12:36 (MIG) testfile
burn#
burn# dmls -l
total 980
-rw-r--r--  1 root      root      1000000 2010-11-05 12:36 (OFL) testfile
```

6. Delete the test file. For example:

```
burn# rm testfile
```


Using an Extremely Large Number of Tapes

This appendix discusses the following:

- "Number of Slots" on page 107
- "Creating Tapes and Relabeling Barcodes in Batches" on page 107

Number of Slots

There must be at least as many slots as there are virtual tapes. By default, the VTL Console provides 678 slots for virtual tapes. For most sites, this is a sufficient number. However, if you want a greater number of tapes, you must adjust the **Number of Slots** field in step 6 of "Create the Library" on page 40 accordingly. You must adjust this number the first time that you create tapes for the library, because you cannot modify the slot number later. The maximum number of slots available varies by the VTL software version:

- 4.5 VTL software has 4096 total number of slots:
 - If you have 1 shelf, it can have 4096 tapes
 - If you have 8 shelves, each shelf can have 512 ($4096/8=512$)
- 7.5 VTL software has 65536 total number of slots:
 - If you have 1 shelf, it can have 65536 tapes
 - If you have 8 shelves, each shelf can have 8192 ($65536/8 = 8192$)

Creating Tapes and Relabeling Barcodes in Batches

The algorithm used to relabel the barcodes so that they refer to the specific RAID set allows for up to 1,296 tapes per RAID set (which is 33,696 tapes per shelf, because there are 26 RAID sets per shelf). For most sites, this is a more than sufficient number. However, if you want an extremely large number of tapes, you must adjust the number of slots as mentioned above and create the tapes in multiple batches, running the `dmreadycopantapes` script ("Relabel the Migration Tapes" on page 54) on the first batch before creating the second batch.

Note: You must adjust the slot number when creating the first batch of tapes; after the first batch has been created, you can no longer adjust the slot number.

Enabling Compression

If your site's data would benefit from software compression, you can enable SGI 400 VTL compression mode (and therefore DMF compression) for each application server (VTL0 and VTL1).

Note: If you are using the 4.5 VTL software, compression carries an overhead CPU penalty and is not appropriate in some cases.

This is not an issue with the 7.5 VTL software because it uses a different compression algorithm.

Do the following:

1. Open the left-hand tree and place the cursor over **Virtual Tape Library System**. For example, for VTL0:

VTL0

> **Virtual Tape Library System**

Right-click and select **Properties**.

2. Check the **Enable Virtual Tape Library compression mode** box, select **Software**, use the default threshold value, and click **OK**. Figure B-1 shows an example.

Note: The 7.5 VTL GUI also displays information about **Migration Threshold** and **Reclamation Threshold**; these values do not apply to DMF and can be ignored.

B: Enabling Compression

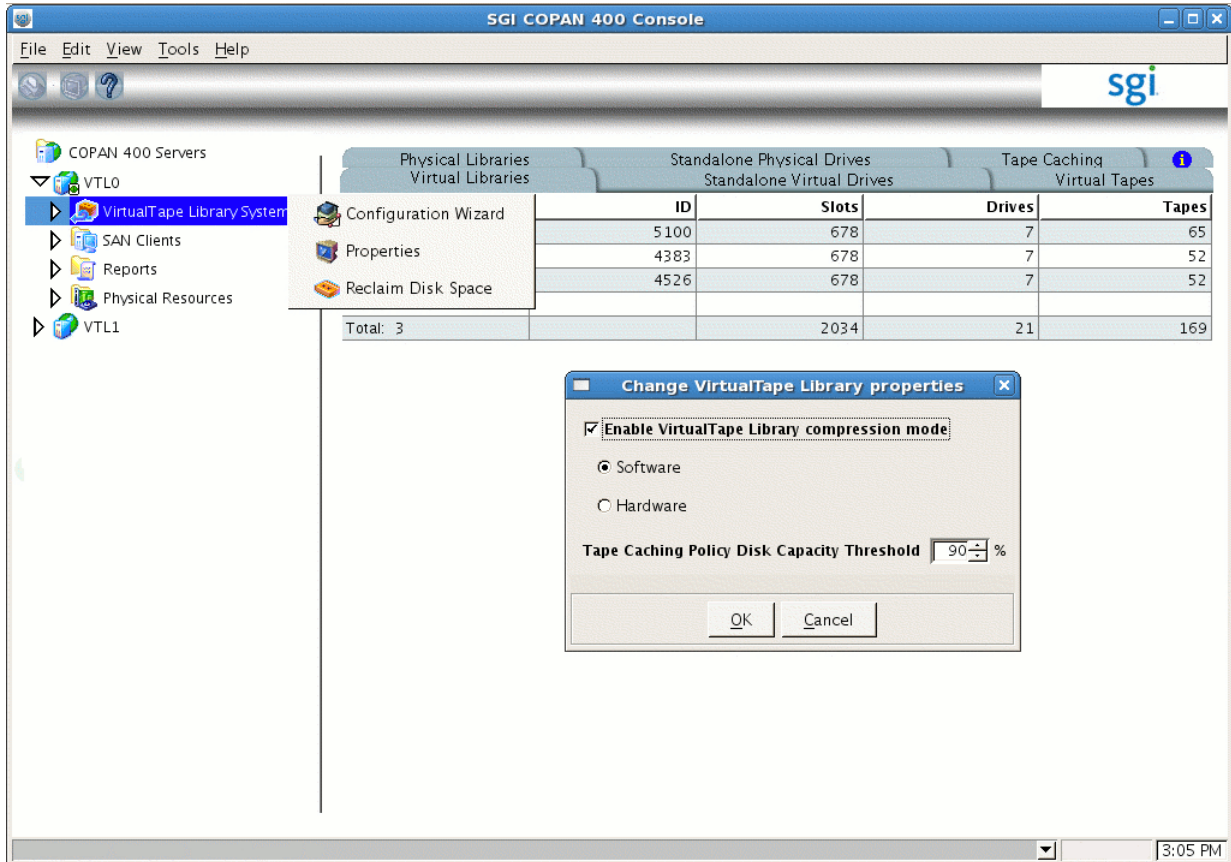


Figure B-1 Enabling Compression

Reclaiming Space for Dumps

If you do not have enough space designated in dump tapes for your backup requirements, you can do the following to reclaim space from an empty migration tape:

1. Verify that the migration tape it is empty by examining the output from the `dmvoladm(8)` command on the DMF server, specifying the volume serial number (VSN, which is the same as the barcode) for the tape:

```
dmfserver# dmvoladm -c "list vsn=VSN"
```

The `DATA LEFT` field should show that there is no data on the tape (0.000000), as in the following example for the migration tape with VSN C00SAW and DMF server burn:

```
burn# dmvoladm -c "list vsn=C00SAW"
-----
VSN          VOLGRP LB      DATA LEFT      DATA   EOT   EOT          WR/FR
WRITTEN  CHUNK  ZONE  HFLAGS          AGE
-----
C00SAW      vg0  a1      0.000000      0.000000      1     1  -----      12m
Database was not modified.
```

2. Lock the tape by setting the operator hold flag to on, which means that it cannot be used for migrating data:

```
dmfserver# dmvoladm -c "update VSN hoa on"
```

For example:

```
burn# dmvoladm -c "update C00SAW hoa on"
Updated 1 record.
```

3. Verify that the tape is not already mounted by verifying that it does not show up in OpenVault status output by using the `ov_stat(8)` command:

```
dmfserver# ov_stat -ld | grep VSN
```

For example:

```
burn# ov_stat -ld | grep C00SAW
burn#
```

If the VSN does appear in the `grep` output, wait for the tape to unmount before continuing on to the next step. It can take up to 10 minutes for the unmount to complete.

4. Reconfirm that the migration tape it is still empty and that it is now locked:

```
dmfserver# dmvoladm -c "list vsn=VSN"
```

The DATA LEFT field should still show that there is no data on the tape (0.000000), and the o hold flag should appear in the HFLAGS field, as in the following example:

```
burn# dmvoladm -c "list vsn=C00SAW"
          DATA    EOT    EOT                WR/FR
VSN      VOLGRP LB   DATA LEFT    WRITTEN  CHUNK  ZONE  HFLAGS    AGE
-----
C00SAW   vg0  a1     0.000000    0.000000    1     1  --o-----  13m
Database was not modified.
```

If the output now shows that there is data left on this tape, see the information about managing tapes and setting the sparse flag in the *DMF 6 Administrator Guide for SGI InfiniteStorage*.

5. Delete the tape from the DMF volume database:

```
dmfserver# dmvoladm -c "delete VSN"
```

For example:

```
burn# dmvoladm -c "delete C00SAW"
Deleted 1 record.
```



Caution: If you do not see the above output, do not continue with this tape. Reset the hold flags to the values the file had in step 1 (for more information, see *DMF 6 Administrator Guide for SGI InfiniteStorage*). Try a different tape or contact SGI Support.

6. Remove the volume from the OpenVault system by using the `ov_vol(8)` command:

```
dmfserver# ov_vol -D -a dmf -v VSN
```


For example:

```
burn# ov_vol -D -a dmf -v C00SAW
Volume deleted:
  volume name = 'C00SAW', application name = 'dmf'
  cartridgeID = 'ooA69k0uDdAABqXI', side = 'SideA', partition = 'PART 1'
```

Take note of the `cartridgeID` field (in this case, `ooA69k0uDdAABqXI`), which will be required for the `cartID` variable in the next step.

7. Remove the cartridge ID entry from the OpenVault catalog by using the `ov_purge(8)` command:

```
dmserver# ov_purge -f -C cartID
```

For example, for cartridge ID `ooA69k0uDdAABqXI`:

```
burn# ov_purge -f -C ooA69k0uDdAABqXI
Deleted partition PART 1
Deleted cartridge ooA69k0uDdAABqXI
```

8. Verify that the tape's VSN is no longer in the OpenVault catalog by examining the output of the `ov_dumptable(8)` command, which should no longer display the cartridge ID for the VSN:

```
dmfserver# ov_dumptable -n -d'|' -c LibraryName,CartridgePCL,CartridgeID SLOT | grep VSN
```

For example, the cartridge ID `ooA69k0uDdAABqXI` does not appear in the following output:

```
burn# ov_dumptable -n -d'|' -c LibraryName,CartridgePCL,CartridgeID SLOT | grep C00SAW
c00|C00SAW|
```

9. Use the VTL Console to delete the migration tape. Open the left-hand tree and place the cursor over the library name. For example, for `SGI-C00`:

```
VTL0
  > Virtual Tape Library System
    > Virtual Tape Libraries
      > SGI-C00
```

Right-click and select **Delete Tape(s)**.

10. Reuse the freed space for dump tapes. See "Create the Dump Tapes" on page 89.

Index

A

- ADMIN_EMAIL, 23
- application server, 2
- archiving concept, 6
- archiving overview, 4
- audits, 100

B

- backup requirements, 15
- backups, 10
- barcode naming convention, 57
- barcode relabeling, 107
- base object, 23

C

- C00 naming convention, 10
- cabinet, 1
- CACHE_MEMBERS, 24
- Celerra, 66
- compression, 109
- configuration summary, 10
- configure virtual dump tapes, 89
- copan directory, 13
- COPAN VTL LS sample configuration in DMF Manager, 25
- copan_ls, 25
- COPAN_VSNS, 23

D

- DATA LEFT, 111

- DATABASE_COPIES, 15
- database_copies, 14
- DCP, 10, 80
- delete tapes, 58
- df, 15
- dmaudit, 100
- dmcheck, 99
- dmdidle, 104
- dmf application in OpenVault, 13
- DMF auditing, 100
- DMF configuration, 21
- DMF configuration file, 23
- DMF license, 22
- DMF Manager, 24
- DMF overview, 4
- /dmf/cache, 14
- /dmf/dumps, 15
- /dmf/home, 14
- /dmf/home/dump_tasks/tapes, 22
- /dmf/journals, 14
- /dmf/journals/database_copies, 14
- /dmf/spool, 14
- /dmf/tmp, 14
- dmls, 104
- dmov_loadtapes, 96
- dmpu, 104
- dmreadycopantapes, 54, 93, 108
- dmtwf, 104
- dmvoladm, 111
- drive control program
 - See "DCP", 10
- drive groups, 26, 26
- DRIVE_GROUPS, 23
- drivegroup object, 23
- dual-state, 7
- dump tape size, 57
- dump tapes, 55, 89

DUMP_FILE_SYSTEMS, 23

E

/etc/lk/keys.dat, 22
EveryoneFC, 63
export library access, 63

F

fast-mount cache
 definition, 7
fastmountcache object, 24
FILE_RETENTION_DAYS, 24
FMC_MOVEFS, 24
FMC_NAME, 24
FREE_VOLUME_MINIMUM, 24
FREE_VOLUME_TARGET, 24
freeing and recalling file data, 7
freeing space, 4
freeing space for dump tapes, 55

G

GROUP_MEMBERS, 23

H

host bus adapter (HBA), 63

I

identifiers, 10
idle DMF, 104
import dump tapes, 94
import migration tapes, 85
initiator WWPN, 63

inodes remain on primary filesystem, 4
integrated data mover functionality, 3
INTERFACE, 24, 29
introduction, 1

L

LCP, 10, 75
library configuration, 31
library control program, 10
library server, 25
libraryserver object, 23
list migrated files, 104

M

MAX_PUT_CHILDREN, 23
merging, 14
migrate a file, 104
migrategroup object, 23
migrating data, 4
migrating file data, 7
migration tape size, 36
migration tapes, 50
mount a dump tape, 97
MOUNT_SERVICE_GROUP, 23
mounting service, 1
mounting tapes, 88
MULTIPLIER, 23

N

naming convention for barcodes (VSNs), 57
naming conventions, 10
network configuration, 12

O

offline, 7
 OpenVault, 1
 OpenVault configuration, 10
 ov_admin, 10, 76, 80
 OV_DEFAULT_CART_TYPE, 86
 ov_dumptable, 113
 ov_mount, 88, 97
 ov_procs, 12
 ov_purge, 113
 ov_scandev, 75, 80
 ov_start, 94
 ov_stat, 85, 95
 ov_stop, 75, 80, 94
 ov_vol, 112

P

parallel data mover node, 3
 Parallel Data Mover Option, 3
 permanent storage, 4
 preliminary requirements, 11

R

RAID technology, 36
 recalling data, 4
 reclaiming space for dumps, 111
 relabel dump tapes, 93
 RESERVED_VOLUMES, 23

S

SERVER_NAME, 23, 29
 SGI-C00 naming convention, 10
 shelf order, 1
 snapshot, 101
 sparse volumes, 14

system overview, 1

T

tapes and slots, 107
 target WWPN, 63
 taskgroup object, 23
 test DMF, 99, 104
 test OpenVault, 88, 97
 theory of archiving, 4

U

/usr/lib/dmf/support/copan, 13
 /usr/lib/dmf/support/dmtwf, 104

V

validate DMF configuration, 99
 /var/lib/dmf/database_copies, 14
 volume serial number, 10
 See "VSN", 10
 volume set addressing, 66
 volumegroup object, 23
 VSN, 10
 VSN naming convention, 57
 VTL application server, 2
 VTL Console, 31
 VTL server, 31
 VTL0, 13
 VTL0 and VTL1, 2

W

WWN, 74
 WWPN, 63

Z

ZONE_SIZE, 23

zoning, 12