EMC® VNX™ Series

Release 7.1

Configuring NDMP Backups on VNX™

P/N 300-013-432 Rev 02

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748-9103
1-508-435-1000
www.EMC.com

Copyright © 1998 - 2012 EMC Corporation. All rights reserved.

Published December 2012

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date regulatory document for your product line, go to the Technical Documentation and Advisories section on EMC Powerlink.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Corporate Headquarters: Hopkinton, MA 01748-9103

Contents

Preface 9

Chapter 1: Introduction 11

System requirements 12
Restrictions 12
Cautions 15
User interface choices 15
Related Information 16

Chapter 2: Concept 19

Basic NDMP backup configuration 21
Single TLU configuration 22
Dual TLU configuration 22
Three-way backup configuration 23
Backup levels with VNX and NDMP 24
NDMP V4 IPv6 Extension for VNX 24
Restore PAX file systems 26
Dynamic drive sharing 26
VTLUs 27
VNX DDAR support feature 27
FileMover file system backup with NDMP 27
File system backup with NDMP using an integrated SnapSure checkpoint 28
Backup/restore with NVB 29
NVB requirements 31
NVB restrictions 32
Filter to exclude directories or files from backup 32

Valid path rules to exclude directories from backup....33
Valid file-pattern rules to exclude files from backup....34
Tape silvering....34
Planning considerations....34
Data Mover-to-TLU cabling examples....35
Implement international character support with NDMP....36
NDMPcopy....37
NDMP last access time....42
Multithreaded backup and restore design....42

Chapter 3: Configuring....45
Configure Data Movers....46
Configure the tape buffer size....46
Configure the read or write block size....47
Configure an NS series Data Mover....47
Configure Data Movers for international character sets....48
Override the autonegotiated NDMP version....49
Configure the Data Mover to increase the NDMP sessions (Optional)....50
Configure the NDMP port range....51
Set up the TLU....53
Configure stand-alone tape drives....56
Configure VTLUs....56
Assign a username and a password to a Data Mover....56
With NIS....57
Without NIS....57
Set environment variables....58

Chapter 4: Managing....63
Manage backup parameters and statistics....64
View backup parameters....64
View a specific parameter....66
View all parameters....67
Modify parameters....69
View backup statistics....70
Reset backup statistics....73
Back up a VDM file system with NDMP....74
Back up a FileMover file system with NDMP....74
Restore a FileMover file system....75

Back up a file system with NDMP by using an integrated SnapSure checkpoint....75
Perform backup and restore with NVB....75
Specify NVB in backup vendor software....76
Perform NVB restores....77
Perform full destructive restore....77
Perform a full destructive restore of a VDM file system....85
File-level restore....85
Exclude directories or files from backup by using filtering....86
Exclude a directory....86
Exclude a file....88
Perform tape silvering....88
Restore with NetBackup....89
Restore with CommVault Galaxy....93
Restore with NetWorker....95

Chapter 5: Troubleshooting....101
EMC E-Lab Interoperability Navigator....102
VNX user customized documentation....102
Troubleshooting checklist....102
Reporting a problem....103
Known problems and limitations....104
Error messages....104
EMC Training and Professional Services....105

Appendix A: Configure your Backup Software....107
Terminology differences....109
Configure Atempo Time Navigator....109
Create the file server....110
Create and configure a library....110
Configure the drives....111
Create a media pool....111
Create the Time Navigator client....112
Configure BakBone NetVault....113
Add an NDMP server....113
Use the automatic device configuration....114
Configure CA BrightStor ARCserve Backup....114
Configure the NDMP devices....115
Configure the nas.cfg file....116

Add a Data Mover to ARCserve Backup....117
Configure CommVault Galaxy....117
Add a client....118
Configure the TLU and tape drives for NDMP operations....118
Test the TLU device....120
Configure EMC Data Manager with NDMP....120
Local backups with EDM....121
Three-way backups with EDM....121
EDM restrictions....122
Install and configure the EDM client....123
EMC NetWorker....124
Configure EMC NetWorker....124
Sample jbconfig script....125
Tape drive buffer sizes....129
Fujitsu-Siemens NetWorker....131
Configure HP OpenView Storage Data Protector....132
Import the NDMP server host....132
Create a media pool....133
Configure the tape library unit....133
IBM Tivoli Storage Manager....135
Configure IBM Tivoli Storage Manager....135
Configure Syncsort Backup Express....141
Configure the Data Mover....141
Configure the libraries attached to the Data Mover....141
Run an NDMP backup of a Data Mover....142
Configure Symantec NetBackup....143
Enable host access to a Data Mover....143
Configure a robot....148
Add a robot....150
Update the volume database....152
Add a drive....152
Add a storage unit....153
Create a backup policy....153
Specify a backup type....154
Configure Symantec Backup Exec....154
Add the NDMP server to the Backup Exec....155
Inventory the newly added TLU....156
Create a backup job....158
Create a restore job....160

Appendix B: Configure CNS and NS for NDMP.....163

Glossary.....165

Index.....169

Contents

Preface

As part of an effort to improve and enhance the performance and capabilities of its product lines, EMC periodically releases revisions of its hardware and software. Therefore, some functions described in this document may not be supported by all versions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.

If a product does not function properly or does not function as described in this document, please contact your EMC representative.

Special notice conventions

EMC uses the following conventions for special notices:

Note: Emphasizes content that is of exceptional importance or interest but does not relate to personal injury or business/data loss.

NOTICE Identifies content that warns of potential business or data loss.

CAUTION Indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.

WARNING Indicates a hazardous situation which, if not avoided, could result in death or serious injury.

DANGER Indicates a hazardous situation which, if not avoided, will result in death or serious injury.

Where to get help

EMC support, product, and licensing information can be obtained as follows:

Product information—For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to EMC Online Support (registration required) at http://Support.EMC.com.

Troubleshooting—Go to EMC Online Support at http://Support.EMC.com. After logging in, locate the applicable Support by Product page.

Technical support—For technical support and service requests, go to EMC Customer Service on EMC Online Support at http://Support.EMC.com. After logging in, locate the applicable Support by Product page, and choose either **Live Chat** or **Create a service request**. To open a service request through EMC Online Support, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Note: Do not request a specific support representative unless one has already been assigned to your particular system problem.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications.

Please send your opinion of this document to:

techpubcomments@EMC.com

1 Introduction

The Network Data Management Protocol (NDMP) allows you to control the backup and recovery of an NDMP server through a network backup application, without installing third-party software on the server. In an EMC VNX, the Data Mover functions as the NDMP server.

This document explains how to perform NDMP backups with VNX and your backup software. It also contains the information you need to cable and configure VNX. Use this and other VNX documentation in conjunction with the specific backup software documentation provided by your vendor. Chapter 2 provides more details.

This document is part of VNX documentation set and is intended for the backup system administrator implementing a backup strategy with one or more VNX. The system administrator should be familiar with backup and restore operations and with the basic administration of VNX.

Topics included are:

- System requirements on page 12
- Restrictions on page 12
- Cautions on page 15
- User interface choices on page 15
- Related Information on page 16

System requirements

Table 1 on page 12 describes the EMC® VNX™ software, hardware, network, and storage configurations.

Note: VNX supports NDMP version 4.0.

Table 1. System requirements

Software	VNX version 7.1 NDMP-compliant backup software running on the NDMP client. *VNX for File Release Notes* contain a current list of the EMC qualified NDMP backup software vendors.
Hardware	VNX
Network	Dedicated network (optional)
Storage	Compatible tape device or tape library unit (TLU)

Restrictions

Note these restrictions when configuring VNX.

General restrictions

- EMC supports NDMP drive sharing only through a storage area network (SAN).
- VNX uses port 10000 for NDMP operations. NDMP three-way backups and restores use dynamic port allocation from the Data Mover.
- If you plan to turn on internationalization, do so before, not after creating the file systems.
- Stripe size on the system can affect backup performance. For optimal backup stripe size settings on EMC Symmetrix® and EMC VNX for block systems, consult the EMC E-Lab™ Interoperability Navigator.
- VNX supports a maximum of eight concurrent NDMP sessions per NDMP-Host Data Mover. The backup client (DMA) should be compatible and configured to support up to eight concurrent NDMP sessions.

- The Data Mover must have a minimum of 8 GB memory to run greater than four and up to eight NDMP sessions.
- For backups with VNX, a media interface adapter (MIA) for device connection to the AUX-0 port is required. For backups with VNX series, a small form-factor pluggable (SFP) adapter is required. Instructions on installing an MIA in VNX, or an SFP in VNX series, are documented in individual setup instructions provided with each system.
- Make sure that no backup or restore jobs are running on the Data Mover before you issue the nas_fsck –aclchk command.
- If a Data Mover is in ASCII mode and you back up a file system with asciifilter off, the filename with international characters cannot be recovered if the asciifilter is turned on during the restore. It is recommended that all Data Movers be configured to use UNICODE rather than ASCII mode.

  Configuring and Managing CIFS on VNX provides information on how to check and change the Data Mover mode. *Using International Character Sets on VNX for File* provides information on enabling Unicode on a new VNX.

Restrictions for TLU connections

- Do not connect the Control Station to the TLU. Only Data Movers can be connected to the TLU.
- Connecting multiple Data Movers to the same TLU is supported.
- If an NDMP-Host Data Mover fails over to its standby, you must physically connect the Data Mover's TLU cable to the standby, and restart the backup. You could also need to reconfigure the backup software with the appropriate address.

NDMP restrictions

- NDMP does not support the Archive Bit method. Only the Date/Time stamp method is supported.
- The special wildcard characters, as well as the include/exclude options, are not supported. To back up multiple directories, include all directories in the file/directory list when creating the backup job. You can use filtering to exclude specific files and directories from a backup, as described in Exclude directories or files from backup by using filtering on page 86.
- NDMP does not follow symbolic links among file systems. Therefore, an NDMP backup does not include a symbolic link's target file system. *Managing a Multiprotocol Environment on VNX* contains more information on symbolic links.

- NDMP backups and restores of iSCSI LUNs are not supported.
- Do not convert a raw file system (rawfs) that is created for NVB full destructive restore to uxfs file system, until the restore operation is completed successfully.

Cautions

If any of this information is unclear, contact your EMC Customer Support Representative for assistance:

- If a failover occurs during a backup, you must directly connect the tape drive to the new Data Mover and restart the backup. You could also need to reconfigure the backup software.
- Do not unmount a target file system during a backup or restore operation.
- If you change the network zoning configuration, where relevant, propagate changes to the consistent binding table and adjust the backup software to reflect the new zoning configuration.
- In VNX for block environment, before you run the server_devconfig -create command, verify that all paths to the Data Mover are active and no LUNs are trespassed.
- You cannot back up tree quotas from the original file system and restore them to a different file system because the action is not supported. To preserve tree quotas when restoring, the target file system must be the same destination as the original file system, and it must have tree quotas enabled on it.

User interface choices

VNX offers flexibility in managing networked storage that is based on your support environment and interface preferences. This document describes how to configure NDMP backups by using the command line interface (CLI). You can also perform some of these tasks by using one of the VNX management applications:

- EMC Unisphere™
- Microsoft Management Console (MMC) snap-ins
- Active Directory Users and Computers (ADUC) extensions

You cannot use the Unisphere software for:

- Creating NDMP user accounts
- Viewing PAX statistics when performing backups
- Configuring server devices

Note: You can use any secure, encrypted, remote login application to perform the NDMP management tasks above by running the appropriate CLI-based command.

For additional information about managing VNX:

- *Learning about VNX*
- Unisphere online help
- Application's online help system in the VNX Documentation on EMC Online Support

The *Installing Management Applications on VNX for File* document includes instructions on launching the Unisphere software, and on installing the MMC snap-ins and the ADUC extensions.

Related Information

For specific information related to the features and functionality described in this document:

- *Learning about VNX*
- *EMC VNX Command Line Interface Reference for File*
- VNX for File Man Pages
- *VNX for File Release Notes*
- *Using VNX Snapsure*
- *Configuring NDMP Backups to Disk on VNX*
- *Using the server_archive Utility for VNX*
- *VNX System Operations*
- *Using FTP, TFTP, and SFTP on VNX*
- *Using VNX FileMover*
- *Using International Character Sets on VNX for File*
- *Parameters Guide for VNX for File*

Other related backup software documents:

BakBone NetVault:

- *NetVault Backup NDMP Plugin User's Guide*
- *NetVault 7.x Administrator's Guide*

CA BrightStor:

- *BrightStor ARCserve Backup for Windows Getting Started*
- *BrightStor ARCserve Backup for Windows Administrator Guide*
- *BrightStor ARCserve Backup for Windows NDMP NAS Options Guide*
- *BrightStor ARCserve Backup for Windows Tape and Optical Library Option Guide*
- *BrightStor ARCserve Backup Storage Area Network (SAN) Option Guide*

CommVault Galaxy:

- *CommVault Galaxy Client Installation and Administration Guide* (VNX NDMP)
- *CommVault Galaxy Pre-Installation Checklist*

EMC Data Manager:

- *Backup of VNX File Server Release Notes*
- *Data Manager Software Release Notes*

EMC NetWorker®:

- *NetWorker Administration Guide* (NDMP chapter)
- *NetWorker Performance Tuning Guide*

HP OpenView Storage Data Protector:

- *HP OpenView Storage Data Protector Administrator's Guide*
- *HP OpenView Storage Data Protector Windows Integration Guide*
- *HP OpenView Storage Data Protector UNIX Integration Guide*
- *HP OpenView Storage Data Protector Installation and Licensing Guide*

IBM Tivoli Storage Manager:

- *IBM Tivoli Storage Manager Administrator's Guide* (Windows or UNIX)

Syncsort Backup Express:

- *Backup Express Installation Guide Release 2.3*
- *Backup Express Interface Guide Release 2.3*
- *Using the Backup Express GUI*
- *Backup Express Jukebox and Device Setup Release 2.3*
- *Backup Express Operations Guide Release 2.3*
- *Backup Express Configuration Guide Release 2.3*

Symantec NetBackup:

- *Veritas NetBackup for NDMP System Administrator's Guide*

- *Veritas NetBackup Installation Guide*
- *Veritas NetBackup System Administrator's Guide*

EMC VNX documentation on EMC Online Support

The complete set of EMC VNX series customer publications is available on EMC Online Support. To search for technical documentation, go to http://Support.EMC.com. After logging in to the website, click **Support by Product** and type **VNX series** in the Find a Product text box. Then search for the specific feature required.

2 Concept

NDMP allows you to control the backup and recovery of an NDMP server through a network backup application, without installing third-party software on the server. On VNX, the Data Mover functions as the NDMP server.

NDMP-based backups are used for high-capacity backups and in environments where multiprotocol support is required. The NDMP protocol separates the control and data transfer components of a backup or restore operation. The actual backups are handled by the Data Mover, which minimizes network traffic.

For successful backup and restore, the length of the path including the filename should not have more than 4,095 bytes.

Note: Configuration options differ depending partly on the backup software and TLU manufacturer.

This section presents an overview of these NDMP configurations:

- Basic NDMP backup configuration on page 21
- Single TLU configuration on page 22
- Dual TLU configuration on page 22
- Three-way backup configuration on page 23
- Backup levels with VNX and NDMP on page 24
- NDMP V4 IPv6 Extension for VNX on page 24
- Restore PAX file systems on page 26
- Dynamic drive sharing on page 26
- VTLUs on page 27
- VNX DDAR support feature on page 27
- FileMover file system backup with NDMP on page 27
- File system backup with NDMP using an integrated SnapSure checkpoint on page 28

- Backup/restore with NVB on page 29
- NVB requirements on page 31
- NVB restrictions on page 32
- Filter to exclude directories or files from backup on page 32
- Valid path rules to exclude directories from backup on page 33
- Valid file-pattern rules to exclude files from backup on page 34
- Tape silvering on page 34
- Planning considerations on page 34
- Data Mover-to-TLU cabling examples on page 35
- Implement international character support with NDMP on page 36
- NDMPcopy on page 37
- NDMP last access time on page 42
- Multithreaded backup and restore design on page 42

Basic NDMP backup configuration

The NDMP architecture uses a client/server model in which the backup software is the NDMP client to the NDMP-Host Data Mover, as shown in Figure 1 on page 21. Backup data flows from the system to the Data Mover to an attached tape library backup device, without traversing the network. Only the backup software's control data, such as scheduling, cataloging, and TLU commands, travel across the network.

The Data Mover maintains a state machine for each NDMP client connection that executes and maintains the backup and restore processes. A maximum of eight simultaneous NDMP sessions are supported on each NDMP-Host Data Mover. Each Data Mover supports up to eight concurrent NDMP sessions at one time. For example, server_2 can run four backup sessions and four restore sessions simultaneously.

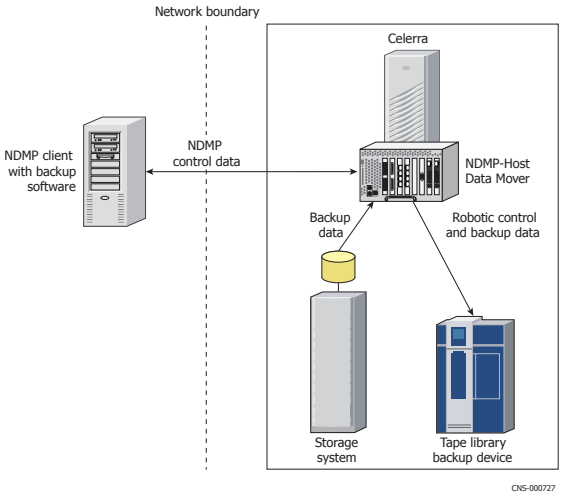

Figure 1. Basic NDMP backup

Note: Optionally, you can configure a Data Mover to serve as a VTLU and provide disk storage for backups, eliminating the need for a physical TLU. *Configuring NDMP Backups to Disk on VNX* explains how to create a VTLU on the Data Mover.

Single TLU configuration

Figure 2 on page 22 demonstrates the first and most common TLU configuration. When a TLU configuration consists of several Data Movers sharing one robot, only one Data Mover needs to be connected to the robot. One Data Mover can execute backups and robot operations. Robotic commands are executed from the backup software.

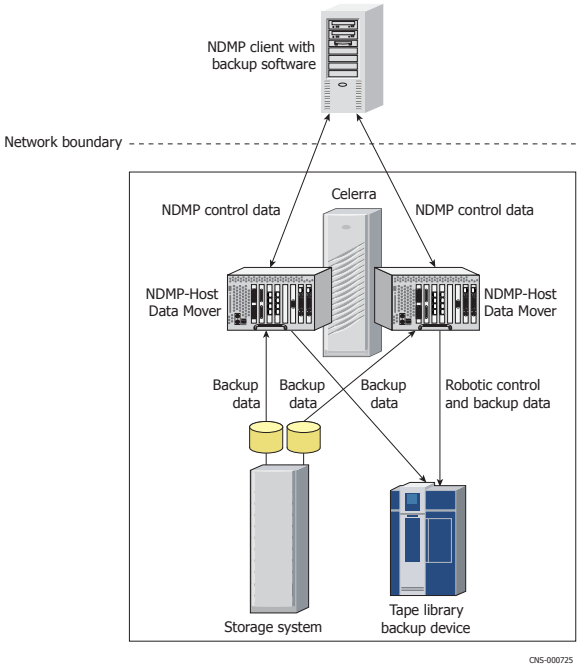

Figure 2. Single TLU NDMP backup

Dual TLU configuration

Figure 3 on page 23 shows two Data Movers connected to two separate storage system volumes with two separate tape-library-backup devices. The NDMP client controls and launches the backup. The Data Mover processes the data and sends it to the tape device.

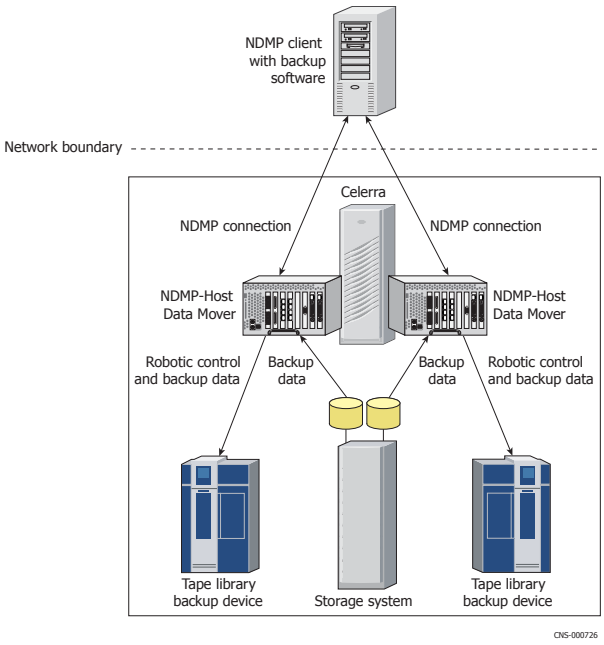

Figure 3. Dual TLU NDMP backup

Three-way backup configuration

You can direct backup data from one Data Mover, which is the NDMP backup host, through a network connection to a remote NDMP tape server to its attached tape backup device, with NDMP version 2.0 and later, and some backup software products provided by your vendor. The NDMP tape server can be a Data Mover or another qualified server. This configuration is known as three-way backup.

Three-way backup involves these hosts:

- NDMP backup client running qualified NDMP backup software
- Data Mover (NDMP data server) containing the data to be backed up (reads from the system and produces a datastream for backup)
- Data Mover or qualified server running the tape service (reads the datastream and writes it to tape)

Note: VNX supports heterogeneous three-way backups with the EMC-qualified backup vendors listed in the E-Lab Interoperability Navigator.

In a three-way backup configuration, backup data and control data travel across the network. This connection can be a dedicated network to avoid network traffic. Figure 4 on page 24 illustrates a Data Mover-to-Data Mover, three-way backup.

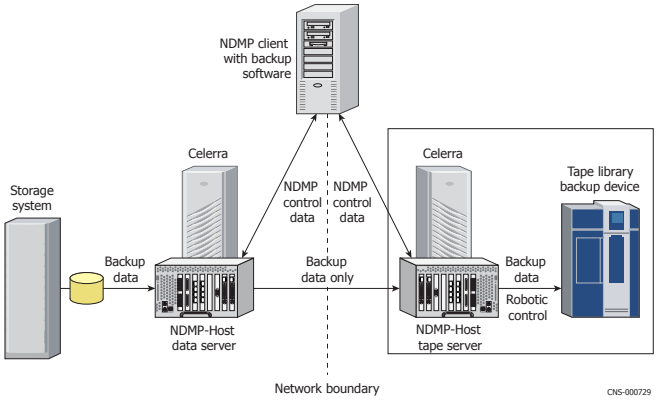

Figure 4. Three-way NDMP backup

Backup levels with VNX and NDMP

A token-based incremental backup uses the BASE_DATE and DUMP_DATE environment variables to invoke full and incremental backups. In a token-based backup, the NDMP client, rather than the Data Mover, maintains the database of backup timestamps, which are used to identify files that should be included in incremental backups. The BASE_DATE environmental variable specifies the incremental count field and the date/time of a previous backup, and this information serves as the base for a current incremental backup. The NDMP client never computes the value of BASE_DATE. Rather, the NDMP client either sets the BASE_DATE by using the DUMP_DATE it receives from the NDMP server after a previous backup of the same path, or it sets the BASE_DATE to zero, thus requesting a full backup.

The DUMP_DATE environment variable specifies the dump-date token for the current backup operation. At the end of a successful backup operation, the NDMP client should retrieve the DUMP_DATE environment variable. It retrieves this variable by issuing an NDMP_DATA_GET_ENV request while the NDMP data server is in a halted state. DUMP_DATE is returned only if BASE_DATE was previously specified by the NDMP client as part of the backup environment. When BASE_DATE is specified, the LEVEL environment is ignored.

NDMP defines 0–9 for dump levels, with level 0 being a full backup. For incremental-backup levels 1–9, only changes since the next available lower-level incremental backup or since a full backup are copied. VNX NDMP supports level 10, which allows any number of incremental backups. It backs up all files modified since the most recent backup, which can be a level 10 or lower, if no level 10 backup was run previously.

NDMP V4 IPv6 Extension for VNX

The NDMP V4 IPv6 Extension for VNX facilitates the NDMP backup by using Internet Protocol version 6 (IPv6) in addition to the existed support of IPv4. When a data connection between two servers is created using TCP/IP, the NDMP requests/responses can include both IPv4 and IPv6 address types.

IPv6 has a vastly larger address space than IPv4. IPv6 uses a 128-bit address, thus supports 2^{128} (about 3.4×10^{38}) addresses. This expansion provides flexibility in allocating addresses, routing traffic, and eliminates the primary need for network address translation (NAT). VNX and DMA needs to be configured to support IPv6 in order to use this feature.

Restore PAX file systems

A restore reverses the backup process shown in Basic NDMP backup configuration on page 21. This operation, initiated by the backup software, retrieves the data from the backup device, sends it to the Data Mover, and restores the data on the system. The benefit of the NDMP backup and restore is that it has no traffic impact on the data network.

Note: VNX database on the Control Station cannot be backed up to tape by using NDMP backup. Rather, this database is backed up automatically each hour by a script running on the Control Station. It can also be backed up manually, as described in *VNX System Operations*, and moved to a safe, remote location by using the File Transfer Protocol (FTP), as described in *Using FTP, TFTP, and SFTP on VNX* .

Restoring space-reduced files from a PAX-based NDMP backup into a file system causes the deduplication state to be set to suspended if the deduplication state is not currently enabled. *Using VNX File Deduplication and Compression* provides more information.

Note: You cannot restore deduplicated files from an NDMP backup performed on a system that is running VNX OE 7.1.47.5 or later to a system running version 5.6.46 or earlier.

Dynamic drive sharing

Several backup software vendors are developing the ability to allow more than one NDMP-Host Data Mover to share single or multiple tape drives. The backup software controls the ability of the NDMP-Host Data Movers to share the tape drives and mediates the negotiation for tape-drive access during simultaneous backups. EMC supports sharing tape drives over a SAN as shown in Figure 5 on page 26.

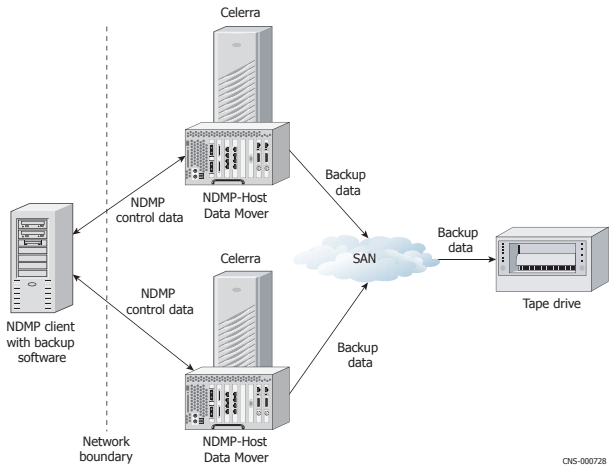

Figure 5. Dynamic drive sharing

VTLUs

A VTLU is a software-based TLU that resides on a Data Mover and emulates a physical TLU. It is also known as NDMP2D. A VTLU allows for compatibility between NDMP and disk-based storage. The NDMP-compliant backup software cannot distinguish between a physical TLU and a VTLU. This allows you to manage the VTLU as if it were a physical TLU. *Configuring NDMP Backups to Disk on VNX* provides more information on VTLUs.

VNX DDAR support feature

The enhanced DAR functionality, called DDAR, optimizes the data recovery operation by allowing the NDMP client to directly access backed-up data anywhere in a tape set without having to traverse the tape set sequentially. This method can save significant time in the restore operation. DDAR is realized only when you use EMC-qualified backup vendor software that fully supports DDAR. The EMC E-Lab Interoperability Navigator lists the EMC-qualified vendors supporting DDAR for this release.

Other improvements include enhancements in the PAX DUMP format incremental backup operation, and the three-way DAR restore operation. With DUMP incremental backups, only the changed files and the directories leading up to them are backed up, and this could save a lot of tape space for a big file system. The three-way DAR performance enhancement changes the way the data is acquired, and can dramatically improve the performance regarding small files. No environmental variable changes are required to realize these enhancements.

FileMover file system backup with NDMP

VNX FileMover-enabled file system, used in conjunction with external policy migration software, automatically migrates infrequently used files to less-expensive secondary storage systems. VNX is the primary data storage system, while the secondary system might be another VNX or a network file server supporting CIFS or NFS.

When a file migrates to secondary storage, the file contents migrate to a file on secondary storage and the original file is converted to a stub file. The stub file contains all the metadata of the file including its actual location on secondary storage to ensure that VNX can locate and recall the file data as needed.

When you back up a FileMover-enabled file system by using NDMP, you have the choice of backing up only the content on primary storage, including stub files. Or you can back up the content on primary and secondary storage, which also includes the content to which the stub files point. By default, only the primary storage content is included in an NDMP backup. The stub files on the primary system that include only file metadata are backed up without the file content. To include the file system content on primary and secondary storage in an NDMP backup, set the environment variable EMC_OFFLINE_DATA=y in the backup software. This variable must be set before you perform a backup to ensure that the backup contains all file system data from primary and secondary storage.

File system backup with NDMP using an integrated SnapSure checkpoint

Creating a checkpoint using EMC VNX SnapSure™ in the backup process lets you back up a consistent, point-in-time image of a live file system without halting activity on the file system during the backup process.

In earlier versions, using a checkpoint for a backup required that you create a checkpoint prior to running NDMP, possibly manage the checkpoint during the backup, and then delete the checkpoint when the backup completed. These manual tasks were required each time you initiated an NDMP backup. The backup failed if the Production File System (PFS) was mounted as read-only.

With version 7.0 and later, backup using integrated checkpoint supports a read-only file system or its subdirectory as the backup root. You can automate the checkpoint creation, management, and deletion activities by setting the SNAPSURE=y environmental variable to y in the qualified vendor backup software. The E-Lab Interoperability Navigator lists the backup software vendors that support the SNAPSURE=y environmental variable. This integrated checkpoint feature can be used with other EMC backup features, such as NVB and File Filtering.

If the backup software does not support the SNAPSURE=y environmental variable, you can set the parameter NDMP.snapsure in the Data Mover to use VNX SnapSure for backup. The default parameter value 0 should be changed to 1. You must ensure that the variable SNAPSURE=n is not set on the backup software. The SNAPSURE=n variable overrides the NDMP.snapsure parameter.

When the SNAPSURE=y variable is set in the backup software, and then each time you run a particular job a checkpoint of the production file system is automatically created and mounted as read-only, before the NDMP backup starts. The checkpoint is automatically used for the backup, allowing production activity to continue uninterrupted on the file system. During the backup process, the checkpoint is automatically managed (for example, autoextended if needed, space permitting). And when the backup completes, regardless of whether the operation succeeded or failed, the checkpoint is automatically deleted.

A checkpoint cannot be backed up by using the integrated checkpoint feature. You can back up only the checkpoint of an IP replication read-only target file system by using the NDMP backup feature. If you back up a replication read-only target file system, NDMP will fail when replication is updating the change.

For the checkpoint to be automatically created and deleted in the backup process, the Data Mover needs to communicate with the Control Station, which is the normal operation.

The checkpoint that is automatically created in the backup process is given a unique, system-defined name that appears in the automaticNDMPCkpts directory, and in the checkpoint list for VNX SnapSure, as follows:

automaticTempNDMPCkpt*<id>-<srcFsid>-<timestamp>*

where:

<id> = checkpoint ID

<srcFsid> = ID of the source file system being backed up using the checkpoint

<timestamp> = date and time of checkpoint creation

Checkpoints with this name are for NDMP use only and must not be unmounted, refreshed, or deleted.

VNX supports a maximum of eight concurrent NDMP sessions for each Data Mover and 2048 checkpoints per file system. The checkpoint that is created in each NDMP session counts toward the maximum number of checkpoints per file system. When the backup process completes, the checkpoint is autodeleted and as with any checkpoint deletion, the checkpoint count is reduced by one. *Using VNX SnapSure* and *Using VNX Replicator (V2)* provide more details on checkpoints.

Each time an NDMP checkpoint is created or deleted, the cmd log is modified with the updated information. The autocreated checkpoint uses the naming convention:

automaticTempNDMPCkpt*<id>-<srcFsid>-<timestamp>*

CAUTION If, for any reason, the checkpoint is not created or mounted, the NDMP backup fails.

When the backup completes, the checkpoint is unmounted and then deleted along with its mount point. The integrated checkpoints are transparent and require no administration.

Backup/restore with NVB

VNX NVB is an EMC-specific type of NDMP backup mechanism that backs up data blocks at a volume level, rather than at a file level. NVB reads a set of disk data blocks in an efficient manner compared to the method used for traditional, file-based backups. NVB works only with EMC-qualified vendor backup software, as listed in the E-Lab Interoperability Navigator. *VNX for File Release Notes* contains late-breaking administrator support information.

Note: For NVB, volume refers to the entire metavolume associated with a file system.

NVB can be used when performing a full backup or an incremental backup. And as discussed later in this section, the mechanism also supports two types of restore: full destructive restore, which delivers the best possible write performance, and file-level restore.

As shown in Figure 6 on page 30, VBB uses the NDMP protocol to interface with backup administration clients, yet it is completely different from the PAX (tar or dump) types of backup. You must specify the VBB backup type before you initiate the backup to invoke VBB functionality. When HIST=y, the environmental variable is enabled in the backup vendor's software. The VBB feature creates a file history that appears similar to one created by the dump backup type.

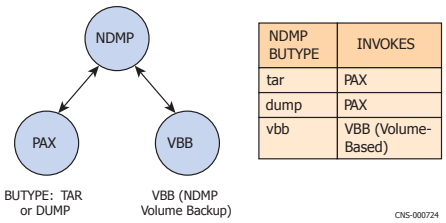

NDMP BUTYPE	INVOKES
tar	PAX
dump	PAX
vbb	VBB (Volume-Based)

Figure 6. Invoking NVB

VNX File Deduplication and Compression-enabled file systems can be backed up by using NVB and restored, using the full destructive restore (FDR) method. Backing up a deduplicated file system does not cause any data reduplication, because NVB operates at the block level (while preserving the history of the files it backs up). The data in the file system is backed up in its reduced form. The benefits of the space-saving storage efficiency realized in the production file system, from compression and file-level deduplication, flow into backups. However, a file-by-file restore of space-reduced files from NVB backups is not supported. It is recommended that a file-by-file restore be performed by using local or remotely replicated checkpoints using VNX SnapSure, and not from tape.

NVB requirements

NVB requires a read-only file system. Any of these options can be used to have a read-only file system as the source for an NVB:

- Use the Backup-with-Integrated-Checkpoints feature, which automatically creates, mounts as read-only, and uses the new checkpoint of the file system for the backup, and then deletes the checkpoint when the backup completes. Back up a file system with NDMP by using an integrated SnapSure checkpoint on page 75 provides the related procedures.
- Manually create a new checkpoint using VNX SnapSure, or specify an existing checkpoint. Checkpoints are created and automounted as read-only file systems. *Using VNX SnapSure* explains the procedure.
- Manually create a new EMC TimeFinder® /FS snapshot, or specify an existing snapshot. Snapshots are created as read-only file systems. *Using TimeFinder/FS, NearCopy, and FarCopy on VNX* explains the procedure.
- Mount the file system as read-only.

NVB restrictions

These restrictions apply when using NVB:

- NVB does not support selectively backing up a subset of the file system. For example, you can back up ufs1, but not ufs1/dir1.
- NVB and the File and Directory Filtering features do not work together. If both are specified, NVB overrides and the filters have no effect.
- NVB and DDAR functionality do not work together.
- If the File History variable is disabled, for example, HIST=n, then only full restores can be performed from an NVB backup. Also, individual files or directories cannot be restored.
- On FileMover-enabled file systems, NVB does not support the NDMP environment variable, EMC_OFFLINE_DATA. In this case, the contents of migrated files, the offline files, on secondary storage must be backed up independently.

Filter to exclude directories or files from backup

Filtering for Backup Clients is a feature introduced with version 5.5 to enable specific directories and files within a file system to be excluded from an NDMP backup session, as needed. The filtering feature is implemented by using one or both of these environmental variables in the qualified backup-vendor software: EMC_EDIR[01-50] and EMC_EFILE[01-50]. File filtering is done based on the UNIX path. You can use these variables to exclude noncritical directories and files, such as .mp3, .wma, and .temp, from the backup session.

The benefits of using file and directory filtering include:

- Shorter backup Windows
- More flexibility to structure a backup session
- Efficient use of tape, that is, space is not consumed by non-critical data

The E-Lab Interoperability Navigator lists the EMC-qualified backup software vendors that currently support the environmental variables used for filtering.

Note: The filtering feature is supported only by NDMP PAX (dump or tar) backup types. Do not use file and directory filtering with the NVB.

Some examples of how the Filtering for NDMP Backup Clients feature can be used:

Example 1: A file system has one directory under /ufs1 that contains a large amount of critical data in a database (/ufs1/db). It also has many other directories that contain non-critical data. You want to back up all the data, but on different schedules. To back up the critical data (/ufs1/db) every day, and then back up the less critical data once a week, you can exclude

the large database directory from the weekly backup by using the EMC_EDIRXX environmental variable as follows:

EMC_EDIR01=/ufs1/db

Example 2: To back up ufs1 and exclude directory /ufs1/db, all .mp3 and .wma files, and all files with the prefix temp, set the variables as follows:

EMC_EDIR01=/ufs1/db

EMC_EFILE01=*.mp3

EMC_EFILE02=*.wma

EMC_EFILE03=temp*

Example 3: To back up ufs1 and exclude any files ending in the doc extension, and any files that begin with ab, set the variables as follows:

EMC_EFILE01=*doc

EMC_EFILE02=ab*

A file named abc.doc would be excluded from this backup.

Valid path rules to exclude directories from backup

Follow these rules to specify the valid path:

- The path must start with a forward slash (/) and the file system name.
- The path can contain a maximum of 4,095 bytes.
- Supported wildcards are: question mark (?), which matches any single character, and asterisk (*), which matches any sequence of zero or more characters.
- Users cannot use the question mark and asterisk wildcards on the same pattern.
- For directory filtering, the question mark and asterisk wildcards can appear only at the end of the string.

Note: The question mark wildcard is not accepted when VNX international character sets are enabled. A message is sent to the server log when this is attempted.

- The pattern-match algorithm is case-sensitive, by default. To make it case-insensitive, use the pax param filter.caseSensitive. The default value 1 should be changed to 0.

Valid file-pattern rules to exclude files from backup

Follow these rules to specify a valid file-pattern:

- The file pattern can contain a maximum of 255 ASCII characters.
- Supported wildcards are: question mark (?), which matches any single character, and asterisk (*), which matches any sequence of zero or more characters.

  Note: The question mark wildcard is not accepted when VNX international character sets are enabled. A message is sent to the server log when this is attempted.

- Users cannot use the question mark and asterisk wildcards on the same pattern.
- For file filtering, valid wildcards can appear at the beginning or end of the string, but not both.
- The pattern-match algorithm is case-sensitive, by default. To make it case-insensitive, use the pax param filter.caseSensitive. The default value 1 should be changed to 0.

Tape silvering

Copying the baseline source file system from the source to the destination site over the IP network can be a time-consuming process. You can use an alternate method by copying the initial checkpoint of the source file system, by backing it up to a disk array or tape drive, and by transporting it to the destination site. This method is known as silvering. *Using VNX Replicator (V2)* and Perform tape silvering on page 88 provide more information.

Note: This special backup is used only for transporting replication data.

Planning considerations

This section includes information on these considerations:

- Data Mover-to-TLU cabling examples on page 35
- Implement international character support with NDMP on page 36
- Multithreaded backup and restore design on page 42

Data Mover-to-TLU cabling examples

This section includes examples of SCSI and Fibre Channel cabling from the Data Mover to the TLU. For more information on cabling a TLU to a backup host, check the backup software documentation provided by your vendor.

SCSI cabling configuration

In Figure 7 on page 35, two SCSI connections connect the NDMP-Host Data Movers to the TLU. One SCSI connection is daisy chained to two tape drives, and the other is daisy chained to two tape drives and a robot. All Data Movers in Figure 7 on page 35 have two direct SCSI connections to VNX.

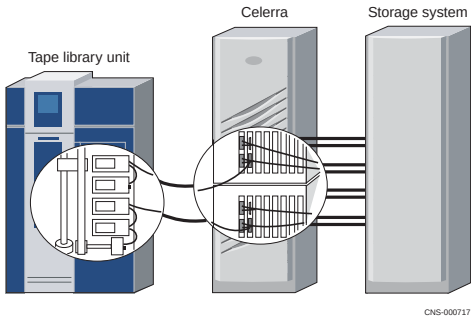

Figure 7. Sample SCSI cabling configuration

Robot connectivity

If a configuration consists of several Data Movers sharing one TLU robot, you need to connect only one of these Data Movers to the robot. Robotic control is configured in the backup software. For more information on robotic control, consult the appropriate backup software documentation provided by your vendor.

Fibre Channel configuration

Figure 8 on page 36 shows the Fibre Channel connection from the system to a switch and from the switch to the top port of each Data Mover. The bottom port of each Data Mover connects to another switch through Fibre Channel, and from that switch to the TLU through Fibre Channel.

In a configuration where SCSI is used to connect to a TLU, the SCSI connection is achieved by connecting the second switch to a bridge through Fibre Channel, and then the SCSI cable connects the bridge to the TLU.

Note: You can use one switch by connecting the system and the TLU to the same switch. Figure 8 on page 36 illustrates a two-switch configuration.

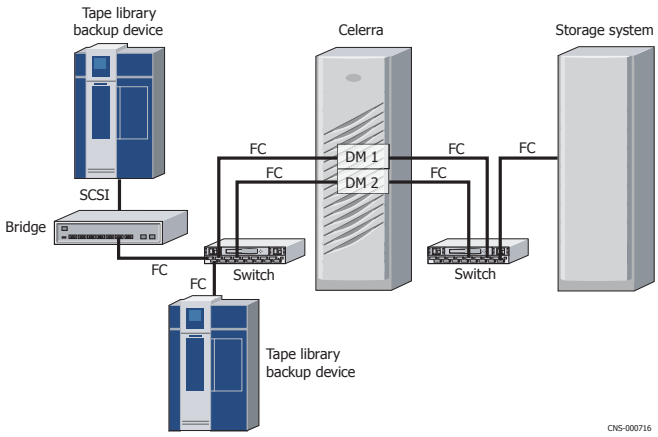

Figure 8. Sample Fibre Channel configuration

Implement international character support with NDMP

VNX supports international character sets if you have Unicode enabled. VNX supports the architecture-independent UTF-8 Unicode format. This section includes these subsections related to performing NDMP backups with Unicode enabled:

- Verify vendor support for international character sets on page 36
- Data backup-type behavior with international character sets on page 37

Note: Unicode, once enabled, cannot be disabled.

Verify vendor support for international character sets

Verify that the EMC-qualified NDMP backup software you chose supports international character sets. The backup software you chose with NDMP determines which of the two VNX implementations to follow:

- If your vendor supports international character sets, follow the vendor instructions for implementation and follow up with the backup software vendor regarding any issues you encounter.

  You must set the Data Mover parameters for NDMP and for international character sets. Configure Data Movers for international character sets on page 48 explains how to set the parameters for international character sets. *Using International Character Sets*

on VNX for File contains more information on how to use international character sets with VNX.

Note: The E-Lab Interoperability Navigator and the backup software vendor documentation list the EMC-qualified backup software vendors that support the use of international character sets.

- If your backup vendor does not support international character sets, choose the best translation configuration file for your environment, as described in *Using International Character Sets on VNX for File*.

  If Unicode is enabled on the VNX, but your backup vendor does not support international character sets, you can set up your environment to work with the best translation configuration file for your environment. The default translation dialect is UTF-8. *Using International Character Sets on VNX for File* provides details about translation files and by choosing the appropriate file for your configuration. You must also set VNX parameters as specified in Configure Data Movers for international character sets on page 48.

Data backup-type behavior with international character sets

Data is backed up onto tape by using the tar, dump, or VBB NDMP backup type. If you have chosen a translation configuration file, the filename is converted to the client encoding and sent to the client as the file history. During an NDMP tar, dump, or VBB backup, if a file was created by using a character not appearing in the code page character set, or a file is found that the translation configuration file cannot translate, the file information that is a filename with a random inode number appended is sent to the backup client to provide file history. However, this filename cannot be used to restore the file. An error log is created and the remaining NDMP backup continues normally. Files with catalog information unable to be backed up can be recovered only through a full restore of the directory.

If a surrogate pair character is used as a folder name in the VNX file system from a CIFS client, or if the CIFS folder name is longer than 255 bytes with multibyte characters, the NDMP restore operation can change the folder name to UTF8ConversionError. UNIX names on VNX are set to UTF8ConversionError as the conversion is not supported. For example, a Japanese character can use 3 bytes per UTF-8 character. In the NDMP restore operation, DAR does not handle a parent folder, and we use UNIX parent directory instead of CIFS parent folder. Filenames can be restored easily. Use non-DAR full NDMP restore or NDMP DDAR restore for such CIFS folders and files.

NDMPcopy

The NDMPcopy utility is an open source program that improves the rate of data transfer. It allows the user to copy file system data within the Data Mover or between different Data Movers that are on the same or different VNX. If the source and destination file systems are

on the same Data Mover, the data is transferred directly from the source to the destination file system without traversing and overloading the NDMPcopy host machine or network. The NDMPcopy utility uses the NDMP to transfer the file system data.

The NDMPcopy utility contacts the NDMP server processes on the source and destination Data Movers. It authenticates the connection, and initiates data transfers between the source and destination file systems. The NDMP host receives status messages only.

The NDMPcopy utility also provides multiprotocol (NFS/CIFS) support. It is recommended to use NDMPcopy for data transfer only. When using NDMPcopy to transfer files between data movers, the destination data mover must be running the same or later VNX operating environment (for File) version as the source data mover. The NDMPcopy utility can affect the scheduling of backup activities in the same source, as it modifies the BackupDates file on the Data Mover. It uses text mode for authentication and so the password will be sent across the network in a non-encrypted format.

Note: NDMPcopy does not require any backup application or tape device.

You can perform data transfers from the source to the destination file system, by using this format:

Data Mover:/*path*

where:

Data Mover = IP address of the Data Mover

path = pathname of the directory to be used for data transfer

Prerequisites

- The NDMPcopy utility should be installed on the host machine.
- The two Data Movers can communicate with each other by using the ping command with hostnames or IP addresses that they need to contact as parameters. If the ping command does not work, the ndmpcopy command will hang. You can use the server_ping and server_route commands to detect and correct network problems.

Examples

- To transfer data from the source to the destination, on the same Data Mover, use this command syntax:

  ```
  # ndmpcopy Data_Mover_IP:/src_filesystem_path Data_Mover_IP:
  /dest_filesystem_path -sa username:password -da username:password
  ```

  where:

  -sa *username:password*: specifies the username and password of the source Data Mover. The source Data Mover authentication is used to verify the network connections to the source Data Mover. You must specify the username and include a colon (:) after the username. If you do not type the password after "username:" in

the command line, the ndmpcopy command will prompt you for the password. Use the -dpass option if the source and destination passwords differ, and if you do not want to type the password in the command line. The username and password should match the username and password that you typed for the source NDMP-Host Data Mover in Assign a username and a password to a Data Mover on page 56.

Note: If the password is not configured on the source or destination Data Movers, the ndmpcopy command will fail.

-da *username:password*: specifies the username and password of the destination Data Mover. The destination Data Mover authentication is used to verify the network connections to the destination Data Mover. Make sure you specify the username and include a colon (:) after the username. If you do not type the password after "username:" in the command line, the ndmpcopy command will prompt you for the password. Use the -dpass option if the source and destination passwords differ, and if you do not want to type the password in the command line. The username and password should match the username and password that you typed for the source NDMP-Host Data Mover in Assign a username and a password to a Data Mover on page 56.

Note: If the password is not configured on the source or destination Data Movers, the ndmpcopy command will fail.

Example:

```
# ndmpcopy 10.6.123.162:/server2_fs1 10.6.123.162:/server2_fs2 -sa
s2_user:s2_password -da s2_user:s2_password
```

Output:

```
Connecting to 10.6.123.162.
Connecting to 10.6.123.162.
10.6.123.162: CONNECT: Connection established.
10.6.123.162: CONNECT: Connection established.
10.6.123.162: LOG: server_archive: emctar vol 1, 16 files, 0 bytes
read, 279236 bytes written
10.6.123.162: HALT: The operation was successful!
Waiting for 10.6.123.162 to halt too.
10.6.123.162: LOG: server_archive: emctar vol 1, 16 files, 279236
bytes read, 0 bytes written
10.6.123.162: HALT: The operation was successful!
The transfer is complete.
```

- To transfer data from the source to the destination file systems that reside on different Data Movers:

  # ndmpcopy *Source_Data_Mover_IP:/src_filesystem_path Destination_Data_Mover_IP:/dest_filesystem_path* **-sa** *username:password* **-da** *username:password*

  Example:

  ```
  # ndmpcopy 10.6.123.162:/server2_fs1 10.6.123.164:/server3_fs1 -sa
  s2_user:s2_password -da s3_user:s3_password
  ```

Output:

```
Connecting to 10.6.123.162.
Connecting to 10.6.123.164.
10.6.123.162: CONNECT: Connection established.
10.6.123.164: CONNECT: Connection established.
10.6.123.162: LOG: server_archive: emctar vol 1, 16 files, 0 bytes
read, 279236 bytes written
10.6.123.162: HALT: The operation was successful!
Waiting for 10.6.123.164 to halt too.
10.6.123.164: LOG: server_archive: emctar vol 1, 16 files, 279236
bytes read, 0 bytes written
10.6.123.164: HALT: The operation was successful!
The transfer is complete.
```

- To avoid typing the password in the command line:

```
# ndmpcopy 10.6.123.162:/server2_fs1 10.6.123.162:/server2_fs2 -sa
s2_user: -da s2_user:
```

Output:

```
Password:
Connecting to 10.6.123.162.
Connecting to 10.6.123.162.
10.6.123.162: CONNECT: Connection established.
10.6.123.162: CONNECT: Connection established.
10.6.123.162: LOG: server_archive: emctar vol 1, 16 files, 0 bytes
read, 279236 bytes written
10.6.123.162: HALT: The operation was successful!
Waiting for 10.6.123.162 to halt too.
10.6.123.162: LOG: server_archive: emctar vol 1, 16 files, 279236
bytes read, 0 bytes written
10.6.123.162: HALT: The operation was successful!
The transfer is complete.
```

- Use the -dpass option, as displayed in the output below, if the source and destination passwords differ, and if you do not want to type the password in the command line:

```
# ndmpcopy 10.6.123.162:/server2_fs1 10.6.123.164:/server3_fs1 -sa
s2_user: -da s3_user: -dpass
```

Output:

```
Source password:
Destination password:
Connecting to 10.6.123.162.
Connecting to 10.6.123.164.
10.6.123.162: CONNECT: Connection established.
10.6.123.164: CONNECT: Connection established.
10.6.123.162: LOG: server_archive: emctar vol 1, 16 files, 0 bytes
read, 279236 bytes written
10.6.123.162: HALT: The operation was successful!
Waiting for 10.6.123.164 to halt too.
10.6.123.164: LOG: server_archive: emctar vol 1, 16 files, 279236
bytes read, 0 bytes written
10.6.123.164: HALT: The operation was successful!
The transfer is complete.
```

- Use the -h option, as displayed in the output below, to get more information about the usage of ndmpcopy and help message:

```
# ndmpcopy -h
```

Output:

```
Usage:
ndmpcopy src_filer:/src/dir dest_filer:/dest/dir
   [-sa none|user:password] [-da none|user:password]
   [-sport ndmp_src_port] [-dport ndmp_dest_port]
   [-dhost ndmp_dest_ip_addr] [-level ndmp_dump_level]
   [-v] [-q] [-dpass] [-h]

Defaults:
src_auth_type      = text
src_auth_user      = root
src_auth_password  =
dest_auth_type     = text
dest_auth_user     = root
dest_auth_password =
ndmp_src_port      = 0 (0 means NDMP default, usually
                        10000)
ndmp_dest_port     = 0 (0 means NDMP default, usually
                        10000)
ndmp_dump_level    = 0 (valid range: 0 - 9)
ndmp_dest_ip_addr  = (no default: user needs to
                      override dest_filer value)
verbosity          = noisy
different_passwords= no
```

You can also make use of these options, if required:

- -dpass: Use this option if the passwords for the source and destination Data Movers are not the same and if you do not want to specify the password on the command line.
- -h: Provides usage information and help message for NDMPcopy.
- -sport (Optional): Port address of the source Data Mover. The default value is 0, which is the NDMP default (10000). If you specify a value other than 0 or 10000, the transfer will fail.
- -dport (Optional): Port address of the destination Data Mover. The default value is 0, which is the NDMP default (10000). If you specify a value other than 0 or 10000, the transfer will fail.
- -level (Optional): The incremental level that can be set for the data transfer. You can set this value from 0 through 9. The default level is 0.

  Note: If the source file system has not been modified after a level 0 copy, all other subsequent levels from 1 through 9 will fail until the source file system is modified.

- -v (Optional): The default value of verbosity is noisy.

NDMP last access time

Last access time (atime) is not retained during NDMP backup/restore and NDMPcopy operations and it is set to the modification time (mtime). After a file migration, the user loses the ability to archive "cold" or "inactive" data. A new environmental variable RETAIN_ATIME is added for the NDMP data stream to preserve the atime.

By default, the RETAIN_ATIME is set to "N". If the variable is set to the default value "N", VNX will restore the file and apply the atime to mtime. When the RETAIN_ATIME is set to "Y", the atime is set as the last access time approximately. VNX then restores the file and applies the retained atime to the file.

To retain the atime during NDMP backup/restore, the data stream needs to be changed. A new field atime is added to the data stream headers. When the RETAIN_ATIME environmental variable is set to "Y", the value is logged in the server log as INFO message:

```
NDMP:4:Name:RETAIN_ATIME Value:Y
```

This feature is supported by the EMC Networker and Symantec NetBackup.

Multithreaded backup and restore design

VNX uses a multithreaded backup and restore design that improves the VNX NDMP backup and restore performance in most environments. The design involves three main thread groups, referred to as NASA, NASS, and NASW. These thread groups process and deliver backup and restore data and metadata. Table 2 on page 42 describes the function of each thread group. Appendix A provides a detailed description of the parameters related to multithreaded backup. The server_pax command, presented in View backup statistics on page 70, displays backup statistics for each thread group.

Table 2. Backup thread groups

Thread	Function
NASA	◆ Sends metadata to the NDMP backup software ◆ Reads file data ◆ Writes data to the buffer pool
NASS	◆ Traverses the file system ◆ Provides metadata for each file or directory
NASW	◆ Obtains data from the data buffer pool ◆ Writes data to a tape or sends it to a remote Data Mover or NDMP server

Multithreaded backup process

When you initiate an NDMP backup, the NASA thread group activates multiple backup threads and puts metadata buffers in the stat buffer pool, which is shared by NASA and NASS. Figure 9 on page 43 shows the components of the multithreaded design.

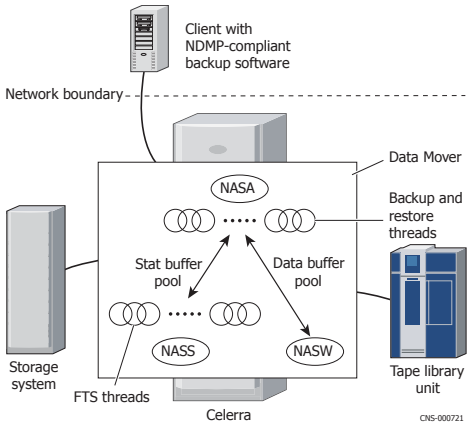

Figure 9. Multithreaded backup and restore

The stat buffer pool contains a user-definable number of buffers determined by the paxStatBuff parameter that transfers file system metadata. Upon receiving metadata buffers, NASS traverses the backup file system or directory and returns metadata to the stat buffer pool. When a metadata buffer is available, a backup thread backs up one directory or file per thread, in sequential order. When a thread completes its portion of the backup, it waits for subsequent threads to complete their backup portions before starting a new one.

Multithreaded restore process

The multithreaded restore process involves the NASA and NASW thread groups. You can define the number of restore threads through the nRestore parameter. Each restore thread is responsible for a single small file or part of a bigger file. When the backup software begins the restore process, the NASA thread activates the restore threads and waits for NASW to deliver data to the data buffer pool. The user-definable ReadBuff parameter determines the number of buffers between the NASA and NASW thread groups.

3 Configuring

The tasks to configure NDMP backups on VNX are:

CAUTION It is recommended to using the default values for most parameters except as otherwise noted in this document because these values optimize performance in most cases. Maximizing all parameter values will not always result in better performance, and could possibly result in Data Mover panic.

- Configure Data Movers on page 46
- Set up the TLU on page 53
- Assign a username and a password to a Data Mover on page 56
- Set environment variables on page 58

Configure Data Movers

The tasks to configure Data Movers are:

1. Configure the tape buffer size on page 46
2. Configure the read or write block size on page 47
3. Configure an NS series Data Mover on page 47
4. Configure Data Movers for international character sets on page 48
5. Override the autonegotiated NDMP version on page 49
6. Configure the Data Mover to increase the NDMP sessions (Optional) on page 50
7. Configure the NDMP port range on page 51

Configure the tape buffer size

The tape buffer size specifies the size of the buffer allocated for reading and writing to the tape. Most backup software vendors work with the VNX default buffer size of 128. Tape drive buffer sizes on page 129 explains how to set the correct parameter value for the tape drive.

EMC NetWorker and Fujitsu-Siemens NetWorker, however, require different buffer sizes for different tape drive types. EMC NetWorker on page 124 contains a partial list of tape buffer values used by NetWorker and Fujitsu-Siemens NetWorker.

Action

To configure the tape buffer size, use this command syntax:
`$ server_param <movername> -facility NDMP -modify bufsz -value <new_value>`
where:
<movername> = name of the Data Mover
<new_value> = value you want to set for the parameter
Example:
To set the bufsz parameter to 256, type:
`$ server_param server_2 -facility NDMP -modify bufsz -value 256`

Note: Parameter and facility names are case-sensitive.

Configure the read or write block size

The read or write block size parameter specifies the maximum allowed PAX buffer size for NDMP read or write. Most backup software vendors work with VNX default read or write block size of 64. However, the EMC Data Manager (EDM) feature requires a bigger block size for the backup or restore process to succeed. Configure EMC Data Manager with NDMP on page 120 provides more details on using EDM.

Action

To configure the read/write block size, use this command syntax:
`$ server_param <movername> -facility PAX -modify readWriteBlockSizeInKB -value <new_value>`
where:
`<movername>` = name of the Data Mover
`<new_value>` = value you want to set for the parameter
Example:
To set the readWriteBlockSizeInKB parameter to 128, type:
`$ server_param server_2 -facility PAX -modify readWriteBlockSizeInKB -value 128`

Note: Parameter and facility names are case-sensitive. The default value must be changed to 128 when using EDM for backup and restore operations.

Configure an NS series Data Mover

From version 5.4 and later, you do not have to configure the Data Mover before backing up the data to a locally attached tape drive. The hardware autodetects and configures for arbitrated loop or switched fabric connections. Further, you do not have to specify the tape backup link speed. By default, the speed will be autonegotiated to the highest mutually acceptable link speed.

Follow this procedure if the autonegotiation fails and you need to configure the link speed manually.

Note: NS series require an MIA for device connection to the AUX-0 port. Your EMC Customer Support Representative can assist in installing the MIA.

Perform the following steps to configure NS series Data Mover:

1. To specify a tape backup link speed manually on an NS series Data Mover, use this command syntax:

```
$ server_param <movername> -facility fcTach -modify linx_speed_aux0
-value <new_value>
```

where:

<movername> = name of the Data Mover

<new_value> = 0x1000 (to set tape backup link speed to 4 GB/s), 0x2000 (to accommodate a 2 GB Fibre Channel-Arbitrated Loop interface), 0x4000 (to accommodate an LC optical 1 GB interface) or 0x8000 (to detect the link speed from controller)

Example:

To set the link speed to accommodate a 2 GB Fibre Channel speed, type:

```
$ server_param server_2 -facility fcTach -modify linx_speed_aux0 -value 0x2000
```

Note: Parameter and facility names are case-sensitive.

2. To restart the Data Mover for the parameter to take effect, use the command syntax:

```
$ server_cpu <movername> -reboot -monitor now
```

where:

<movername> = name of the Data Mover

Example:

To restart server_2, type:

```
$ server_cpu server_2 -reboot -monitor now
```

Configure Data Movers for international character sets

Use this procedure to configure the Data Mover for international character sets. Implement international character support with NDMP on page 36 and *Using International Character Sets on VNX for File* provide more information on support and behavior of international character sets.

This section explains how to specify a dialect in VNX. You will need to specify a dialect if:

- Your backup software does not support Unicode, and a backup is done with Unicode enabled on VNX. The filename must then be converted.
- VNX is in internationalization mode and you are restoring a backup image of an ASCII-mode Data Mover where a non-Latin1 dialect was used.

Note: During a backup, if a file is found that the translation configuration file cannot translate a particular file, an entry is created in the server_log file. For example, if a file was created by using a character not appearing in the code page's character set, or if an NDMP dialect does not match the dialect used to create the file; then an entry is created in the server log. When restoring, you cannot perform a single-file restore on a file containing an invalid (inode_<number>) filename. To restore the file, you must restore the entire directory in which the file resides.

Perform the following steps to configure Data Movers for international character sets:

1. To determine the dialect you want to assign to the Data Mover, use these guidelines:
   - If Unicode was enabled during the backup and your backup software does not support Unicode, use the same dialect as the Data Mover client.
   - If you are restoring a backup image of an ASCII-mode Data Mover where a non-Latin1 dialect was used, and VNX is in internationalization mode, specify the dialect used on the ASCII-mode Data Mover.
2. To set the dialect, use the command syntax:

   ```
   $ server_param <movername> -facility NDMP -modify dialect
   -value <new_value>
   ```

   where:

   <movername> = name of the Data Mover

   <new_value> = value that matches the dialect on the Data Mover client

   Example:

   To set the dialect to big5, type:

   ```
   $ server_param server_2 -facility NDMP -modify dialect -value big5
   ```

   Output:

   ```
   server_2 : done
   ```

Note: Parameter and facility names are case-sensitive.

Override the autonegotiated NDMP version

Complete this procedure if you want to force VNX to use an earlier version of NDMP, such as version 3.

By default, some servers and NDMP devices autonegotiate to use the latest NDMP version supported by the Data Mover and the NDMP backup software. It is recommended to use this default setting. If necessary, you can override the process by specifying the latest NDMP version. Overriding the NDMP version can help you isolate a problem by determining

whether it occurs with different NDMP autonegotiated versions. Autonegotiation between VNX and the NDMP client determines the highest NDMP version used.

Note: Parameter and facility names are case-sensitive. The maxProtocolVersion parameter setting overrides the autonegotiation.

Action

To override the autonegotiated NDMP version, use this command syntax:
`$ server_param <movername> -facility NDMP -modify maxProtocolVersion -value <new_value>`
where:
<movername> = name of the Data Mover
<new_value> = value you want to set for the parameter
Example:
To set the NDMP version to 3, type:
`$ server_param server_2 -facility NDMP -modify maxProtocolVersion -value 3`

Output
`server_2 : done`

Configure the Data Mover to increase the NDMP sessions (Optional)

VNX supports a maximum of eight NDMP sessions for each NDMP-Host Data Mover. By default, up to four NDMP sessions are performed at a time. You can configure the Data Mover to change the value of the NDMP sessions from four up to eight.

Use the following procedures to:

- View the value of concurrent NDMP sessions on page 50
- Configure the NDMP concurrent sessions on page 51

View the value of concurrent NDMP sessions

Action

To view the value of concurrent NDMP sessions, use this command syntax:
`$ server_param {ALL |<mover>} -facility NDMP -info concurrentDataStreams`
To view the value of concurrent NDMP sessions, type:
`$ server_param server_2 -facility NDMP -info concurrentDataStreams`

Output
server_2: name = concurrentDataStreams facility_name = NDMP default_value = 4 current_value = 4 configured_value = user_action = reboot DataMover change_effective = reboot DataMover range = (1,8) description = number of concurrent data streams

Configure the NDMP concurrent sessions

Action
To change the value of concurrent sessions, use this command syntax: $ **server_param** {*ALL*|*<movername>*} **–facility NDMP –modify concurrentDataStreams –value** *<number_of_sessions>* where: *<movername>* = name of the Data Mover *<number_of_sessions>* = number of concurrent NDMP sessions to be performed Example: To change the value of the concurrent NDMP sessions to eight, type: **$ server_param server_2 –facility NDMP –modify concurrentDataStreams –value 8** Note: If the Data Mover has less than 8 GB memory, you cannot set the value of the concurrentDataStreams parameter to more than four. The server_param command will fail with the error message, "8 is not in range (1, 4)".
Output
server_2: You must reboot server_2 for concurrentDataStreams changes to take effect.

Configure the NDMP port range

You can set the NDMP port range in order to limit the NDMP data connection listening port, to support firewall configurations. To specify the port range use the portRange parameter. It is recommended to set the NDMP port range by using dynamic or private ports (49152-65535) of IANA. You can set the minimum and maximum port range as per your requirement. The required minimum value of the port width is 32, however it is recommended to set the width to at least 100. Setting the port range does not imply that these ports are reserved for NDMP connection. When the port range is set, NDMP selects a random port number from the specified port range as the data connection listening port.

The port range feature does not support NAT. It uses the existing server_param management interface and is independent of the DMA. Configuring the NDMP port range allows the Data Mover to use a limited range of TCP ports that correspond to the ports for which the firewall allows inbound connections.

Note: You could need to configure the DMA to allow specific ports for outbound connections.

Use the following procedures to:

- View the port range on page 52
- Configure the port range on page 52

View the port range

Action

To view the value of the NDMP port range, use this command syntax:
`$ server_param {ALL|<movername>} -facility NDMP -info portRange`
where:
<movername> = name of the Data Mover
Example:
To view the value of the NDMP port range, type:
`$ server_param server_2 -facility NDMP -info portRange`

Output

`server_2 :`
`name = portRange`
`facility_name = NDMP`
`default_value = 1024-65535`
`current_value = 1024-65535`
`configured_value =`
`user_action = none`
`change_effective = immediate`
`range = 1024-65535`
`description = Port range for NDMP data connection listening`

Configure the port range

Action
To set the value of the NDMP port range, use this command syntax: `$ server_param {ALL|<movername>} -facility NDMP -modify portRange -value 50000-50100` where: *<movename>* = name of the Data Mover

Action
Example: To set the value of the NDMP port range, type: **$ server_param server_2 -facility NDMP -modify portRange -value 50000-50100**
Output
server_2 : done

Note: You can also use Unisphere to modify and display the NDMP portRange parameter. In Unisphere, select Settings and click Data Mover Parameters. See the Unisphere help for more information.

Set up the TLU

Setting up the TLU includes the steps for connecting a physical TLU and retrieving and recording physical device names. *Configuring NDMP Backups to Disk on VNX* provides information on configuring virtual tape library emulation software and disk storage on a Data Mover.

You need the device names for the robot and drives when configuring the backup software. VNX addresses a TLU device component as follows:

<SCSI_controller> <SCSI_ID> <SCSI_LUN>

where:

<SCSI_controller> = chain number of the controller

<SCSI_ID> = SCSI ID (target) of the controller

<SCSI_LUN> = logical unit number (LUN) of the controller

Note: For backups with VNX, an MIA adapts the Fibre Channel AUX-0 port from an HSSDC copper interface to an LC optical interface. For backups with VNX series, an SFP adapter is used. Instructions on how to install an MIA in VNX, or an SFP in VNX series, are documented in individual setup guides provided with each system.

Perform the following steps to set up the TLU:

1. To halt each Data Mover to be connected to the TLU and confirm it has halted from the Control Station, use this command syntax:

   $ **server_cpu** *<movername>* **-halt -monitor now**

   where:

   <movername> = name of the Data Mover

   Example:

   To halt the Data Mover running on server_2, type:

```
$ server_cpu server_2 -halt -monitor now
```

2. Type /nas/sbin/getreason and ensure that the status is powered off.
3. Cable each Data Mover to the TLU. Data Mover-to-TLU cabling examples on page 35 provides more information.
4. Turn on the TLU and verify that it is online.
5. To restart each Data Mover connected to the TLU and confirm it has restarted, use this command syntax:

   $ **server_cpu** *<movername>* **-reboot -monitor now**

   where:

   <movername> = name of the Data Mover

   Example:

   To restart server_2, type:

   ```
   $ server_cpu server_2 -reboot -monitor now
   ```

   Output:

   ```
   server_2: reboot in progress 0.0.0.0.0.1.1.3.3.3.3.3.3.3.3.4. done
   ```

6. To verify that the Data Mover can recognize its TLU device, use this command syntax:

   $ **server_devconfig** *<movername>* **-probe -scsi -nondisks**

   where:

   <movername> = name of the Data Mover

   Note: The server_devconfig command discovers only physical devices. It does not discover virtual devices.

   Example:

   To verify that server_2 can recognize its TLU devices, type:

   ```
   $ server_devconfig server_2 -probe -scsi -nondisks
   ```

7. To save the TLU devices of the Data Mover to the VNX database, use this command syntax:

   $ **server_devconfig** *<movername>* **-create -scsi -nondisks**

   where:

   <movername> = name of the Data Mover

   CAUTION In VNX for block environment, before you run the server_devconfig -create command, verify that all paths to the Data Mover are active and no LUNs are trespassed. Running this command while paths are inactive causes errors in the Data Mover configuration file.

   Example:

To save the device information to the database, type:

```
$ server_devconfig server_2 -create -scsi -nondisks
```

8. To list the device addresses, use this command syntax:

   $ **server_devconfig** *<movername>* **-list -scsi -nondisks**

   where:

   <movername> = name of the Data Mover

   Example:

   To list the devices for server_2, type:

   ```
   $ server_devconfig server_2 -list -scsi -nondisks
   ```

   Output:

   ```
   server_2 :
   Scsi device table
   name addr type info
   jbox1 c1t0l0 jbox ATL P1000 62200501.21
   tape2 c1t4l0 tape QUANTUM DLT7000 245Fq_
   tape3 c1t5l0 tape QUANTUM DLT7000 245Fq_
   ```

9. Record the device name information as shown in Table 3 on page 56. It shows the type of configuration information to record. You will need the device names when you configure your backup software.

10. Perform this procedure for each Data Mover attached to the TLU.

Configure stand-alone tape drives

When configuring stand-alone tape drives, it is recommended to prefix the VNX device name with c in your backup software. This ensures that tape compression is used during a backup. For example, to use tape compression for the device C1T0L0, type cC1T0L0 as the device name. To turn off compression, prefix the device name with "u". When no prefix is specified, the default compression setting for the tape drive is used.

Table 3. Sample device name information

Data Mover	TLU SCSI port	Resulting device name	Device component
server_x	chain 1, target 0, LUN 0	C1T0L0	Robot
server_y	chain 1, target 4, LUN 0	C1T4L0	Tape Drive 1
server_z	chain 1, target 5, LUN 0	C1T5L0	Tape Drive 2

Configure VTLUs

To configure VNX for NDMP backups to disk, configure a VTLU, and then associate storage with the VTLU. *Configuring NDMP Backups to Disk on VNX* contains information about configuring VTLUs. VTLU backups require the same VNX configuration settings as NDMP backups to tape. Depending on your backup software vendor, you will need to complete some or all of these tasks that are described in these sections:

- Configure the Data Mover on page 141
- Assign a username and a password to a Data Mover on page 56
- Set environment variables on page 58

Assign a username and a password to a Data Mover

NDMP configuration for the backup software and VNX requires setting a trusted username and password for each NDMP-Host Data Mover. The methods to set the username and password are:

- With NIS on page 57
- Without NIS on page 57

With NIS

For Network Information Service (NIS) installations, the username and password must match the username and password in the NIS database.

Note: The same username and password should be typed when you configure the backup software.

NIS must be enabled on every NDMP-Host Data Mover.

Note: Use the text password method only to create an NDMP account in the NIS environment. Also, ensure that the NDMP.md5 parameter is set to 0 (zero) on all Data Movers in this environment. The *Parameters Guide for VNX for File* provides information about the NDMP.md5 parameter.

Without NIS

Perform the following steps to assign a username and password to a Data Mover without NIS:

1. To assign a user account name and password to one or more Data Movers, log in to the Control Station as nasadmin and switch user to root by typing:

   ```
   $ su
   ```

2. Type the root password when prompted.

3. To create an account, use the appropriate command syntax, as follows:

   Text method:

   ```
   # /nas/sbin/server_user <movername> -add -password <username>
   ```

   Note: Before using the text method, change the value of the NDMP.md5 parameter to 0.

   MD5 password encryption method:

   ```
   # /nas/sbin/server_user <movername> -add -md5 -password <username>
   ```

   where:

   <movername> = name of the Data Mover to which you want to assign a user account and password

   <username> = name of the user account you want to assign to the Data Mover

   Note: For security reasons, it is recommended that you use the MD5 password encryption method.

Output:

```
Creating new user user_name
User ID: 1000
Group ID: 1000
Home directory:
Changing password for user user_name
New passwd:
Retype new passwd:
server_2 : done
```

Note: In the output, the two mandatory fields, User ID (UID) and Group ID (GID), are integers. VNX uses UNIX-style UIDs and GIDs to record the ownership of files and directories. The UID of the root user is 0. The password you assign to the Data Mover can contain between six and eight characters.

Examples:

To create an account by using the text method, type:

```
# /nas/sbin/server_user server_2 -add -password user_name
```

To create an account by using the MD5 encryption method, type:

```
# /nas/sbin/server_user server_2 -add -md5 -password user_name
```

4. Repeat steps 1, 2, and 3 for each NDMP-Host Data Mover.

Set environment variables

Table 4 on page 58 lists environment variables supported by VNX. Refer the backup vendor documentation for more information about environment variable requirements, as backup vendors use and implement variables differently.

Table 4. Environment variables

Variable	Possible values	NDMP spec default value	Description
BASE_DATE	Obtained from the NDMP client for the previous back-up operation or set to 0 (zero) for a full backup.	N/A	Incremental count and date/time of previous backup operation. Note: This variable works in conjunction with DUMP_DATE. When BASE_DATE is specified, LEVEL is ignored. Backup levels with VNX and NDMP on page 24 offers more details.
DIRECT	y/n	n	Improves restore results. Must be enabled to invoke DAR functionality.

Table 4. Environment variables *(continued)*

Variable	Possible values	NDMP spec default value	Description
DUMP_DATE	Obtained from the NDMP server for the current backup operation only if BASE_DATE is specified by the NDMP client as part of the environment.	N/A	Incremental count and date/time of current backup operation. Note: This variable works in conjunction with BASE_DATE. Backup levels with VNX and NDMP on page 24 shows how to use DUMP_DATE in a backup session.
EMC_EDIR[01-50]	<path name> as /fs/directory	blank	Specifies the directory to exclude from the NDMP backup session. Exclude directories or files from backup by using filtering on page 86 provides details on using this directory filter. By default, the maximum number of directory filters is five. You can increase this to a maximum of 50 by using the filter.numDirFilter parameter. Modify parameters on page 69 provides more details.
EMC_EFILE[01-50]	<filename> or <file type>	blank	Specifies the file, or a type of file, to exclude from the NDMP backup session. Exclude directories or files from backup by using filtering on page 86 shows how to use this file filter. By default, the maximum number of file filters is five. You can increase this to a maximum of 50 by using the filter.numFileFilter parameter. Modify parameters on page 69 provides more details.
EMC_OFFLINE_DATA	y/n	n	By default, VNX backs up the attributes associated with migrated files. Y specifies that VNX backup migrated data as well as attributes. This variable applies when using NDMP to back up VNX FileMover-enabled file system. Back up a FileMover file system with NDMP on page 74 and *Using VNX FileMover* contain more details on VNX FileMover.
FILESYSTEM	<pathname>	N/A	The path to be backed up.

Table 4. Environment variables *(continued)*

Variable	Possible values	NDMP spec default value	Description
HIST	y/n	n	Determines if a file history is created. Set HIST= y before you perform a backup to be able to perform a direct-access restore operation later.
LEVEL	0–10	0	Dump level. Backup levels with VNX and NDMP on page 24 explains how to set backup levels.
OPTION	NT LK MI/MD/MM	NT for tar and dump	Saves NT attributes. Follows symbolic links. Sets collision policy for localization during restore.
OVERWRITE	y/n	y	Determines whether a file should be restored or skipped, if a file of the same name exists on the destination location. The OVERWRITE variable can be set using your backup application. If the value of this variable is set at the time of backup, it will be applicable to all the restores performed from that backup. This value can be changed, at the time of restore, if it is allowed by the backup application.
PREFIX	<pathname>	N/A	PREFIX can be used in place of FILESYSTEM for backward compatibility.
RECURSIVE	y/n	y	Recursively restores any given directory. In version 5.5 or later, recursive must be enabled to affect expected restore functionality.
SNAPSURE	y/n	n	Determines if a checkpoint created using VNX SnapSure is automatically created, mounted, used for the backup, and deleted when the backup completes. This must be set to y for these automated activities to occur. Once set to y, the automation occurs every time you run that particular job. Back up a FileMover file system with NDMP on page 74 explains how to use this feature.

Table 4. Environment variables *(continued)*

Variable	Possible values	NDMP spec default value	Description
TYPE	tar or dump or vbb	N/A	The backup type. Note: Dump format does not cross file system boundaries.
VLC	y/n	n	Specifies an NDMP Volume-Level Copy to tape. You must use a read-only file system for a volume-level copy. The VLC option can be used for the first fs_copy when configuring EMC VNX Replicator. Note: This feature has not been qualified with backup software vendors.
VBB	y/n	n	Specifies backup type to VBB.
RETAIN_ATIME	Y/N	N	By default, the RETAIN_ATIME is set to "N". VNX restores the file and apply the atime to mitime. When the RETAIN_ATIME is set to "Y", the atime is set as the last access time approximately. VNX restores the file and applies the retained atime to the file. NDMP last access time on page 42 provides more information.

4

Managing

The tasks to manage NDMP backups are:

- Manage backup parameters and statistics on page 64
- Back up a VDM file system with NDMP on page 74
- Back up a FileMover file system with NDMP on page 74
- Restore a FileMover file system on page 75
- Back up a file system with NDMP by using an integrated SnapSure checkpoint on page 75
- Perform backup and restore with NVB on page 75
- Perform NVB restores on page 77
- Exclude directories or files from backup by using filtering on page 86
- Perform tape silvering on page 88

Manage backup parameters and statistics

The tasks to manage backup and restore parameters and view backup and restore statistics are:

- View backup parameters on page 64
- View a specific parameter on page 66
- View all parameters on page 67
- Modify parameters on page 69
- View backup statistics on page 70
- Reset backup statistics on page 73

View backup parameters

Action

To view backup parameters and their values, use this command syntax:
`$ server_param <movername> -facility PAX -list`
where:
<movername> = name of the Data Mover
Example:
To view a list of PAX parameters on server_2, type:
`$ server_param server_2 -facility PAX -list`

Output

```
param_name              facility  default  current  configured
dump                    PAX             0        0
allowVLCRestoreToUFS    PAX             0        0
checkUtf8Filenames      PAX             1        1
paxStatBuff             PAX           128      128
readWriteBlockSizeInKB  PAX            64       64
filter.numFileFilter    PAX             5        5
filter.numDirFilter     PAX             5        5
noFileStreams           PAX             0        0
scanOnRestore           PAX             1        1
filter.caseSensitive    PAX             1        1
nRestore                PAX            16       16
writeToArch             PAX             1        1
writeToTape             PAX             1        1
```

Note
where: param_name = name of the parameter facility = facility to which the parameter belongs default = default value for the parameter current = value set on the Data Mover configured = value set that has not taken effect yet because some user action, such as a Data Mover restart, is required.

Note: The configured and current values can differ when you use the server_param command to configure the value of a parameter that is different from the default value. Furthermore, changing the parameter requires user action such as a Data Mover restart to take effect. Parameter and facility names are case-sensitive.

View a specific parameter

Action
To view detailed information on a specific parameter, use this command syntax: `$ server_param {<movername>|ALL} -facility <facility_name> -info <param_name>` where: *<movername>* = name of the Data Mover (or keyword ALL for all Data Movers) *<facility_name>* = name of the facility to which the parameter belongs *<param_name>* = name of the parameter you want to view Example: To view detailed information on the paxStatBuff parameter on server_2, type: **$ server_param server_2 -facility PAX -info paxStatBuff**

Note: Parameter and facility names are case-sensitive.

Output

`name = paxStatBuff`
`facility_name = PAX`
`default_value = 128`
`current_value = 64`
`configured_value = 64`
`user_action = none`
`change_effective = immediate`
`range = (1,512)`
`description = Number of buffers in stat queue per backup session`

Note
The output in this example shows the paxStatBuff parameter value has been increased from the default value to 128. Output definitions: facility_name = facility to which the parameter belongs default_value = value set as the default current_value = value set on the Data Mover configured_value = value configured with the server_param command user_action = action necessary for the parameter change to take effect change_effective = when the parameter change takes effect range = range of possible values

Note
description = description of the parameter (*Parameters Guide for VNX for File* provides more information.)

View all parameters

Action

To view all parameters for a particular Data Mover, use this command syntax:
`$ server_param <movername> -facility <facility_name> -info -all`
where:
<movername> = name of the Data Mover
<facility_name> = name of the facility to which the parameter belongs
Example:
To view all PAX parameters on server_2, type:
`$ server_param server_2 -facility PAX -info -all`

Output

```
name = paxStatBuff
facility_name = PAX
default_value = 128
current_value = 128
configured_value =
user_action = none
change_effective = immediate
range = (1,512)
description = Number of buffers in stat queue per backup session

name = nRestore
facility_name = PAX
default_value = 8
current_value = 8
configured_value =
user_action = none
change_effective = immediate
range = (1,16)
description = Number of worker threads per restore session

name = readWriteBlockSizeInKB
facility_name = PAX
default_value = 64
current_value = 64
configured_value =
user_action = reboot DataMover
change_effective = reboot DataMover
range = (64,256)
description = Maximum allowed PAX buffer size for NDMP read/write in
kilobytes
```

Note: Consult the example presented in View a specific parameter on page 66 to view detailed information on a specific parameter, for an explanation of the categories displayed in this output.

Modify parameters

- Parameters allow you to increase the number of threads and data blocks used during the backup and restore processes. However, more threads and data blocks use more of the Data Mover memory normally dedicated to the backup and restore processes. Therefore, using more threads and data blocks can negatively impact other processes running on the Data Mover.

CAUTION

- Parameter changes made while backups are running on a Data Mover take effect only after any backups on the Data Mover complete.

Action

To modify a parameter, use this command syntax:
`$ server_param {<movername>|ALL} -facility <facility_name> -modify <param_name> -value <new_value>`
where:
<movername> = name of the Data Mover (or keyword ALL for all Data Movers)
<facility_name> = name of the facility to which the parameter belongs
<param_name> = name of the parameter
<new_value> = new value you want to set for the parameter. Numeric values can be given as either a decimal value or in hexadecimal format (0x followed by one to eight hexadecimal characters). To specify a null string, type double quotation marks (“ “).
Note: Parameter and facility names are case-sensitive.
Example:
To set the paxStatBuff parameter to 128, type:
`$ server_param server_2 -facility PAX -modify paxStatBuff -value 128`

Output
`server_2 : done`

View backup statistics

The procedure to view backup statistics describes how to use the server_pax command to view statistics related to your file system data and backup performance.

Note: When you restore data to a location different from its original location, the new path for the data does not appear in the backup statistics. The server_pax command only supports viewing PAX-based backup statistics, not statistics associated with NVBs.

Action

To view statistics on a backup session in progress, use this command syntax:
`$ server_pax <movername> -stats -verbose`
where:
<movername> = name of the Data Mover
Example:
To view statistics for a backup running on server_2, type:
`$ server_pax server_2 -stats -verbose`

The output is grouped by thread group, which performs various functions within the multithreaded backup and restore process. Multithreaded backup and restore design on page 42 provides more information.

Output

```
************** SUMMARY PAX STATS ****************
---- NASS STATS ----
** nass thid 0 **
Total file processed: 16
throughput: 2 files/sec
Total nass wait nasa count: 13
Total nass wait nasa time: 236 msec
Total time since last reset: 8 sec
fts_build time: 0 sec
getstatpool: 0 buffersputstatpool: 128 buffers

---- NASA STATS ----
** nasa thid 0 (BACKUP) **
Backup root directory: /s4fs1/adm
Total bytes processed: 104937747
Total file processed: 16
throughput: 11 MB/sec
average file size: 6404KB
Total nasa wait nass count: 1
Total time since last reset: 8 sec
Total nasa wait nass time: 13 msec
Tape device name: /dev/c0t5l0
0 size file processed: 3
1 -- 8KB size file processed: 8
8KB+1 -- 16KB size file processed: 0
16KB+1 -- 32KB size file processed: 0
32KB+1 -- 64KB size file processed: 1
64KB+1 -- 1MB size file processed: 2
1MB+1 -- 32MB size file processed: 1
32MB+1 -- 1GB size file processed: 1
1G more size file processed: 0

---- NASW STATS ----
nasw00 BACKUP (in progress)
Session Total Time: 00:00:08 (h:min:sec)
Session Idle Time: 00:00:00 (h:min:sec)
KB Tranferred: 98406 Block Size: 64512 (63 KB)
Average Transfer Rate: 12 MB/Sec 42 GB/Hour
Average Burst Transfer: 12 MB/Sec 42 GB/Hour
Write Block Counters: 1562/0 (List/Direct)
__Point-in-Time__ (over the last 10 seconds):
Rate=8 MB/Sec Burst=11 MB/Sec Idle=11 msec/sec
Get Pool: 63 buffers Put Pool: 0 buffers
Compression page retrieved 00:00:08 (h:min:sec) ago:
ReadC=2.1 WriteC=2.1 Read=310692 KB Written=408327 KB
```

Note
When you specify verbose mode, information on file size distribution is included in the output. where: NASS STATS = thread responsible for traversing the file system and providing metadata for each directory and/or file Total file processed = total number of files and/or directories for which metadata was processed Total NASS wait NASA count = number of times NASS waited for NASA Total NASS wait NASA time = amount of time NASS waited for NASA Total time since last reset = time since the last reset; a reset occurs automatically after a backup completes fts_build time = time spent building the file system or directory tree getstatpool = if the value is consistently 0, then NASA could be slowing down the backup putstatpool = if the value is consistently 0, then NASS could be slowing down the backup NASA STATS = thread responsible for writing file header information, reading file data, and writing to the buffer Backup root directory = directory being backed up Total bytes processed = bytes backed up since last reset or start of current backup Total file processed = number of files backed up since start or reset of current backup Throughput = how fast NASA processed data Average file size = average size of file for current backup Total nasa wait nass count time = number of times NASA waited for NASS Total nasa wait nass time = amount of time NASA waited for NASS Total time since last reset = amount of time since the backup statistics were reset; a reset occurs automatically after a backup completes Tape device name = target device for the backup data File size statistics = statistics on the size of files backed up since the start or reset of the current backup NASW STATS = thread responsible for getting data from the buffer pool, writing it to tape, or sending it to a remote Data Mover Session total time = total time of current session Session idle time = idle time for current session

Note
KB transferred = total KB transferred Average transfer rate = per-second and per-hour transfer rate for current session's data Average burst transfer = burst transfer rate in MB/s and GB/h Write block counters (List/Direct) = scatter/gather write count _Point-in-time_ (over the last 10 seconds) = information on data processed during a 10 second point in time Rate = transfer rate in MB/s Burst = burst transfer rate in MB/s Idle = amount of time NASW was idle in msec Get pool = number of buffers in get pool; if the value is consistently 0, then NASA and NASS could be slowing down the backup Put pool = number of buffers in put pool; if the value is consistently 0, then the tape could be slowing down the backup Compression rate retrieved = compression rate ReadC = read compression rate at the tape device WriteC = write compression rate at the tape device Read = amount of data read in KB Written = amount of data written in KB

Reset backup statistics

Statistics are compiled for each backup session and automatically reset when a backup completes.

Action

To reset statistics for a backup session in progress, use this command syntax:
`$ server_pax <movername> -reset`
where:
<movername> = name of the Data Mover
Example:
To reset the statistics for a backup running on server_2, type:
`$ server_pax server_2 -reset`

Output
`server_2 : done`

Back up a VDM file system with NDMP

This procedure explains only the NDMP-specific guidelines for Virtual Data Mover (VDM). *Configuring Virtual Data Movers on VNX* provides more information on VDM.

A VDM enables you to administratively separate CIFS servers and their associated resources, like file systems, into virtual containers. These virtual containers allow administrative separation between groups of CIFS servers, enable replication of CIFS environments, and allow the movement of CIFS servers from Data Mover to Data Mover. In a CIFS environment, when you use NDMP to back up a VDM configured file system, you must specify the full mount path of the file system.

Perform the following steps to backup a VDM file system with NDMP:

1. To retrieve the full mount path of a file system, use this command syntax:

   `$ server_mount <movername> | grep <fs_name>`

   where:

   <movername> = name of the Data Mover

   <fs_name> = name of the file system for which you want the full path

   Example:

   To retrieve the full mount path of file system ufs1, type:

   ```
   $ server_mount server_2 | grep ufs1
   ```

   Output:

   ```
   ufs1 on /root_vdm_1/ufs1 uxfs,perm,rw
   ```

   Note: /root_vdm_1, in this example, indicates that the file system is a VDM-configured file system. To back up this file system by using NDMP, type /root_vdm_1/ufs1 as the full path in the backup software.

2. In the NDMP backup software, type the pathname revealed in step 1. For example:

   ```
   /root_vdm_1/ufs1
   ```

   Note: Your backup software vendor's documentation provides more information on specifying file system pathnames in your NDMP backup software.

Back up a FileMover file system with NDMP

To include the file content on secondary storage, set this NDMP environment variable in your backup software:

```
EMC_OFFLINE_DATA=y
```

Note: When you use the default setting EMC_OFFLINE_DATA=n, the content on secondary storage is not backed up.

Restore a FileMover file system

When you restore a FileMover file system that includes content from the primary and secondary storage systems, the behavior is the same as with any other file system restore: the entire file system is restored as online files. To restore a file system as online content, the file system must be backed up with the environment variable EMC_OFFLINE_DATA=y set in the backup software.

When an NDMP backup is performed with the environment variable EMC_OFFLINE_DATA=n, which is the default configuration, only the primary storage content is backed up. When you restore such a file system, the stub files are restored, and they continue to serve as pointers to access the file content on secondary storage. The file content remains on secondary storage as long as its associated policy dictates. *Using VNX FileMover* provides more information on the FileMover feature.

Back up a file system with NDMP by using an integrated SnapSure checkpoint

To integrate a checkpoint into the NDMP backup process and automate the checkpoint creation, management, and deletion tasks upon initiation of backups, set this NDMP environment variable in your backup software:

```
SNAPSURE=y
```

Note:

- If you use the default setting SNAPSURE=n in the backup software, no automated checkpoint activity occurs.
- If the backup software does not support the SNAPSURE=y environmental variable, you can set the parameter NDMP.snapsure in the Data Mover to use VNX SnapSure for backup. The default parameter value 0 should be changed to 1.

Perform backup and restore with NVB

Backup/restore with NVB on page 29 provides information about backing up and restoring volumes with NVB:

- Specify NVB in backup vendor software on page 76 provides information about specifying NVB in vendor backup software.

- Perform NVB restores on page 77 provides information about performing NVB restores.

Specify NVB in backup vendor software

Some backup software vendors support the use of backup types, others support the use of environmental variables, and some vendors support both. Therefore, the action required to initiate a VBB depends on your backup software vendor. The EMC E-Lab Interoperability Navigator provides a current list of EMC-qualified backup software vendors for NVB and their supported implementation methods.

This section shows an example of implementing NVB by specifying VBB as the backup type and VBB as an environmental variable in NetWorker and Symantec NetBackup. Your backup vendor documentation provides instructions on implementing backup types and environmental variables.

Note: Always ensure that the HIST=y and UPDATE=y environmental variables are set to create the index required for future direct-access restore operations. Do not use file and directory filters with NVB. They will have no effect.

Specify NVB in EMC NetWorker

Action
To set the backup type, in the Client Properties window, at the **Application information** field, type: **`nsrndmp_save -Type vbb`** Note: In the NetWorker software, the backup type, VBB, must be typed in lowercase. Otherwise the backup immediately fails. Or use this environmental variable: VBB=y

Specify NVB in Symantec NetBackup

Action
To set the backup type, type: **`SET type=vbb`** Or use this environmental variable: **`SET VBB=y`**

Perform NVB restores

The types of restore that NVB supports are:

- Full destructive restore
- File-level restore

Perform full destructive restore on page 77 and File-level restore on page 85 provide more information.

Perform full destructive restore

When performing a full destructive restore, the NVB feature requires the target file system be of equal size or greater than the file system backed up.

When you use the backup image created by using NVB to perform a full destructive restore, the restore occurs at the volume level. Large sequential writes are issued to the volume as data blocks are read from tape and written directly to the specified volume. All existing data on the target file system is destroyed. This type of restore ensures the best possible data block write performance.

Note: Perform a full destructive restore of a VDM file system on page 85 describes the procedure to do a full destructive restore of a VDM file system.

Perform the following steps for a full destructive restore:

1. To complete a full NVB of a PFS by using the NVB backup mechanism in the backup-vendor software, type the following. In this example, the PFS is ufs1:

   ```
   $ nas_fs -s ufs1
   ```

   Output:

   ```
   total = 73855 avail = 69133 used = 4721 ( 6%) (sizes in MB)
   (blockcount = 153600000)
   volume: total = 75000 (sizes in MB) (blockcount = 153600000)
   ```

2. To use Automatic Volume Management (AVM) in order to create the file system needed for the full destructive restore, type this command to get the pool information of the PFS, which is needed in the process:

   ```
   $ nas_fs -i ufs1
   ```

   Output:

```
id = 1150
name = ufs1
acl = 0
in_use = True
type = uxfs
worm = off
volume = v1772
pool = clar_r5_performance
member_of = root_avm_fs_group_3
rw_servers= server_2
ro_servers=
rw_vdms =
ro_vdms =
auto_ext = no,virtual_provision=no
stor_devs = APM00043306037-000C,APM00043306037-
000F,APM00043306037-0008,APM00043
306037-000B
disks = d19,d11,d17,d9
disk=d19 stor_dev=APM00043306037-000C addr=c16t1l7
server=server_2
disk=d19 stor_dev=APM00043306037-000C addr=c32t1l7
server=server_2
disk=d19 stor_dev=APM00043306037-000C addr=c0t1l7
server=server_2
disk=d19 stor_dev=APM00043306037-000C addr=c48t1l7
server=server_2
disk=d11 stor_dev=APM00043306037-000F addr=c0t1l10
server=server_2
disk=d11 stor_dev=APM00043306037-000F addr=c48t1l10
server=server_2
disk=d11 stor_dev=APM00043306037-000F addr=c16t1l10
server=server_2
disk=d11 stor_dev=APM00043306037-000F addr=c32t1l10
server=server_2
disk=d17 stor_dev=APM00043306037-0008 addr=c16t1l3
server=server_2
disk=d17 stor_dev=APM00043306037-0008 addr=c32t1l3
server=server_2
disk=d17 stor_dev=APM00043306037-0008 addr=c0t1l3
server=server_2
disk=d17 stor_dev=APM00043306037-0008 addr=c48t1l3
server=server_2
disk=d9 stor_dev=APM00043306037-000B addr=c0t1l6
server=server_2
disk=d9 stor_dev=APM00043306037-000B addr=c48t1l6
server=server_2
disk=d9 stor_dev=APM00043306037-000B addr=c16t1l6
server=server_2
disk=d9 stor_dev=APM00043306037-000B addr=c32t1l6
server=server_2
```

Note: In this output example, the pool is clar_r5_performance.

3. To create a raw file system (rawfs) to be used for the NVB full destructive restore (FDR), type:

```
$ nas_fs -name newufs1 -type rawfs -create samesize=ufs1 pool= clar_r5_performance
```

Note: The raw file system should be the same size as the PFS. The easiest method to ensure equivalency is to use the samesize option of the nas_fs command. If you do not use the samesize option, ensure that the rawfs is at least the same size as the PFS.

Output:

```
id = 1151
name = newufs1
acl = 0
in_use = False
type = rawfs
worm = off
volume = v1775
pool = clar_r5_performance
member_of = root_avm_fs_group_3
rw_servers=
ro_servers=
rw_vdms =
ro_vdms =
auto_ext = no,virtual_provision=no
stor_devs = APM00043306037-000E,APM00043306037-
000D,APM00043306037-000A,APM00043
306037-0009
disks = d20,d10,d18,d8
```

4. To create the mount point for the raw file system, type:

```
$ server_mountpoint server_2 -c /newufs1
```

Output:

```
server_2 : done
```

5. To mount the raw file system, type:

```
$ server_mount server_2 newufs1 /newufs1
```

Output:

```
server_2 : done
```

Note: When you list the file system, it appears as unmounted, as shown below. This is expected behavior.

```
$ server_mount server_2
```

Output:

```
server_2 :
root_fs_common on /.etc_common uxfs,perm,ro
ufs1 on /ufs1 uxfs,perm,rw
newufs1 on /newufs1 rawfs,perm,rw,<unmounted>
```

6. To record the volume ID of the raw file system, type:

```
$ nas_fs -i newufs1
```

Output:

```
id = 1151
name = newufs1
acl = 0
in_use = False
type = rawfs
worm = off
volume = v1775
pool = clar_r5_performance
member_of = root_avm_fs_group_3
rw_servers=
ro_servers=
rw_vdms =
ro_vdms =
auto_ext = no,virtual_provision=no
stor_devs = APM00043306037-000E,APM00043306037-
00D,APM00043306037-000A,APM000430
306037-0009
disks = d20,d10,d18,d8
```

The volume ID in this output example is v1775. This ID without the leading "v" is required for the restore.

Note: The next step depends on the backup software you are using. For NetWorker, perform steps 7 through 9. For Symantec NetBackup, perform steps 10 through 12. After completing the required step, proceed to step 13.

---

7. NetWorker procedure:

   If you are performing a full destructive restore by using NetWorker software, perform steps 7 through 9.

   In UNIX:

   a. From the **NetWorker Recover** window, click the **Options** drop-down list and select **Recover SaveSets**.

   b. Select the PFS to be recovered and click **Start Recover Operations**.

   In Windows:

   a. From the **NetWorker Recover Client** menu, select **Save Set Recovery**. The **Save Sets** window appears.

b. Select the PFS to be recovered and click **Recover Options**.

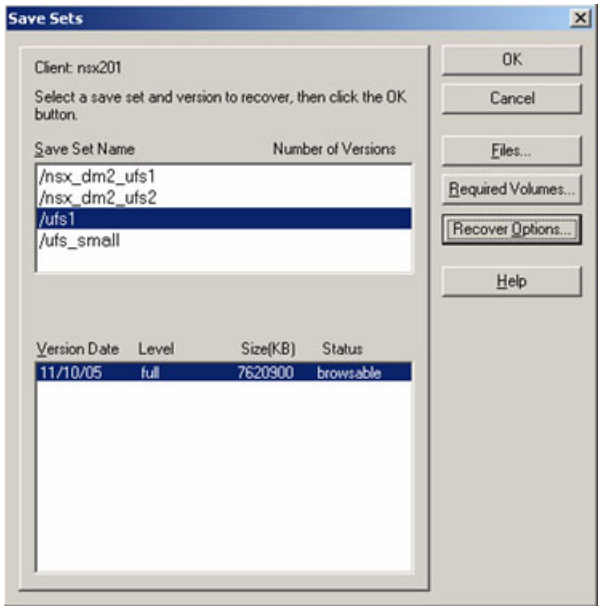

8. In the Relocate recovered data to this path field, specify the rawfs path by using this command syntax:

   `/.celerra_vol_ID`

   where:

   `ID` = the volume ID of the rawfs

   Example:

   To specify the rawfs created earlier, type:

   `/.celerra_vol_1775`

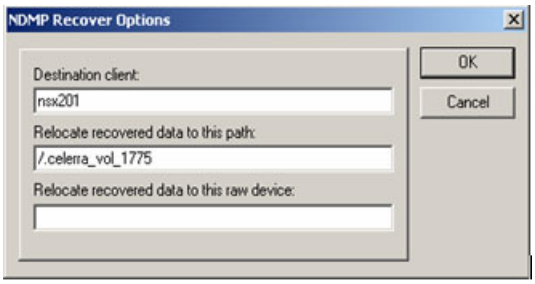

9. When the restore is complete, the client displays the message:

   `nsmdmp_recover: Successfully done`

10. Symantec NetBackup procedure:

If you are performing a full destructive restore by using Symantec NetBackup software, perform steps 10 through 12.

From the **Backup, Archive, and Restore - NetBackup** window, select the file system to restore. Select the checkbox at the left of the file system name.

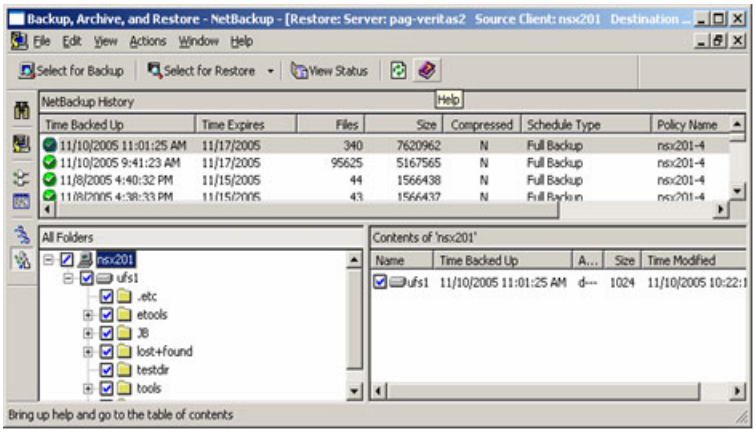

11. From the **Restore Marked Files** window, select **Restore everything to a different location**.

    Type **`/.celerra_vol_ID`**, where ID is the volume ID, without the leading "v" of the rawfs created earlier, and click **Start Restore**.

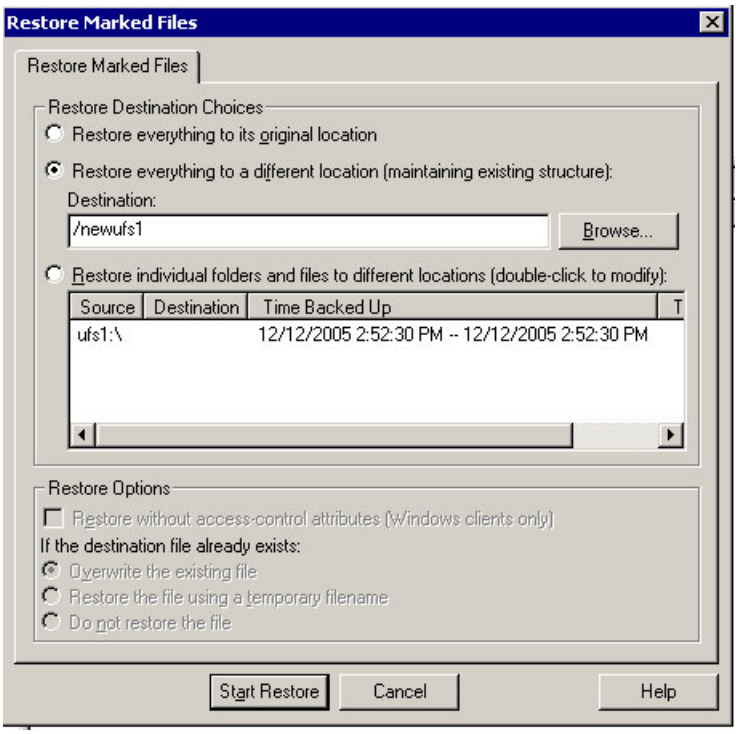

12. When the restore completes, the **View Status** window displays this information.

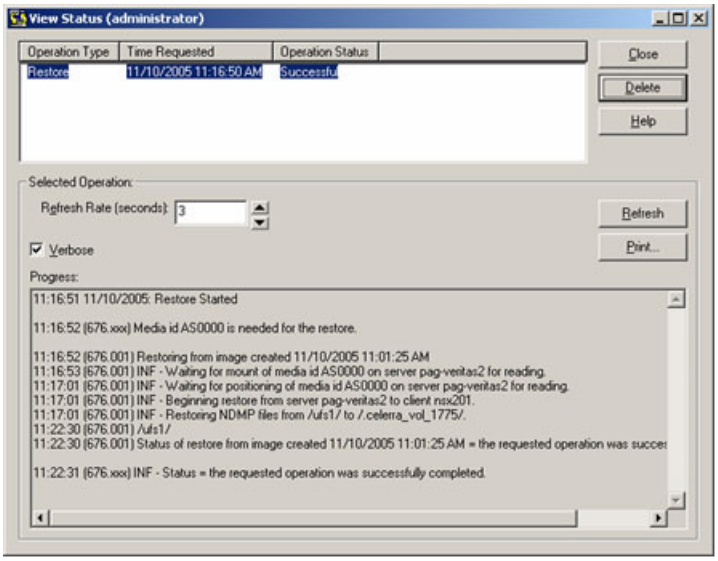

13. To convert the file system from a rawfs to a uxfs, when the restore is complete, type:

```
$ nas_fs -Type uxfs newufs1 -Force
```

Output:

```
id = 1151
name = newufs1
acl = 0
in_use = True
type = uxfs
worm = off
volume = v1775
pool = clar_r5_performance
member_of = root_avm_fs_group_3
rw_servers= server_2
ro_servers=
rw_vdms =
ro_vdms =
auto_ext = no,virtual_provision=no
stor_devs = APM00043306037-000E,APM00043306037-
000D,APM00043306037-000A,APM00043
306037-0009
disks = d20,d10,d18,d8
disk=d20 stor_dev=APM00043306037-000E addr=c16t1l9
server=server_2
disk=d20 stor_dev=APM00043306037-000E addr=c32t1l9
server=server_2
disk=d20 stor_dev=APM00043306037-000E addr=c0t1l9
server=server_2
disk=d20 stor_dev=APM00043306037-000E addr=c48t1l9
server=server_2
disk=d10 stor_dev=APM00043306037-000D addr=c0t1l8
server=server_2
disk=d10 stor_dev=APM00043306037-000D addr=c48t1l8
server=server_2
disk=d10 stor_dev=APM00043306037-000D addr=c16t1l8
server=server_2
disk=d10 stor_dev=APM00043306037-000D addr=c32t1l8
server=server_2
disk=d18 stor_dev=APM00043306037-000A addr=c16t1l5
server=server_2
disk=d18 stor_dev=APM00043306037-000A addr=c32t1l5
server=server_2
disk=d18 stor_dev=APM00043306037-000A addr=c0t1l5
server=server_2
disk=d18 stor_dev=APM00043306037-000A addr=c48t1l5
server=server_2
disk=d8 stor_dev=APM00043306037-0009 addr=c0t1l4
server=server_2
disk=d8 stor_dev=APM00043306037-0009 addr=c48t1l4
server=server_2
disk=d8 stor_dev=APM00043306037-0009 addr=c16t1l4
server=server_2
disk=d8 stor_dev=APM00043306037-0009 addr=c32t1l4
server=server_2
```

14. The file system remains mounted on the Data Mover, but it is now automatically mounted read/write, as follows:

    `$ server_mount server_2`

    Output:

```
server_2 :
root_fs_common on /.etc_common uxfs,perm,ro
ufs1 on /ufs1 uxfs,perm,rw
newufs1 on /newufs1 uxfs,perm,rw
```

15. You can verify the contents of the new file system from a CIFS or NFS client. To display the size and usage of the new file system from the Control Station, type:

    ```
    $ nas_fs -s newufs1
    ```

    Output:

    ```
    total = 73855
    avail = 63501 used = 10354 (14%) (sizes in MB) (blockcount = 153600000)
     volume:
    total = 75000 (sizes in MB) (blockcount = 153600000)
    ```

    Note: After the restore is done, the size of the rawfs will appear the same as the PFS. To view the actual raw file system size, extend it by a minimum of 1 MB, and then list the raw file system.

Perform a full destructive restore of a VDM file system

Following are steps:

1. Create a raw file system the same size as the VDM PFS.
2. Mount the rawfs on the physical Data Mover, that is, server_2, not the VDM.
3. Perform a FDR by using the same directions as for a non-VDM file system, described in Perform full destructive restore on page 77.
4. Permanently unmount the rawfs from the physical Data Mover.
5. Create a mount point on the VDM.
6. Mount the rawfs on the VDM.
7. Convert the rawfs to a uxfs.

File-level restore

The restore time for performing a file-level restore by using an NVB is comparable to the time required to do the same by using a PAX backup.

The skipDedupFiles parameter of the vbb facility controls whether to skip the deduplicated files that are encountered during a file-level restore, thereby restoring at least the non-deduplicated files. If this parameter is set to 0, the restore is aborted when a deduplicated file is encountered. If it is set to 1, an error is logged and the restore process skips all

deduplicated files and restores any subsequent non-deduplicated files. The *Parameters Guide for VNX for File* provides more information.

Temporary space requirements for a file-level restore

VNX requires temporary space to process metadata during the file-level restore process when an NVB serves as the backup source. The space needed is usually less than 2.5 percent of the overall backup image. By default, the restore target file system is used for the temporary space. Ensure that the file system has sufficient space for the temporary metadata before initiating the file-level restore.

VNX names the temporary space in the root of the file system.vbbtemp. It reads metadata about all the files from the tape into this space, creates and restores the requested files to the appropriate locations, and then cleans up the space when the restore completes. If the target file system becomes full during the restore, an error results.

Note: The NVB and the NVB full destructive restore processes do not require temporary space to process metadata. Files are guaranteed to be fully restored only when the restore process is complete. During the restore process, a file can exist, but all of its data may not be restored unlike PAX, which restores entire files at once. When an NVB is the backup source, entire trees are restored. The completeness of any file is not guaranteed until the entire restore process has finished.

Exclude directories or files from backup by using filtering

Filter to exclude directories or files from backup on page 32 provides information about excluding specific files or directories from being backed up. Exclude a directory on page 86 describes the procedure to exclude a directory from backup, and Exclude a file on page 88 describes the procedure to exclude a file from backup.

Note: The filtering feature is supported only by NDMP PAX (dump or tar) backup types. Do not use file and directory filtering with the NVB.

Exclude a directory

You can increase the number of directory filters, up to 50, by using the filter.numDirFilter parameter. Modify parameters on page 69 provides more information.

If you are using the VNX internationalization feature with a character set other than UTF-8, use the filter.dialect parameter to translate the pattern received to UTF- 8. The default value of the parameter is ''. After the pattern is converted to a UTF-8 pattern, the valid path rules are applied, and if the criteria are met, the pattern is accepted.

Action
To exclude a specific directory from the NDMP backup, type:

Action
EMC_EDIR*<nn>=/<valid directory pattern>* where: *<nn>* = an ASCII number from 01 to 50 that helps you keep track of the total number of variable entries. By default, five variable entries are allowed per NDMP backup session for directory filtering. *<valid directory pattern>* = a pattern that identifies the specific type of directory to be excluded from the backup. Example: To exclude the directory ufs1/dir1 from the NDMP backup, type: **EMC_EDIR01=/ufs1/dir1** To exclude all directories that begin with "d ? under /ufs2, type: **EMC_EDIR02=/ufs2/d***

Exclude a file

You can increase the number of directory filters, up to 50, by using the filter.numDirFilter parameter. Modify parameters on page 69 provides more information.

If you are using the VNX internationalization feature with a character set other than UTF-8, use the filter.dialect parameter to translate the pattern received to UTF-8. The default value of the parameter is ''. After the pattern is converted to a UTF-8 pattern, the valid path rules are applied, and if the criteria are met, the pattern is accepted.

Action
To exclude a specific file from the NDMP backup, type: **EMC_EFILE***<nn>*= *<valid file pattern>* where: *<nn>* = an ASCII number from 01 to 50 that helps you keep track of the total number of variable entries. By default, five variable entries are allowed per NDMP backup session for file filtering. *<valid file pattern>* = a pattern that identifies the specific type of file to be excluded from the backup. Example: To exclude all MP3 files from the NDMP backup, type: **EMC_EFILE01=*.mp3** To exclude all files that begin with "temp ? from the backup, type: **EMC_EFILE02=temp*** To exclude all WMA files from the backup, type: **EMC_EFILE03=*.wma**

Perform tape silvering

Use tape silvering to restore the backup that is created on the source site by using the following software:

- Restore with NetBackup on page 89
- Restore with CommVault Galaxy on page 93
- Restore with NetWorker on page 95

Restore with NetBackup

Following are steps:

1. Physically place the tape volumes onto the destination TLU and then inventory the TLU to add new tape volumes to the NetBackup unit:
   a. Start the **NetBackup Administration Console**. In a Windows environment, click **Start**, and select **Programs ➤ Symantec NetBackup ➤ NetBackup Administration Control**.
   b. From the left panel, select **Devices** and then click **Robot**.
   c. From the right panel, right-click the appropriate robot and then select **Inventory Robot** from the menu. The **Robot Inventory** window appears.
   d. Under **Inventory Operation**, select **Update volume configuration** and then click **Start**.
   e. Under **Results**, when the message **Volume configuration successfully updated** appears, click **Close**.
   f. From the left panel, click **Media** and verify that the tapes are added in the media database.
2. Perform a Phase 1 import to scan the tape volumes containing the backup image and read the backup headers:
   a. Select **NetBackup Management** and then click **Catalog**.
   b. Click the **Search** tab and then select **Import** from the **Action** list box.
   c. From the **Actions** menu, click **Initiate Import**. The **Initialize Import** dialog box appears.
   d. Type the required information (for example, the tape volume ID of the first tape containing the backup image) in the **Media ID** field, and click **OK**. The **Confirm Initiate Import** window appears.

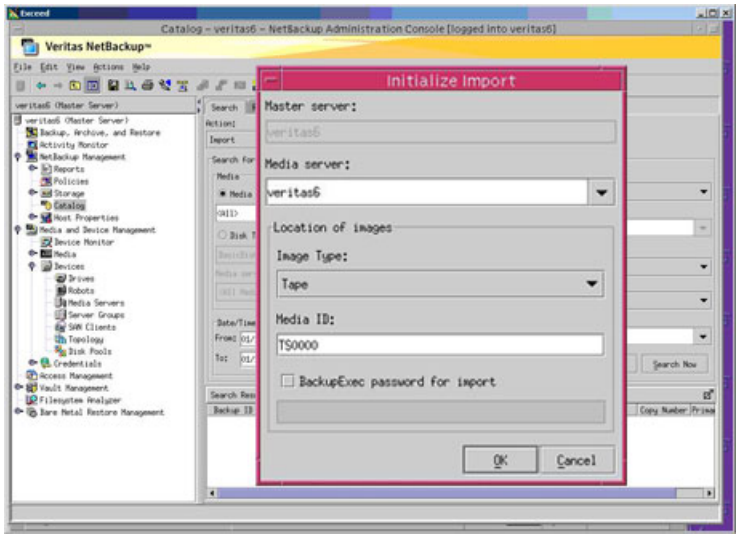

e. Click **OK** to confirm the initialization process.

f. Click the **Results** tab to monitor the import process until it is completed successfully. Continue the import process until all the tape volumes are imported successfully.

3. Verify the backup images created in Phase 1 import:

   a. Click the **Search** tab on the right panel, and then select **Import** from the **Action** list box.

   b. From the **Policy type** list box, click **NDMP** and then click **Search Now**.

   c. When the search is complete, the backup images are listed in the **Search Results** area. Select the required backup image to perform Phase 2 import.

4. Perform a Phase 2 import to update the catalog with the required information:

   a. From the **Actions** menu, click **Import**. The **Confirm Import** window appears.

   b. Click **OK** to confirm the import.

   c. Click the **Results** tab and monitor the import process until it has completed successfully.

Note: Follow steps 5 through 10 to perform the restore operation.

---

5. To perform the restore operation, from the left panel, select **Backup, Archive and Restore**.

6. From the right panel, click the **Restore Files** tab. The **Specify NetBackup Machines and Policy Type** window appears:

   a. In the **Source client for restores** list box, select the original server as the source client for the restore operation. If the source client does not appear in the client list, click **Edit Client List** and enter the source client name.

b. In the **Destination client for restores** list box, select the required destination server for the restore operation.

c. In the **Policy type for restores** list box, select **NDMP**. Click **OK** to close the **Specify Net Backup Machines and Policy Type** window.

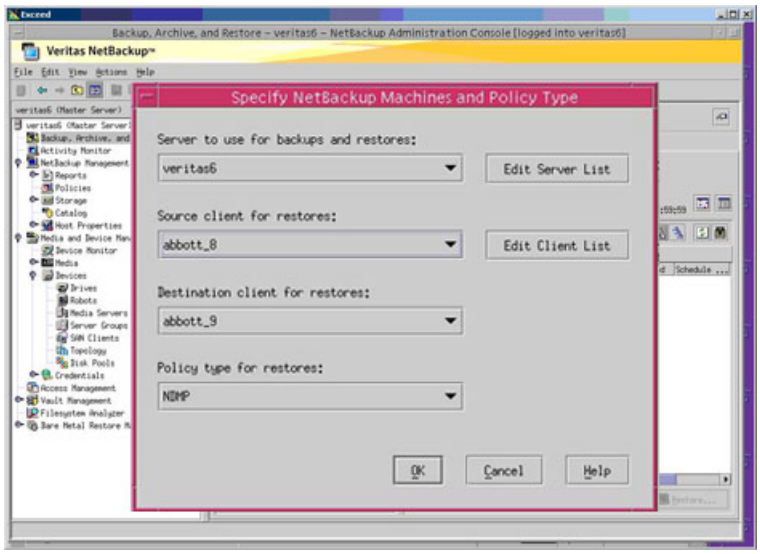

7. Under the **Directory Structure**, select the backup image to be restored, and then click **Restore**. The **Restore Marked Files** window appears.

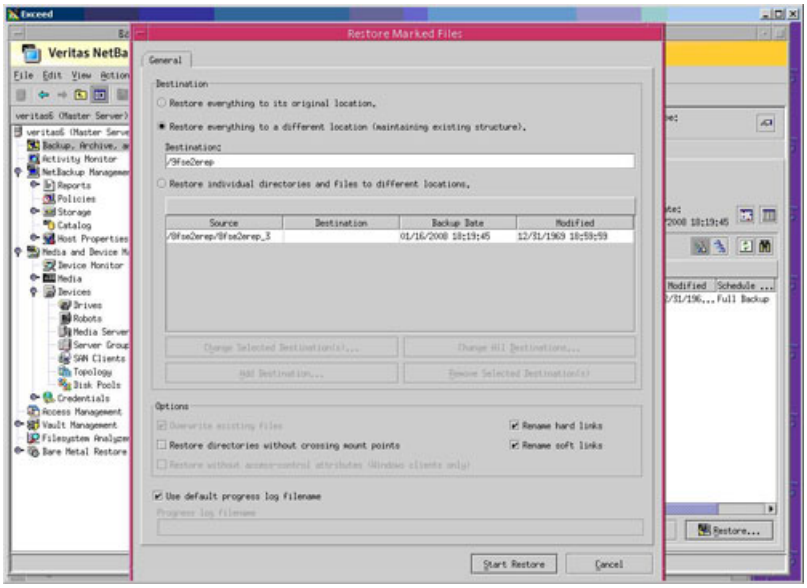

8. Select the **Restore everything to a different location (maintaining existing structure)** option. In the **Destination** field, type the appropriate restore location as the destination file system for replication.

9. Click **Start Restore** to start the restore operation. The **View Progress** window appears. Click **OK** to view the progress of the restore operation.

10. From the left panel, select **Activity Monitor** and then double-click the listed job in the right panel. Click **Close** in the **Job Details** window, when the restore operation is complete.

Restore with CommVault Galaxy

Perform the following steps to restore with CommVault Galaxy:

1. Create a subclient with the path:

   /<mountpoint>/.ts

   where:

   <mountpoint> = path to mount the file system on VNX.

   Specify the mount point for the source file system in the backup definition with the postfix /.ts. For example, if the primary file system is mounted on mountpoint replication_pfs, the backup definition appears as:

   Backup file system: /replication_pfs/.ts

   Note: The .ts is the default value for the server_param NDMP tapeSilveringStr. You can change this value at any time. The source file system and checkpoints must be mounted on the NDMP Data Mover.

2. To create a file system replication session with the tape_copy option, use this command syntax:

   $ **nas_replicate -create** *<name>* **-source -fs** *<fsName>* **-destination -pool** *<dstStoragePool>* **-interconnect** *<name>* **-overwrite_destination -tape_copy**

   where:

   <name> = name of the replication session

   <fsName> = name of the source file system

   <dstStoragePool> = name of the destination storage pool

   Example:

   $ nas_replicate -create e2_g2_2 -source -fs newfs -destination -pool clar_r5_economy -interconnect E2_G2 -overwrite_destination -tape_copy

   Output:

   OK

3. Run a full backup of the newly created subclient.

4. To capture the catalog information of the client for tape silvering, type:

   C:\Program Files\CommVault\Galaxy\Base> qoperation capture -c *<client name>* **-dfn** *<db file name>* **-af n** *<answer file name>*

   Example:

```
C:\Program Files\CommVault\Galaxy\Base>qoperation capture -c eg2_DS  -dfn
c:\Galaxy_CCM\eg2_db -afn c:\Galaxy_CCM\eg2_afile_capture
```

Output:

```
Capture successful
```

5. Copy the captured files to the destination site and then modify the answer file by entering the required information.

6. To merge the catalog information of the client in the Galaxy catalog at the destination site, type:

   C:\Program Files\CommVault\Galaxy\Base>qoperation merge -dfn *<db filename>* **-afn** *<modified answer file name>*

   Example:

   ```
   C:\Program Files\CommVault\Galaxy\Base>qoperation merge -dfn C:\Galaxy_CCM\eg2_db
   -afn C:\Galaxy_CCM\eg2_afile_merge
   ```

   Output:

   ```
   Merge successful
   ```

7. Export media from the source library.

8. Import media in the destination library.

9. Browse the tape silvering backup in the destination setup to initiate the index retrieval job. When the job is complete, you can view the backup.

10. Select the <fs_#> folder for restore and specify target server and path.

11. To start replication, when the restore process is complete, type:

    $ **nas_replicate -start** *<name>* **-overwrite_destination**

    where:

    <name> = name of the replication session

    Example:

    ```
    $ nas_replicate -start e2_g2_2 -overwrite_destination
    ```

    Output:

    ```
    OK
    ```

Restore with NetWorker

Perform the following steps to restore with NetWorker :

1. Ensure that there is a connection between the source and destination VNX systems.

   On the source VNX Control Station, type:

   ```
   nas_cel -name LA -create -ip 228.222.200.200 -passphrase pass1
   ```

   On the destination VNX Control Station, source to root, type:

   ```
   nas_cel -name NY -create -ip 192.168.100.100 -passphrase pass1
   ```

   Note: Ensure that the passphrase is the same on the source and destination sites and that it is 6 to 25 characters in length.

   Output:

   ```
                Source              Destination
   Celerra      NY                  LA
   IP address   192.168.100.100     228.222.200.200
   ```

2. Create a Data Mover Interconnect between the source and destination Data Movers.

   For example, if the source file system is mounted on server_2 of the source VNX, and the destination file system is mounted on server_3 of the destination VNX, create a Data Mover Interconnect from server_2 of NY to server_3 of LA, and vice versa.

   On the source Control Station, type:

   ```
   nas_cel -interconnect -create NYs2-LAs3 -source _server server_2
   destination_system LA -destination _server server_3 -source  _interfaces ip=
   228.222.200.200 -destination _interfaces
   ip= 192.168.100.100
   ```

   Create a corresponding return interconnect from LA to NY :

   ```
   nas_cel -interconnect -createLAs3 -Nys2 -source_server  server_3 -destination
   _system NY -destination _server server_2 -source_interfaces ip= 192.168.100.100
   ```

3. Start the replication session by using the tape_copy option.

   Note: Ensure that you start the replication session by using the tape_copy option for the source file system, so that a common base checkpoint of the file system is backed up to tape. If it is not selected, the backup tape will not have the necessary objects to complete the tape silvering process and it will fail.

   If you are using the Unisphere software to create a new replication session, ensure that you select the tape_copy option. If you use CLI, type:

   ```
   nas_replicate -create NY_src_fs_to_LA -source -fs NY_fs -destination -pool
   clalr_r5_performance -interconnect NYs2-Las3 -tape_copy
   ```

4. To verify that the replication session is configured, type:

   ```
   # nas_replicate -list
   ```

   Output:

   ```
   Name      Type  LocalMover Interconnect  Celerra Status
   NY_src_fs file  server_2   ->NYs2-LAs2   LA      Info 26045317
   _to_LA    system                                 7429: stopped
   ```

5. Back up the source file system by using tape silvering method. Configure EMC NetWorker on page 124 provides the procedure to set up NetWorker:

   a. In the NetWorker application, create an NDMP client for NDMP backup.

   b. Click the **Apps and Modules** tab, and then type the username and password.

   c. Type the following backup command:

      ```
      nsrndmp_save -T ts
      ```

   d. Application information:

      - HIST=y
      - UPDATE=y
      - SNAPSURE=y (optional)
      - TS=y

      Note: Make sure that the NDMP environment variable for tape silvering is set to TS=y.

   e. Start the backup operation. Verify that the `Starting Tape Silvering (VSB) backup. [IsVsb]` message appears in the server_log of the Data Mover running the backup sroerver. All server_log messages will have the VSB session id in the log messages from here on.

   f. To verify that the backup is completed and savesets are created, run the mminfo command at command prompt.

      Example:

      ```
      C:\>mminfo -avV
      ```

      Output:

      ```
      volume      client     size level name
      ssid     save time    date     time    browse  clretent
      first     last file  rec volid     total fl
      000057L2     server_2    440 GB  full /NY_fs
      3785502540  1201811201  1/31/2008 3:26:41 PM 2/29/2008 1/31/2009
      1048283320320 1499348067764 6  0 4171314420 1499348067765 tbN
      000061L2     server_2    1023 GB  full /NY_fs
      3785502540  1201811201    1/31/2008 3:26:41 PM 2/29/2008 1/31/2009
      0 1048283320319 3    0 4188091594 1499348067765 hbN
      bootstrap       eng252222   194 KB  full bootstrap
      3735193758  1201834142    1/31/2008 9:49:02 PM 2/29/2008 1/31/2009
      0     198399 3735193758 0 10553920    198400 cr
      ```

From the output the two tape volumes, 000061L2 and 000057L2, were used in the backup process, where:

000061L2 = header browse able NDMP (hbn) tape to mount on the tape volume

000057L2 = tail browse able NDMP (tbn) tape to mount on the tape volume

6. Transport the tapes to the destination site.

7. Merge the database of the backup application. Mount the tapes onto the tape library unit of the destination data center. The DMA server at the destination site does not have the catalog for these new tapes. Scan the tape and read the catalog information.

   Physically mount the backup tapes (delivered from the source site) onto the destination NDMP tape library unit.

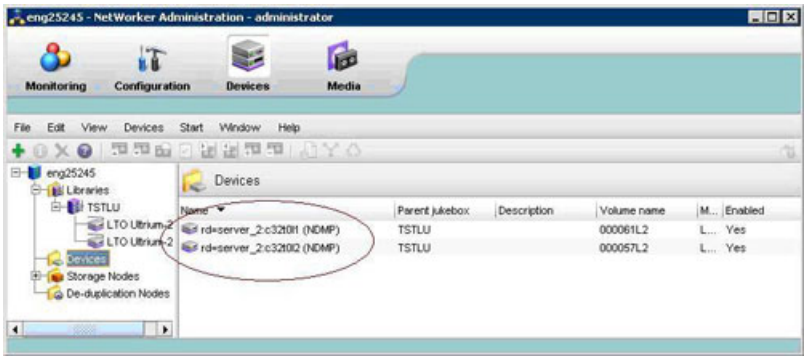

8. From the destination DMA server prompt, run the NetWorker scanner command:

   `scanner -m "rd=server_2:c32t0l1 (NDMP)"`

   where:

   `server_2:c32t0l1 (NDMP)` = the NDMP TLU that is configured

   Output:

```
C:\>scanner -m "rd=server_2:c32t0l1 (NDMP)"
8909:scanner: using 'rd=server_2:c32t0l1 (NDMP)' as the device name
NDMP Service Debug: The process id for NDMP service is 0xdefe4a9c
8936:scanner: scanning LTO Ultrium-2 tape 000061L2 on rd=server_2:c32t0l1
```

9. To obtain the saveset id (SSID) from the backup tape, type:

   `mminfo -avV` *`<tape name>`*

   where:

   `<tape name>` = name of the tape that is mounted on the NDMP TLU

   Example:

   `mminfo -avV 00061L2`

   Output:

```
volume   client   size    level name ssid save time date
time browse clretent first  last file rec volid  total fl
000061L2 server_2 1042 MB full/1G 3850629539 1199829345 1/8/2008
4:55:45 PM 2/9/2008 1/9/2009 1067113652 3 0 4286834478 1067113653
crN
```

From the output the saveset id (SSID) is 3850629539.

10. To run inventory on the destination TLU slots that contain the tape 00061L2, on the destination DMA server, use this command syntax:

    nsrjb -IEvvv -S *<slot_number>*

    where:

    <slot_number> = slot number of destination TLU to run inventory

    Example:

    nsrjb -IEvvv -S 1-2

11. To find the client ID from the index that was scanned, use this command syntax:

    mminfo -aS -q "ssid=*<ssid>***"**

    where:

    <ssid> = saveset id obtained from the backup tape

    To find the client ID from the index, type:

    mminfo -aS -q "ssid=3850629539"

    Output:

    ```
    ssid=3850629539 savetime=1/8/2008 4:55:45 PM (1199829345)
    server_2:/1G level=full sflags=vrNF size=1067113653 files=2
    insert=1/9/2008 create=1/8/2008 complete=1/8/2008 browse=2/9/2008
    2:41:08 PM retent=1/9/2009 2:41:08 PM
    clientid=27aab1da-00000004-4783e62c-4783e62b-00010000-ac18fcde
    *DIRECT: y, y;
    *FILESYSTEM: /1G;
    *HIST: y, y;
    ```

    From the output, the client ID is 27aab1da-00000004-4783e62c-4783e62b-00010000-ac18fcde.

12. To create a new backup client:

    a. In the **Properties** dialog box, click the **Globals (1 of 2)** tab and type the client ID in the **Client ID** field.

    b. Type the server name of the hosts file in the **Aliases** field.

c. Type the appropriate information for the rest of the tabs in the NeWorker client.

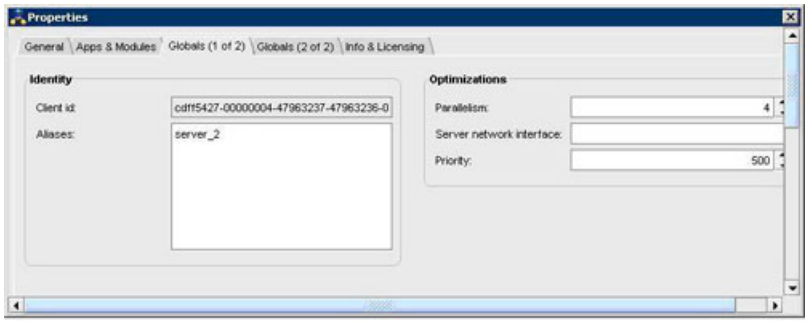

Note: If a backup client having the same name already exists, delete and re-create it.

13. To restore the saveset id from the backup tape, use NetWorker or type:

nsrndmp_recover -s *<server name>* **-c** *<client name>* **-S** *<ssid>* **-v -m** *<location for restore>*

Example:

```
nsrndmp_recover -s win_server01 -c server_2 -S 3850629539 -v -m server_2::/NY_fs
```

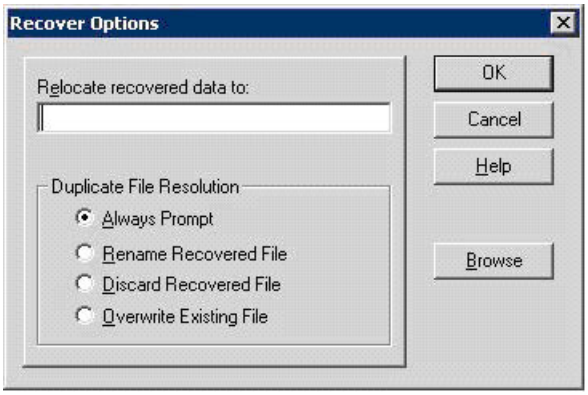

Note: Ensure that the destination directory is set up correctly. If not, from the Options menu in NetWorker, click Recover Options, and then type the new directory name in the Recover Options dialog box.

14. Start restore to the destination file system and verify that the restore is completed successfully by using the server log of the destination Data Mover. Search entries with VSB labels in it.

Example:

```
2008-01-26 00:11:00: VSB: 6: Restore data verified[VsbRestoreTask:
:run, sessionId = 1]
2008-01-26 00:11:01: VSB: 6: Snap refresh done[VsbRestoreTask:
:run,1]
2008-01-26 00:11:01: VSB: 6: Successfully applied replication
signature to version[VsbTask::ApplyVersionSignature, sessionId=1]
2008-01-26 00:11:01: VSB: 6: Restore succeeded[VsbSession::Restore
OneVersion, 1]
2008-01-26 00:11:01: VSB: 6: Total restore data :323350 MB
[VsbRestoreTask::printStat, sessionId=0]
2008-01-26 00:11:01: VSB: 6: Restore time : 27572 Sec
[VsbRestoreTask::printStat, sessionId=1]
2008-01-26 00:11:01: VSB: 6: Restore speed is : 11 MB/s
[VsbRestoreTask::printStat, sessionId=1]
```

15. To start the replication session, from the source Control Station, use this command syntax:

    # **nas_replicate -start** *<replication_session_name>* **overwrite _destination**

    where:

    <replication_session_name> = name of the replication session

    Example:

    To start the replication session, type:

    # nas_replicate -start NY_src_fs_to_LA -overwrite_destination

    Output:

    ```
    OK
    ```

16. To verify that the replication session has started successfully, type:

    # nas_replicate -list

    Output:

    ```
    Name         Type       Local Mover Interconnect Celerra  Status
    NY_fs-to-41  filesystem server_2    <--41s2-40s2 eng25240 Runing
    ```

5 Troubleshooting

As part of an effort to continuously improve and enhance the performance and capabilities of its product lines, EMC periodically releases new versions of its hardware and software. Therefore, some functions described in this document may not be supported by all versions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.

If a product does not function properly or does not function as described in this document, contact your EMC Customer Support Representative.

Topics included are:

- EMC E-Lab Interoperability Navigator on page 102
- VNX user customized documentation on page 102
- Troubleshooting checklist on page 102
- Reporting a problem on page 103
- Known problems and limitations on page 104
- Error messages on page 104
- EMC Training and Professional Services on page 105

EMC E-Lab Interoperability Navigator

The EMC E-Lab™ Interoperability Navigator is a searchable, web-based application that provides access to EMC interoperability support matrices. It is available on EMC Online Support at http://Support.EMC.com. After logging in, in the right pane under **Product and Support Tools**, click **E-Lab Navigator**.

VNX user customized documentation

EMC provides the ability to create step-by-step planning, installation, and maintenance instructions tailored to your environment. To create VNX user customized documentation, go to: https://mydocs.emc.com/VNX.

Troubleshooting checklist

Review this checklist for tips on troubleshooting your NDMP backup:

- Is the backup or restore spanning multiple tapes? If yes, are more tapes available when the job spans to the next tape?
- Verify the TLU physical connections and correct any loose connections or incorrectly labeled hardware. Set up the TLU on page 53 provides more information.

  Verify the server_log to ensure that the parameters executed successfully. To read the server_log for a Data Mover, use this command syntax:

  ```
  $ server_log <movername> | grep param
  ```

  where:

  <movername> = name of the specified Data Mover

  Configure the Data Mover on page 141 provides more information on setting parameters.
- Are the necessary parameters for NDMP and PAX set on the Data Mover? Configure the Data Mover on page 141 offers more information.
- Ensure that the length of path including the filename of the source file does not have more than 4,095 bytes.
- Confirm in the E-Lab Interoperability Navigator that VNX, backup software version, and hardware are supported for NDMP backups with VNX.
- Have you created the usernames and passwords on all NDMP-Host Data Movers? Assign a username and a password to a Data Mover on page 56 provides more details. Have you used the same usernames and passwords when configuring Data Mover access in the backup software?

- Have the TLU device names been recorded correctly and typed during configuration of the backup software? Set up the TLU on page 53 and your backup software vendor documentation provide more information.
- If your backup vendor supplies utilities for verifying the physical and logical paths between the backup server and the TLU devices, confirm that these paths are valid.
- Test a backup or restore by using VNX server_archive utility. By eliminating the NDMP protocol, you are able to eliminate possible problems and more easily troubleshoot issues. *Using the server_archive Utility on VNX* explains how to use this utility.
- If using the SNAPSURE=y environmental variable (or manually by using a checkpoint for the backup), and the checkpoint in use is refreshed, unmounted, or deleted, the backup fails and a message is sent to the event log. (*Using VNX SnapSure* provides information on troubleshooting checkpoint creation failure.)
- If a timeout occurs while creating a checkpoint, an UNMOUNT_AND_DELETE event is sent to the Control Station. The backup fails and the failure is recorded in the server_log. If the checkpoint is actually created after timeout, it will be deleted by the event.

If the backup still does not work correctly after you verify these things, contact the backup software vendor for further support.

Reporting a problem

Problem Resolution Roadmap for VNX identifies troubleshooting resources and processes. In addition to the general information in Problem Resolution Roadmap for VNX, gather these backup-specific information before reporting a problem to the EMC Support Center:

- At what stage did the problem occur?
- What is the version and patch number of your backup software?
- Is your backup software installed on a UNIX or Windows platform?
- What is the IP and hostname of the system running the backup software?
- Are you backing up CIFS files, NFS files, or both?
- What model is the TLU?
- Does your environment have a two-way or three-way backup configuration?

Known problems and limitations

Table 5 on page 104 describes known problems that can occur when using NDMP backups and presents workarounds.

Table 5. Known problems and workarounds

Known problem	Symptom	Workaround
A full restore recovers only data from the last full backup. The restore does not include subsequent incremental backup data.	You run a full restore, and incremental backup data is not included with the restored data.	Perform a separate restore of the incremental backup data.

Error messages

All event, alert, and status messages provide detailed information and recommended actions to help you troubleshoot the situation.

To view message details, use any of these methods:

- Unisphere software:
  - Right-click an event, alert, or status message and select to view Event Details, Alert Details, or Status Details.
- CLI:
  - Type nas_message -info <MessageID>, where <MessageID> is the message identification number.
- *Celerra Error Messages Guide*:
  - Use this guide to locate information about messages that are in the earlier-release message format.
- EMC Online Support:
  - Use the text from the error message's brief description or the message's ID to search the Knowledgebase on EMC Online Support. After logging in to EMC Online Support, locate the applicable **Support by Product** page, and search for the error message.

EMC Training and Professional Services

EMC Customer Education courses help you learn how EMC storage products work together within your environment to maximize your entire infrastructure investment. EMC Customer Education features online and hands-on training in state-of-the-art labs conveniently located throughout the world. EMC customer training courses are developed and delivered by EMC experts. Go to EMC Online Support at http://Support.EMC.com for course and registration information.

EMC Professional Services can help you implement your system efficiently. Consultants evaluate your business, IT processes, and technology, and recommend ways that you can leverage your information for the most benefit. From business plan to implementation, you get the experience and expertise that you need without straining your IT staff or hiring and training new personnel. Contact your EMC Customer Support Representative for more information.

Appendix A

Configure your Backup Software

This appendix lists restrictions and configuration steps for the various backup software vendors EMC-qualified for NDMP backups with VNX. The E-Lab Interoperability Navigator is the definitive source for information on NDMP backup software vendor qualification.

These guidelines apply to all backup software vendors:

- Use the same password used in Assign a username and a password to a Data Mover on page 56.
- Use this format for device names, which VNX uses to address a TLU device:

  <SCSI_controller> <SCSI_ID> <SCSI_LUN>

  where:

  <SCSI_controller> = chain number of the controller

  <SCSI_ID> = SCSI ID (target) of the controller

  <SCSI_LUN> = logical unit number (LUN) of the controller

- Configure port 10000 for NDMP operations. VNX uses dynamic port allocation from the Data Mover for NDMP three-way backups and restores.
- Use the Data Mover name as the hostname, for example, server_2.

Topics included are:

- Terminology differences on page 109
- Configure Atempo Time Navigator on page 109
- Configure BakBone NetVault on page 113
- Configure CA BrightStor ARCserve Backup on page 114
- Configure CommVault Galaxy on page 117
- Configure EMC Data Manager with NDMP on page 120
- EMC NetWorker on page 124

- Fujitsu-Siemens NetWorker on page 131
- Configure HP OpenView Storage Data Protector on page 132
- IBM Tivoli Storage Manager on page 135
- Configure Syncsort Backup Express on page 141
- Configure Symantec NetBackup on page 143
- Configure Symantec Backup Exec on page 154

Terminology differences

Table 6 on page 109 lists terminology differences that exist between this document and the third-party documentation.

Table 6. Terminology differences

EMC term	Equivalent third-party terms
Tape library unit (TLU), jukebox, tape drive	Library (refers to any backup storage device), autochanger handle, autochanger
TLU robot	Robotic arm, arm changer
NDMP-Host Data Mover	VNX filer, NAS filer, application server
VNX	Client computer, VNX filer

Configure Atempo Time Navigator

The procedures for configuring Time Navigator are based on Atempo Time Navigator version 4.2. These procedures describe how to configure NDMP backups by using VNX and Atempo Time Navigator.

For information on how to set environment variables for Atempo Time Navigator, see Set environment variables on page 58.

The tasks to configure Atempo Time Navigator are:

1. Create the file server on page 110
2. Create and configure a library on page 110
3. Configure the drives on page 111
4. Create a media pool on page 111
5. Create the Time Navigator client on page 112

Create the file server

1. From the **Time Navigator Administration** application, select **Platform ➤ New ➤ Host** to open the **New Host** window.
2. In the **General** tab, type the following information:
   - **Type list**: Select the file server type.
   - **Host Name**: Type the Data Mover name, for example, server_2.
3. In the **Protocol** tab, the NDMP box is selected. Choose the NDMP version and type the following information:
   - **User**: Type **ndmp**
   - **Password**: Type the password that you used in Assign a username and a password to a Data Mover on page 56.
   - **Confirm**: Retype the password.
4. In the **Advanced** tab, the **Backup Master** option is automatically selected. Select **Server** if the NDMP client is installed on the **Time Navigator** server, else select **Host** and provide the name of the Time Navigator client acting as the Backup Master in the **Master Name** field.
5. Click **OK**.
6. Right-click the host icon you have just created, and select **Enable**.

Create and configure a library

1. From the **Time Navigator Administration Console**, select the NDMP host that you created in Create the file server on page 110.
2. Select **Devices ➤ Library ➤ New** and then select a library type to open the **New Library** window.
3. Type the **Library** name and the **Device Descriptor**. The device descriptor uses the format:

   <SCSI_controller> <SCSI_target> <LUN>

   Set up the TLU on page 53 provides more information.
4. Associate the previously created drives to the library.

5. Click **OK**. Time Navigator performs a reinitialization of the library. You can also use the **Time Navigator Device Detection Wizard** to detect and help you configure your library. In the **Time Navigator Administration Console**, select **Devices** and then select **Device Detection Wizard**.

Configure the drives

1. In **Time Navigator Administration Console**, select the file server host (Data Mover) that you created in Create and configure a library on page 110.

2. Select **Devices ➤ Library ➤ New**, and then select a drive type to open the **New Drive** window.

3. Type the appropriate information in the fields. For device descriptor use the format:

   <SCSI_controller> <SCSI_target> <LUN>

   Set up the TLU on page 53 provides more information.

4. Click **OK**.

5. Right-click the drive icon and select **Enable**. Repeat the procedure if you are creating more than one drive.

Note: You can also use the **Time Navigator Device Detection Wizard** to detect and configure your drives. From the **Time Navigator Administration Console**, select **Devices** and then select **Device Detection Wizard**.

Create a media pool

1. From the **Time Navigator Administration Console**, select **Storage ➤ MediaPool ➤ New** to open the **New Media Pool** window.

2. Type the necessary information and ensure that you associate the media pool to the previously created drives.

3. Click **OK**.

Create the Time Navigator client

1. From the **Time Navigator Administration Console**, select the host where you want to create the **Time Navigator Client** for NDMP application.

   Note: Do not select the NDMP server, as the host should be able to run the tina_daemon, which the NDMP server cannot do.

2. Select **Platform ➤ New ➤ Application**.
3. Select the NDMP application.
4. Click **OK**. The **New Application** window appears.
5. Type the following information:
   - **Host Name**: Name of the host where Time Navigator resides. The same system you selected in step 1.
   - **Application Name**: Name of the NDMP application.
   - **Comment**: Comment concerning the application you are creating.
   - **NDMP Server**: The NDMP-Host Data Mover you want to back up.
   - **NDMP User Name**: Type ndmp, which is also the username you entered for the NDMP-Host Data Mover in Assign a username and a password to a Data Mover on page 56.
   - **NDMP Password**: Password you typed for the NDMP-Host Data Mover in Assign a username and a password to a Data Mover on page 56.
   - **Local User**: Name of an operating system user who can access the host where the NDMP application resides.
   - **Password**: Password for the username you typed in the previous field.
6. Click **OK**.
7. Right-click the NDMP application and select **Enable**.

Configure BakBone NetVault

The procedures for configuring BakBone NetVault are based on NetVault version 7.1. They describe how to configure your BakBone NetVault software for NDMP backups with VNX.

When performing NDMP backups with the VNX and BakBone NetVault, version 6.3 or later of the NDMP Application Module (APM) must be installed on the client running the NetVault backup software.

The EMC E-Lab Interoperability Navigator provides a list of qualified APM versions. NetVault documentation provides information on installing the NDMP APM.

The tasks to configure BakBone NetVault are:

- Add an NDMP server on page 113
- Use the automatic device configuration on page 114

Add an NDMP server

1. From the main **NetVault** window, double-click the **Backup** icon. The **NetVault Backup** window appears.
2. Double-click the BakBone client name.
3. Right-click the NDMP client and select **Add Server**.
4. In the **NDMP server** dialog box, type the required information:
   - **Name**: Name of the NDMP-Host Data Mover.
   - **Address**: IP address of the NDMP-Host Data Mover.
   - **Port**: NDMP server port number, which you must set as **10000**.
   - **Account**: The username, which must match the one you set for the NDMP-Host Data Mover.
   - **Password**: Password for the username account, which must match the password you set for the NDMP-Host Data Mover.
5. Click **OK**.

Use the automatic device configuration

NetVault 7.1 automatically locates the devices on the network. After you add the NDMP server, complete the procedure in this section, which describes how to recognize and configure a device by using the NetVault wizard. BakBone NetVault documentation provides information about how to manually recognize and configure devices in the following steps:

1. Double-click **Device Management** to open the **NetVault Device Management** window.
2. From the **Add** menu, select **Add Library**. The **Add Library** window appears.
3. From the **Library Selection** tab, double-click **BakBone client**.

   Under **Ndmp Tape**, select the correct drive to which you want the drive-bay value assigned.
4. Double-click the library you want to configure.
5. Click the **Drive Selection** tab.
6. Double-click **BakBone client**.
7. By using the **Select for drive bay**, select or type a valid drive-bay value.
8. Under **Ndmp Tape**, select the correct drive to which you want the drive-bay value assigned.
9. Repeat steps 7 and 8 to assign a drive-bay value to each possible drive in that library.
10. Click the **Configuration** tab.
11. Right-click the library, and select **Scan for shared drives**.
12. Click **OK**, and then click the **Save** details icon in the top-left corner.
13. In the **Library added successfully** dialog box, click **OK**, and then close the **Add Library** window.

Configure CA BrightStor ARCserve Backup

CA BrightStor ARCserve Backup configuration procedures are based on CA BrightStor ARCserve Backup version 11.1:

1. Configure the NDMP devices on page 115
2. Configure the nas.cfg file on page 116
3. Add a Data Mover to ARCserve Backup on page 117

Note: Filters are unsupported when performing NDMP backups with VNX and CA BrightStor ARCserve Backup. If you are using the version 5.4.21.0 or later, ensure that the NDMP parameter v4OldTapeCompatible is set to 0. If you plan to use Dynamic Drive Sharing, set the NDMP scsiReserve parameter to 0 on the Data Mover running version 5.1 or later, where devices are connected. Modify parameters on page 69 explains how to change parameter settings.

Configure the NDMP devices

1. From the **BrightStor ARCserve Backup** home page, click **Configuration** to expand the Configuration subtopics.
2. In **Configuration**, click **Device Configuration Wizard**.
3. On the **Welcome** screen, select **Windows Server** and click **Next**.
4. In **Options**, select **NAS Servers** and click **Next**.

   This message appears:

   ```
   Tape Engine service will be stopped in order to allow device
   configuration. Would you like to continue?
   ```

5. Click **Yes**, which stops any other backups that are running.
6. In **NAS Server Configuration,** click **Add** and select **NAS Server**.
7. In the **Device Configuration** dialog box, in **Server Configuration**, click **New Server** and type the name of the Data Mover.
8. In the **User** column, type the username that you have created for the NDMP-Host Data Mover.
9. Double-click **Password** and type the same password that you have created for the NDMP-Host Data Mover in the **Password** field.
10. Type the password again in **Confirm password**, and click **OK**.
11. Select **Continue** to configure libraries and click **Next**.

    Tape Library Option appears, which can take a few minutes.
12. Ensure that the proper devices appear in the **Library devices** box and click **Next**.
13. In **Assign cleaning tape,** click **Properties**, select the properties appropriate for your environment, and click **OK**.
14. Click **Next**.

15. Select or clear **Library Quick Initialization**, depending on your configuration, and click **Next**.

16. Click **Finish**, and then click **Exit**.

17. Click **Yes** and then start the tape engine when the following message appears:

    ```
    Exit Device Configuration?
    ```

Configure the nas.cfg file

You can configure the nas.cfg file to perform partial-volume backups. This backup type requires that you type the paths in the configuration file.

1. Open the nas.cfg file. This file is located in the *<base_install>*\NAS Option folder.

2. On the first line, type the name of the Data Mover.

3. On separate lines, type each absolute path, starting with the logical device name. Use this format:

   <movername> /volume1/folder1/folder2 /volume2/folder3/folder4

   where:

   <movername> = name of the Data Mover

   Note: To end each server configuration, type a semicolon.

   Example:

   server_2 /volume1/folder1/folder2/ volume2/folder3/folder4

4. Save the file.

5. To run a backup, use the Backup Manager to configure and submit a backup. Select a Data Mover (VNX) as your source and the tape device as the target. The tape device can be connected to the same VNX or to another VNX.

Add a Data Mover to ARCserve Backup

1. From the **BrightStor ARCserve Backup** home page, in **Quick Start**, click **Backup**.
2. Click **Source**, right-click **NAS Servers**, and select **Add Machine/Object**.
3. In **Add Server**, type the name of the Data Mover.
4. Select the appropriate settings for your configuration, click **Add** and then click **Close**.

   The Data Mover name should appear under the NAS Servers heading on the **BrightStor ARCserve Backup** home page.
5. Double-click the Data Mover you added in step 4, and then in the **Security** dialog box, type the username; then type the password.

   Note: You must use the same username and the same password you typed in Assign a username and a password to a Data Mover on page 56.

Configure CommVault Galaxy

CommVault Galaxy is not compatible with the Volume-Level Copy (VLC) feature, which is an environment variable setting. This restriction applies when performing NDMP backups with VNX and CommVault Galaxy.

The procedures herein are based on an autodetect configuration of CommVault Galaxy version 7.0.

Note: If you have a license for SCSI-3 Reserve with CommVault Galaxy version 5.9 or later, and you plan to use Dynamic Drive Sharing with CommVault Galaxy, set the NDMP scsiReserve parameter to 0 on the Data Mover, running version 5.1 or later, where devices are connected. Modify parameters on page 69 explains how to change this parameter.

The tasks to configure CommVault Galaxy are:

1. Add a client on page 118
2. Configure the TLU and tape drives for NDMP operations on page 118
3. Test the TLU device on page 120

Add a client

1. From **CommCell Browser**, right-click **Client Computers** and select **New Client**. The **Add New Client** window appears.
2. Select **NAS Client** and click **OK**. The **Add NDMP Server** window appears.
3. Specify the attributes:
   - **NDMP server Hostname**: Hostname of the NDMP server.
   - **NDMP Login**: Type the same username you created for the NDMP-Host Data Mover.
   - **NDMP Password**: Type the same password you created for the NDMP-Host Data Mover.
4. Confirm the password, and then click **Detect**.

   The system detects the EMC VNX as vendor and displays the firmware revision, which is the VNX software running on the NDMP server. It also displays the listen port, which appears as 10000.
5. If all the information is correct, click **OK**. If not, click **Cancel**.

   The newly added server appears in CommCell Browser.

Configure the TLU and tape drives for NDMP operations

1. From **CommCell Browser**, select **Tools ➤ Control Panel**. Select and double-click **Library and Drive Configuration**.

   The **Library and Drive Configuration** window appears.
2. Select the **MediaAgents** window.
3. Select the media agent to which you want to associate the library, and click **Add** to move the selection to the **Selected MediaAgents** side. Click **OK** to close the window.
4. From the **Galaxy Library** window, click **Start ➤ Detect/Configure Devices**. The **Detect Library** window appears.
5. Click **NDMP Devices**, and click **OK**. The **Select NDMP servers to Detect** window appears with a list of NDMP servers.
6. If the desired NDMP servers appear in this list, select them, click **Add** to move them to the selected NDMP servers side, and click **OK**. A confirmation window appears asking if you want to continue. Go to step 10.

If the desired NDMP server does not appear in the list, click **Update NDMP Host List** to refresh the server list and try again.

The **NDMP server List** window appears.

7. Click **Add**. The **Add NDMP server** window appears.

8. Type the **NDMP server Hostname**.

   Type the **NDMP Login**, which should be the same username you created in Assign a username and a password to a Data Mover on page 56.

   Confirm the password, and then click **Detect**.

9. The system detects the EMC VNX as vendor, and displays the firmware revision, which is the VNX software running on the NDMP server. It also displays the listen port, which appears as 10000.

   If all the information is correct, click **OK**. If not, click **Cancel**. The **NDMP server List** window reappears, listing the newly added server.

10. Repeat steps 7 and 8 for each new server you want to add. After all the required servers are added, click **Close**. The **Select NDMP Servers to Detect** window reappears. Click **OK.**

    A confirmation window appears asking if you want to continue.

11. If so, click **Yes**. The **Processing** window appears showing the progress and time to complete. This can take time, depending on the number of devices to detect.

    The **Log** window appears showing the log of the operation.

12. Click **Close**. An information window appears stating: “Right-click on a particular device to continue with its configuration." Click **OK**.

13. From the **Library and Drive Configuration** window, click **Libraries** to view the newly added libraries beside the yellow library icon. The question marks mean the library is not yet fully configured.

14. Expand the libraries to view the details and ensure that all required tape drives are shown under the library.

15. Configure the library by right-clicking the library and selecting **Configure**. Complete the information requested.

16. Click **Start ➤ Exit** and close the control panel.

Test the TLU device

1. After you configure the library, ensure that the library is accessible to CommVault Galaxy.

   From **Control Panel**, select and double-click **Library and Drive Configuration**. The **Library and Drive Configuration** window appears.

2. Select the **MediaAgents** window.

3. Select the media agent to which you want to associate the library and click **Add** to move the selection to the **Selected MediaAgents** side. Click **OK** to close the window.

   The newly added library appears.

4. From the **Galaxy Library** window, right-click the library and select **Validate**.

   The **Confirm Validate** window appears stating: `Validating a library will mount a media in each drive and can take several minutes per drive. Are you sure you want to validate drives in library.`

5. Click **Yes**. The **Processing** window appears showing the progress of validation, and then the **Log** window appears.

6. Click **Close.**

   If the validation fails, a window appears displaying the error message. Click **OK**. Verify the log and click **Close.**

7. To end the session, click **Start ➤ Exit**.

Configure EMC Data Manager with NDMP

EMC Data Manager (EDM) is a high-performance backup solution optimized for high-end UNIX and Windows NT environments. EDM offers centralized management of scheduling, monitoring, reporting, and configuring for all NDMP backups. It initiates data movement that is then performed by the NDMP software on VNX, thus distributing resource usage and freeing the EDM for other backup tasks. File server clients are never affected by backup processing, allowing users uninterrupted operation.

You can configure EDM to have NDMP send data directly to a TLU attached to VNX, or to use the the network to transport data from one file server to another, and then to a TLU attached to the secondary file server. With EDM and NDMP, data flows from the system to the Data Mover, and then directly to the VNX-attached TLU through ultra fast-wide SCSI or Fibre Channel connections. Network congestion is minimized because the data path and control path are separated. Backup can occur locally, from VNX to the tape drives, while management occurs centrally from a single EDM backup server.

With EDM support for VNX:

- EDM backs up all VNX native attributes, including NFS and CIFS/NT attributes and ACLs.
- Data moves from disk to local tape through the NDMP server, eliminating the need to move the backup data over the network.

EDM allows you to use one of two hardware configurations with VNX:

- Local backups
- Three-way backups

Local backups with EDM

In the local backup, data transfer occurs between the Data Mover and the TLU over ultra fast-wide SCSI or Fibre Channel connections. This configuration is achieved by attaching drives from a TLU to each NDMP-Host Data Mover in the configuration. Each NDMP host runs all NDMP services such as data and tape/mover. Figure 10 on page 121 illustrates an EDM NFS backup.

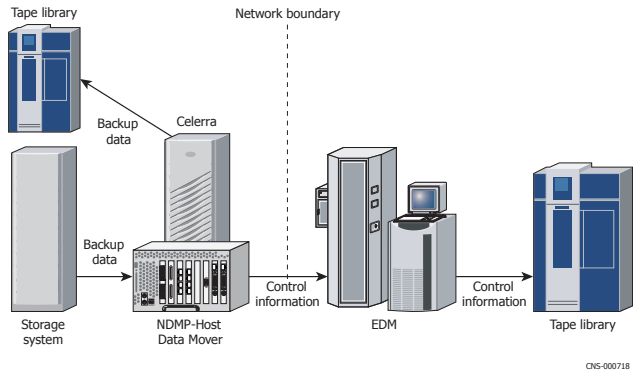

Figure 10. EDM NFS backup

Three-way backups with EDM

In the centralizing EDM configuration shown in Figure 11 on page 122, you can transfer all backup data to a single TLU even if the data is served from more than one VNX. Bandwidth is limited to the bandwidth available on the LAN that is used to transfer the data.

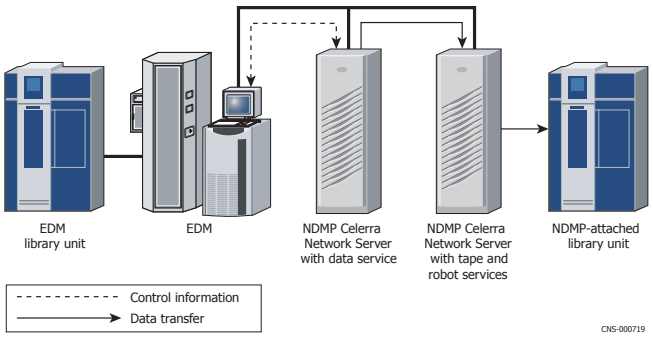

Figure 11. Three-way backups with EDM

EDM restrictions

These restrictions apply when performing NDMP backups with VNX and EDM:

- Configuring the same NDMP-attached TLU on more than one EDM is not supported. This causes unpredictable results and can cause data loss.
- The bufsz parameter setting on VNX must be at least 128 KB, which is the default. Configure the tape buffer size on page 46 explains how to change this parameter.
- The readWriteBlockSizeInKB parameter setting on VNX must be at least 128. Configure the read or write block size on page 47 explains how to change this parameter.
- Stand-alone drives are unsupported, drives must reside in a TLU.
- The TLU robot must be connected to the EDM only through VNX. Connecting the NDMP TLU robot directly to the EDM system is unsupported.
- The EDM-attached TLU used to back up the EDM and to perform network, Symmetrix Path, and Symmetrix Connect backups and restores cannot be the same TLU used to perform NDMP backups. Do not attempt to share the TLU used to back up the EDM and perform other types of EDM backups with the NDMP solution.
- Incremental backups from VNX cannot detect filename changes or deleted files. Therefore, multiple files can be restored if filenames were changed or files were deleted.
- EDM version 5.1 or later is required for NDMP backups with this version of VNX.

This section assumes the robot is directly attached to the NDMP-Host Data Mover. The EDM documentation noted in Related Information on page 16 provides more details on this procedure.

To configure your TLUs:

1. Run **lmconfig**.
2. When the **EMC Library Manager Configuration Tool** menu appears, select the **AUTOCONFIG** option, and then note the information that follows.

Note: The LIST option does not display NDMP devices in the unconfigured devices list because they are not locally attached.

Install and configure the EDM client

Follow the EDM documentation to install and configure the client.

Before you can perform NDMP backups, EDM must have access to each Data Mover involved in the backup process.

Perform the following steps to install and configure the EDM client:

1. Type a trusted account username and password in the **EDM Configuration Wizard** when prompted. The username and password should match those assigned to each NDMP-Host Data Mover.

2. Run **lmconfig** to configure your TLUs. The **EMC Library Manager Configuration Tool** appears.

3. Select **AUTOCONFIG** and note the information that follows.

   Note: The **EDM Configuration Wizard** automatically locates and tests the TLU connections.

EMC NetWorker

These restrictions apply when performing NDMP backups with VNX and EMC NetWorker:

- The bufsz parameter setting, which is called block size in NetWorker documentation, must be set based on the tape drive vendor. Tape drive buffer sizes on page 129 lists the settings for different tape drives. Configure the Data Mover on page 141 shows how to set the bufsz parameter.
- VNX allows only one device name per robot and one device name per tape drive. Use this device name when configuring the robot or a tape drive for the NDMP backup software. The name for this device should be an entity that can be resolved through the network, such as server_2:c1t0l0, where:

  server_2 = customer-assigned Data Mover name c1t0l0 = robot or tape device name

- Ensure that the HIST and DIRECT environmental variables are enabled prior to the backup to make DAR available for file recovery.

Configure EMC NetWorker

The procedures for configuring NetWorker for NDMP backups with VNX are based on NetWorker version 7.3.x or later:

1. From a Windows backup client running NetWorker, type:

   ```
   <networker_install_path>\bin\directory
   ```

   From a UNIX backup client running NetWorker, proceed to step 2.

2. At the prompt, type:

   ```
   # jbconfig
   ```

   To complete the jbconfig script, you need:

   - The password for the NDMP-Host Data Mover where the tape device is attached. This is the password you set in Assign a username and a password to a Data Mover on page 56.
   - The SCSI device name. Set up the TLU on page 53 provides information on accessing device names.

3. When prompted, select the **Autodetect NDMP SCSI Jukebox** option.

4. Complete the interactive jbconfig script. Use the Sample jbconfig script on page 125 as an example.

Sample jbconfig script

The NetWorker configuration script looks similar to the example in this section. Within the script, bold text indicates user entry. Brackets indicate default values. You can press Enter to accept a default value, or you can type the value:

```
Script started on Mon Aug 27 15:09:13 2007

            # jbconfig

            Jbconfig is running on host ohnetworker (SunOS 5.8),

              and is using ohnetworker as the NetWorker server.
                     1) Configure an AlphaStor Library.

                     2) Configure an Autodetected SCSI Jukebox.

                     3) Configure an Autodetected NDMP SCSI Jukebox.

                     4) Configure an SJI Jukebox.

                     5) Configure an STL Silo.

            What kind of Jukebox are you configuring? [1] 3

            Enter NDMP Tape Server name: ? scarlos5


            Communicating to devices on NDMP Server 'scarlos5', this
may take a while...

            14484:jbconfig: Scanning SCSI buses; this may take a
while ...

            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0

            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0

            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0

            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0

            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0

            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0

            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0

            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0
```

```
NDMP Service Debug: The process id for NDMP service is
0xbbb633c0

            Installing 'Standard SCSI Jukebox' jukebox - scsidev@0.0.0.

            What name do you want to assign to this jukebox device?
Garbage


            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0

            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0

            15814:jbconfig: Attempting to detect serial numbers on
the jukebox and drives ...

            15815:jbconfig: Will try to use SCSI information returned
by jukebox to configure drives.

            Turn NetWorker auto-cleaning on (yes / no) [yes]?

            The following drive(s) can be auto-configured in this
jukebox:

             1> 4mm @ 0.0.1 ==> c0t0l1 (NDMP)

             2> 4mm @ 0.0.2 ==> c0t0l2 (NDMP)

             3> 4mm @ 0.0.3 ==> c0t0l3 (NDMP)

             4> 4mm @ 0.0.4 ==> c0t0l4 (NDMP)

             5> 4mm @ 0.0.5 ==> c0t0l5 (NDMP)

             6> 4mm @ 0.0.6 ==> c0t0l6 (NDMP)

            These are all the drives that this jukebox has reported.

            To change the drive model(s) or configure them as shared
or NDMP drives,

             you need to bypass auto-configure. Bypass auto-configure?
(yes / no) [no] yes


            Is (any path of) any drive intended for NDMP use? (yes
/ no) [no] yes


            Is any drive going to have more than one path defined?
(yes / no) [no] no


            Please enter the device path information in one of the
following formats:

            /dev/rmt/1cbn --for local path or

            host:device-path --for remote node or NDMP device(s)
or
```

```
            host:drive-letter:directory path --for Windows disk file

               After you have entered a device path, you will be
prompted for an NDMP

            user name for that path's host. If this device path is
not an NDMP device,

            press the enter key to advance to the next device path.
 For NDMP devices,

            you need to enter the user name and password the first
time we encounter

            that NDMP host. Pressing the enter key for the NDMP user
name for any

            subsequent device path on the same host will set the
user name and password

            to those defined the first time. You will not be prompted
for the password

            in such a case.

            Drive  1, element 1, system device name = c0t0l1,

                  local bus, target, lun value = 0.0.1,

                      TQ2RvdO200 model 4mm

            Drive path ? [scarlos5:c0t0l1]


            Is this device configured as NDMP? (yes / no) [no]yes


            Drive  2, element 2, system device name = c0t0l2,

                  local bus, target, lun value = 0.0.2,

                      9nrZwdO200 model 4mm

            Drive path ? [scarlos5:c0t0l2]

            Is this device configured as NDMP? (yes / no) [no]yes
```

```
Drive  3, element 3, system device name = c0t0l3,
                local bus, target, lun value = 0.0.3,
                    pTP7ydO200 model 4mm
            Drive path ? [scarlos5:c0t0l3]
            Is this device configured as NDMP? (yes / no) [no]yes

            Drive  4, element 4, system device name = c0t0l4,
                  local bus, target, lun value = 0.0.4,
                      FpegzdO200 model 4mm
            Drive path ? [scarlos5:c0t0l4]
            Is this device configured as NDMP? (yes / no) [no]yes

            Drive  5, element 5, system device name = c0t0l5,
                  local bus, target, lun value = 0.0.5,
                      VVCoAdO200 model 4mm
            Drive path ? [scarlos5:c0t0l5]
            Is this device configured as NDMP? (yes / no) [no]yes

            Drive  6, element 6, system device name = c0t0l6,
                  local bus, target, lun value = 0.0.6,
                      bs1xBdO200 model 4mm
            Drive path ? [scarlos5:c0t0l6]
            Is this device configured as NDMP? (yes / no) [no]yes

            Only model 4mm drives have been detected.
            Are all drives in this jukebox of the same model?  (yes
/ no) [yes] yes

            NDMP Service Debug: The process id for NDMP service is
0xbbb633c0
            Jukebox has been added successfully
            The following configuration options have been set:
            > Jukebox description to the control port and model.
            > Autochanger control port to the port at which we
found it.
```

```
            > Networker managed tape autocleaning on.

            > Barcode reading to on.

            > Volume labels that match the barcodes.

            > Slot intended to hold cleaning cartridge to 32.
 Please insure that a

                    cleaning cartridge is in that slot

            > Number of times we will use a new cleaning cartridge
to 25.

            > Cleaning interval for the tape drives to 6 months.

            You can review and change the characteristics of the
autochanger and its

                    associated devices using the NetWorker Management
Console.

            Would you like to configure another jukebox? (yes/no)
[no]no


            # sync;sync;

            # exit

            script done on Mon Aug 27 15:16:31 2007
```

Tape drive buffer sizes

Table 7 on page 129 lists the buffer size settings, called block size in NetWorker documentation, for different types of tape drives. You must set the bufsz parameter on VNX by using the appropriate value for your tape drive type. Configure Data Movers on page 46 explains how to set the bufsz parameter.

Table 7. Tape drive buffer size settings

Tape drive type	Buffer size (block size) setting	Tape drive type	Buffer size (block size) setting
3480	32 KB	8 mm Mammoth-3	256 KB
3570	256 KB	9840	256 KB
3590	384 KB	9940B	256 KB
3592	384 KB	9940	256 KB
4890	32 KB	9840b	256 KB

Table 7. Tape drive buffer size settings *(continued)*

Tape drive type	Buffer size (block size) setting	Tape drive type	Buffer size (block size) setting
4 mm	32 KB	9840C	256 KB
4 mm 4 GB	32 KB	9490	64 KB
4 mm 8 GB	32 KB	DLT	96 KB
4 mm 12 GB	32 KB	DLT1	32 KB
4 mm 20 GB	32 KB	DLT7000	128 KB
8 mm	32 KB	DLT8000	96 KB
8 mm 5 GB	32 KB	DLT-S4	128 KB
8 mm 7 GB	32 KB	DLT-V4	128 KB
8 mm AIT	9840C	DLT VS160	128 KB
8 mm AIT-2	64 KB	SDLT	128 KB
8 mm AIT-3	192 KB	SDLT320	128 KB
8 mm AIT-4	128 KB	SDTL600	128 KB
8 mm AIT-5	128 KB	DLT 20 GB	32 KB
8 mm Mammoth-2	192 KB	DST	1168 KB
DFT-2	384 KB	DTF	384 KB
DST (NT)	992 KB	4mm DAT160	32 KB
Optical	32 KB	SAIT - 1	64 KB
magnetic	32 KB	TS 1120	384 KB
file	32 KB	TZ85	32 KB
adv_file	128 KB	TZ86	32 KB
himt	32 KB	TZ87	32 KB
LTO Ultrium	64 KB	TZ88	32 KB

Table 7. Tape drive buffer size settings *(continued)*

Tape drive type	Buffer size (block size) setting	Tape drive type	Buffer size (block size) setting
LTO Ultrium-2	64 KB	TZ89	256 KB
LTO Ultrium-3	128 KB	TZ90	256 KB
LTO Ultrium-4	128 KB	TZS20	64 KB
LTO Accelis	128 KB	TKZ90	384 KB
QIC	32 KB	VHS	32 KB
QIC 13GB	32 KB	VXA	64 KB
SD3	256 KB	VXA - 2	64 KB
SLR	64 KB	VXA - 172	64 KB
T10000	256 KB	VXA - 320	64 KB
Travan 10	32 KB	CentricStor	256 KB
TK50	32 KB	generic	32 KB
TK70	32 KB	generic128	128 KB
4m DAT72	32 KB	generic256	256 KB

Fujitsu-Siemens NetWorker

EMC NetWorker on page 124 offers more information.

Configure HP OpenView Storage Data Protector

The tasks to configure HP OpenView Storage Data Protector software (based on version 5.5) are:

1. Import the NDMP server host on page 132
2. Create a media pool on page 133
3. Configure the tape library unit on page 133

Import the NDMP server host

1. In the **HP OpenView Storage Data Protector Manager**, switch to the clients context.
2. In the **Scoping** pane, expand **Data Protector Cell**.
3. Right-click **Clients**, and select **Import Client**.
4. In the **Import Client** window, type the name of the NDMP-Host Data Mover you want to import, and select **NDMP server**.

   Note: The HP Data Protector server should be able to map this Data Mover name to an IP address.

5. Click **Next**.
6. In the **Import NDMP Host** window, type the specific import parameters:
   - **NDMP Type**: Select VNX.
   - **Port**: Type `10000`
   - **Username**: Use the same username that you used in Assign a username and a password to a Data Mover on page 56.
   - **Password**: Use the same password that you used for the NDMP-Host Data Mover in Assign a username and a password to a Data Mover on page 56.
7. Click **Finish** to import the NDMP-Host Data Mover.

Create a media pool

Data Protector requires that the NDMP devices have dedicated media pools. Before you create an NDMP device, create a dedicated media pool to be used exclusively with this device as follows:

1. From the **Devices and Media** context, use the **Add Media Pool** wizard to create a media pool. Online help offers more information.

   Note: Data Protector Free pools are not supported with the Data Protector NDMP integration.

2. Format the media by browsing for the specific device, media pool, or library slot in the **Devices and Media** context.

   Note: Devices are automatically formatted when you select them for backup if the loose media allocation policy is defined for the media pool or the global variable InitOnLoosePolicy is set to 1. The value is 0.

3. Right-click the device, media pool, or library slot and select **Format**. Online help offers more information.

Configure the tape library unit

If your TLU robot is attached to the NDMP-Host Data Mover, you need to configure the TLU for the Data Mover. Use this procedure to configure a TLU whose robot is attached to the NDMP-Host Data Mover. If you are not using a TLU, then configure a stand-alone drive:

1. Switch to the **Devices and Media** context.

2. Expand **Environment**, right-click **Devices**, and select **Add Device**. The **Add Device** wizard appears.

3. Specify the device name and then click **Next**. You can type the description as follows:

   - **Device Type**: Select SCSI II Library.

     Note: To configure a stand-alone drive, select Standalone.

   - **Interface Type**: Select NDMP-VNX.
   - **Client**: Select the NDMP client system running the Data Protector software.
   - **NDMP server**: Select the Data Mover where the robotic device is attached.

4. Follow the wizard and type the library SCSI ID for the robotic device attached to the Data Mover.

   The SCSI address has this format:

   <SCSI_controller> <SCSI_ID> <SCSI_LUN>

   where:

   <SCSI_controller> = chain number of the controller

   <SCSI_ID> = SCSI ID (target) of the controller

   <SCSI_LUN> = logical unit number (LUN) of the controller

   Example:

   C1T0L0

   Set up the TLU on page 53 provides more information.

5. Specify the slots you want to use with the Data Mover and click **Next**.

6. Select the media type used in the library and click **Next**.

7. Click **Finish** to configure the device. Click **Yes** to configure the drives in the library.

8. Specify the drive name. You also have the option of typing the description for the drive. Click **Next**.

9. Follow the wizard and type the information about the drive's SCSI address.

   The SCSI address has this format:

   <SCSI_controller> <SCSI_ID> <SCSI_LUN>

   where:

   <SCSI_controller> = chain number of the controller

   <SCSI_ID> = SCSI ID (target) of the controller

   <SCSI_LUN> = LUN of the controller

   Example:

   C1T0L0

   Set up the TLU on page 53 provides more information.

   Note: Do not change the drive index number.

   ---

10. Click **Next**, and then specify the media and media pools information.

11. Click **Yes** to create another drive or **No** to finish creating drives for the library.

    The *HP OpenView Storage Data Protector Administrator's Guide* and online help provide more information on creating a drive.

IBM Tivoli Storage Manager

These restrictions apply when performing NDMP backups with VNX and IBM Tivoli Storage Manager:

- You must use a Windows Tivoli Storage Manager client to browse an NDMP backup. Windows is the only client type that allows you to perform a file-by-file restore with Tivoli Storage Manager.
- Supports only two-way backup.
- Does not support environment variables.
- You cannot run concurrent NDMP backups on a Data Mover where the target backup devices are DLT drives on the same SCSI bus, and are attached to an IBM 2108 SAN Data Gateway (all models) or an ADIC Pathlight 5000. Tivoli Storage Manager supports concurrent NDMP backups of Data Movers to direct SCSI-attached DLT drives.

Configure IBM Tivoli Storage Manager

This procedure is based on IBM Tivoli Storage Manager version 5.2, 5.3, or 5.4 and requires you to log in to the SUN TSM Admin client before you begin.

Perform the following steps to configure IBM Tivoli Storage Manager:

1. To define a library, use the command syntax:

   $ **define library** *<library_name>* **libtype=scsi shared=no**

   where:

   <library_name> = name you want to assign the library

   Example:

   To define a library, type:

   $ define library ndmplib libtype=scsi shared=no

2. To define the device class, use the command syntax:

   $ **define devclass** *<device_class>* **library=***<library_name>* **devtype =nas mountretention=0 estcapacity=** *<media_capacity>*

   where:

   <device_class> = name you want to assign the device class.

   <library_name> = library name you assigned in step 1.

   <media_capacity> = estimated capacity of the media. This value depends upon a tape's capacity.

Note: For NDMP operations, the mountretention value must be 0.

Example:

To define the device class, type:

```
$ define devclass nasclass library=ndmplib devtype=nas mountretention=0
estcapacity=70g
```

3. To define a tape storage pool, use this command syntax:

   $ **define stgpool** *<storage_pool>* *<device_class>* **maxscratch=30 dataformat=celerradump**

   where:

   <storage_pool> = name you want to assign the storage pool.

   <device_class> = device class name you assigned in step 2.

   Example:

   To define a tape storage pool, type:

   ```
   $ define stgpool naspool nasclass maxscratch=30 dataformat=celerradump
   ```

4. To create a policy domain for the NDMP-Host Data Mover, from the Tivoli Storage Manager software, use the command syntax:

   $ **define domain** *<policy_domain>* **description='***<policy_description>***'**

   where:

   <policy_domain> = name you want to assign the policy domain.

   <policy_description>= user-defined description of the NDMP policy.

   Example:

   To create a policy domain for the NDMP-Host Data Mover, type:

   ```
   $ define domain nasdomain description='policy domain for NDMP-Host Data Mover'
   ```

5. To create a policy set in the domain you created in step 4, use the command syntax:

   $ **define policyset** *<policy_domain>* *<policy_set>*

   where:

   <policy_domain> = policy domain you assigned in step 4.

   <policy_set> = name you want to assign the policy set.

   Example:

   To create a policy set, type:

   ```
   $ define policyset nasdomain naspolicy
   ```

6. To define a management class, use this command syntax:

   $ **define mgmtclass** *<policy_domain> <policy_set> <management_class>*

   where:

   <policy_domain> = policy domain you assigned in step 4.

   <policy_set> = policy set name you assigned in step 5.

   <management_class> = management class name you specify.

   Example:

   To define a management class, type:

   $ define mgmtclass nasdomain naspolicy nasmgtclass

7. To assign the management class as the default for the policy set, use this command syntax:

   $ **assign defmgmtclass** *<policy_domain> <policy_set> <management_class>*

   where:

   <policy_domain> = policy domain you assigned in step 4.

   <policy_set> = policy set you assigned in step 5.

   <management_class> = name you want to assign the management class.

   Example:

   To assign the management class as the default for the policy set, type:

   $ assign defmgmtclass nasdomain naspolicy nasmgtclass

8. To define a backup pool and set the allowable TOC size, use this command syntax:

   $ **define volume backuppool** *<pool_name>* **f=***<size>*

   where:

   <pool_name> = name of the backup pool.

   <size> = maximum size in MB of a backup's TOC.

   Example:

   To define a backup pool and set the allowable TOC size, type:

   $ define volume backuppool bkpool f=1000

9. To define a backup policy group in the default management class, use this command syntax:

   $ **define copygroup** *<policy_domain> <policy_set> <management_class>* **destination=***<storage_pool>* **verexits=***<number_of_backup_versions>* **tocdestination=***<TOC_storage_pool>*

   where:

   <policy_domain> = domain policy you assigned in step 4.

<policy_set> = policy set you assigned in step 5.

<management_class> = management class you assigned in step 6.

<storage_pool> = storage pool you assigned in step 3.

<number_of_backup_revisions> = the number of backup versions you want to retain in your table of contents (index).

<TOC_storage_pool> = TOC destination.

Note: Ensure that the destination for the copy group is assigned the appropriate backup format (celerradump). Step 3 provides more information.

Example:

To define a backup policy group in the default management class, type:

```
$ define copygroup nasdomain naspolicy nasmgmtclass destination=naspool verexits=10
tocdestination=backuppool
```

10. To activate the policy set, use this command syntax:

    $ **activate policyset** *<policy_domain> <policy_set>*

    where:

    <policy_domain> = policy domain you assigned in step 4.

    <policy set> = policy set you assigned in step 5.

    Example:

    To activate the policy set, type:

    ```
    $ activate policyset nasdomain naspolicy
    ```

11. To register the Data Mover with the Tivoli Storage Manager server, use this command syntax:

    $ **register node** *<movername> <password>* **domain=***<policy_domain>* **type=nas**

    where:

    <movername> = name of the Data Mover you want to register. This server must be accessible from the Tivoli Storage Manager client for your setup to work properly.

    <password> = Data Mover password you assigned in Assign a username and a password to a Data Mover on page 56.

    <policy_domain> = policy domain you assigned in step 5.

    Example:

    To register the Data Mover with the Tivoli Storage Manager server, type:

    ```
    $ register node server_2 admin domain=nasdomain type=nas
    ```

12. To define each NDMP-Host Data Mover, use this command syntax:

    $ **define datamover** *<movername>* **type=nas hladdress=***<movername_IP_address>* **lladdress=10000 userid=user_name password=***<password>* **dataformat=celerradump**

    where:

    <movername> = Data Mover name you registered with the Tivoli Storage Manager server in step 11.

    Note: The Data Mover name must match the name you used in the previous step to register the Data Mover with the Tivoli Storage Manager server.

    <movername_IP_address> = Data Mover IP address.

    lladdress = NDMP port number, which must be 10000.

    userid = NDMP username you assigned in Assign a username and a password to a Data Mover on page 56.

    <password> = Data Mover password you assigned in Assign a username and a password to a Data Mover on page 56.

    Note: You must specify 10000 as the lladdress.

    Example:

    To define each NDMP-Host Data Mover, type:

    ```
    $ define datamover server_2 type=nas hladdress=192.0.0.2 lladdress=10000
    userid=user_name password=admin dataformat=celerradump
    ```

13. To define a path to a library, use this command syntax:

    $ **define path** *<movername>* *<library_name>* **srctype=datamover desttype =library device=***<device_name>*

    where:

    <movername> = name of the Data Mover connected to the TLU.

    <library_name> = library name you assigned in step 1.

    <device_name> = TLU device name used by VNX. This name is determined by the SCSI controller, SCSI ID, and SCSI LUN. Set up the TLU on page 53 explains how to obtain a device name.

    Example:

    To define a path to a library, type:

    ```
    $ define path server_2 ndmplib srctype=datamover desttype=library device=c2t0l0
    ```

14. To show slots to reveal the element numbers you require in the next step, use this command syntax:

$ **show slots** *<library_name>*

where:

<library_name> = library name you assigned in step 1.

Example:

To show slots to reveal the element numbers, type:

```
$ show slots ndmplib
```

15. To define the tape drives for NDMP operations and the paths to those drives, use this command syntax:

    $ **define drive** *<library_name> <drive_name>* **element=***<element_number>*

    where:

    <library_name> = library unit name you assigned in step 1.

    <drive_name> = drive name, which can be any name you want to assign.

    <element_number> = element addresses, which step 14 revealed.

    Examples:

    To define the tape drives for NDMP operations and the paths to those drives, type:

    ```
    $ define drive ndmplib eisen1 element=128
    ```

    ```
    $ define drive ndmplib eisen2 element=129
    ```

    Note: Repeat this step for each drive in your library.

16. To define a path for the drive, use this command syntax:

    $ **define path** *<movername> <drive_name>* **srctype=datamover desttype =drive library=***<library_name>* **device=***<device_name>*

    where:

    <movername> = name of the Data Mover connected to the TLU.

    <drive_name> = drive name you assigned in step 15.

    <library_name> = library name you assigned in step 1.

    <device_name> = device name used by VNX. Set up the TLU on page 53 provides more information on locating device names.

    Examples:

    To define a path for the drive, type:

    ```
    $ define path server_2 eisen1 srctype=datamover desttype=drive library=ndmplib
    device=c2t2l0
    ```

    ```
    $ define path server_2 eisen2 srctype=datamover desttype=drive library=ndmplib
    device=c2t3l0
    ```

Configure Syncsort Backup Express

Syncsort Backup Express configuration procedures are based on using VNX and Syncsort Backup Express version 2.3:

1. Configure the Data Mover on page 141
2. Configure the libraries attached to the Data Mover on page 141
3. Run an NDMP backup of a Data Mover on page 142

Configure the Data Mover

1. From the **Syncsort Backup Express** home page, select **Configure** ➤ **Enterprise**.
2. Select the enterprise, and then select the Node Group to which you want to add the Data Mover. Click **Add Node**.
3. When the page refreshes, type the **Logical Node Name**, **Resolvable Node Name** or **IP address**, **Client node**, **Port 10000**, **User Name**, and the **Password** that you created for the NDMP-Host Data Mover.
4. Click **TEST** to verify that the correct authentication with the Data Mover, and click **Apply**.

Configure the libraries attached to the Data Mover

1. Start the **Syncsort Backup Express Device Configuration Wizard** and provide the login information. Click **Next**.
2. Select the Data Mover to which the TLU is attached and click **Next**. When the system finishes querying the Data Mover for the list of libraries (status = Done), click **Next**.
3. Select the library you want to configure. Click **Next**. The **Add Jukebox** window appears.
4. Configure the requested parameters and click **Next**.
5. To configure the tape drives, select **Add Device**. Select the controller Path and specify the Device Bay number and Device name. Click **OK**.
6. Repeat step 5 for each tape drive you want to configure.
7. Click **Exit**.

Run an NDMP backup of a Data Mover

1. From the **Syncsort Backup Express** home page, select **BACKUP** and click **NDMP**. The **NDMP Backup** window appears.

2. Select the Enterprise. From the **Node Group** category select the Data Mover you want to back up. (This is the Data Mover you configured in Configure Data Movers on page 46.)

3. From the **Backup** menu, select **Define Source Options**. The **Define Source Options** window appears.

4. Set the desired environment variables for the selected backup job:

   Note: Syncsort Backup Express offers two ways to enable the EMC NVB feature:

   - Set the environment variable VBB=yes for the specific backup job.
   - Select vbb as the backup type for the Data Mover.

   This step shows how to enable VBB by using both methods, and explains the reasons for using each. Backup/restore with NVB on page 29 provides more details on the NVB feature.

   a. Use this procedure to set the desired environment variables, including the NVB variable, for the selected backup job.

      In the Additional NDMP Environment field of the **Define Source Options** window of the selected backup job, type the desired NDMP environment variables (comma separated).

      Example:

      SNAPSURE=y, VBB=y, HIST=y

      It is recommended to enable VBB by using the VBB variable (as shown above, with the other variables for the backup job) because specifying the backup method at the backup-job level makes it easier to manage different backup methods on the same Data Mover.

   b. Optionally, you can enable NVB at the Data Mover level, as follows: From the Syncsort Backup Express home page, select **Configure** and click **Enterprise**. Select the Enterprise, Node Group, and Data Mover name that you want to back up. Next, from the **Backup Type** list, select **vbb**. At this point, every NDMP backup job run with the selected Data Mover will be performed by using the NVB method.

      Use of this method makes it difficult to troubleshoot a situation where a VBB type is set at the Data Mover level, and a non-VBB-supported configuration is specified when defining a specific backup job. (Examples of a nonsupported configuration include backing up a subdirectory of the file system, or backing up a file system that is not read-only.)

5. Click **OK**. Select **File** and click **Save Job**.

6. Select **BACKUP** and click **Run**.

Configure Symantec NetBackup

The procedures for configuring Symantec NetBackup are based on Symantec NetBackup versions 5.0, 5.1, 6.0, and 6.5. They include steps for configuring backup devices in Media Manager, the Symantec NetBackup graphical user interface. The procedures for configuring devices for Media Manager are performed from the NetBackup Administration Console. When configuring Symantec NetBackup for NDMP operations, you must:

1. Configure the system devices.
2. Specify the information that Media Manager, the Symantec NetBackup graphical user interface, requires to control the devices.

Each procedure highlights some of the VNX-specific information, but does not cover the configuration in detail. You cannot use the Device Configuration wizard to configure devices for NDMP. You can use the Device Monitor to verify that the devices were set up properly. A successfully configured device is in an “Up" state.

Also, ensure that the HIST environmental variable is enabled before performing a backup, and enable DAR before a restore to make DAR available for file recovery.

To configure Symantec NetBackup:

1. Enable host access to a Data Mover on page 143
2. Configure a robot on page 148
3. Add a robot on page 150
4. Update the volume database on page 152
5. Add a drive on page 152
6. Add a storage unit on page 153
7. Create a backup policy on page 153
8. Specify a backup type on page 154

Enable host access to a Data Mover

To perform NDMP backups, NetBackup must have access to the Data Mover as the NDMP host. To authorize host access to the Data Mover, perform the following procedure from the NetBackup Administration Console on your NetBackup for NDMP master server.

You can enable host access to the Data Mover by:

- Using the CLI on page 144

- Using the NetBackup Administration Console on page 145

Note: Use the NetBackup Administration console for versions 6.0 and later.

Using the CLI

Perform the following steps to enable host access to a Data Mover using the CLI:

1. To enable host access to a Data Mover, use the command syntax:

   set_ndmp_attr -insert -auth *<movername>* **user_name**

   where:

   <movername> = customer-assigned name for the NDMP-Host Data Mover connected to the TLU. Symantec NetBackup NDMP accesses a Data Mover by its NIS, DNS (Domain Name Server), or /etc/hosts name.

   user_name = username under which Symantec NetBackup NDMP accesses the Data Mover.

   For Windows NT/2000 systems, the set_ndmp_attr command is located in the <install_path>\Volmgr\bin directory. By default, the location is C:\Program Files\Veritas\Volmgr\ bin.

   For a UNIX system, the command is located in the /usr/openv/volmgr/bin/ directory.

   Example:

   On Windows NT/2000, where the Data Mover name is server_2:

   <install_path>\Volmgr\bin\set_ndmp_attr -insert -auth server_2 user_name

   On UNIX, where the Data Mover name is server_2:

   /usr/openv/volmgr/bin/set_ndmp_attr -insert -auth server_2 user_name

   Output:

   ```
   Password: password
   Re-enter Password: *****
   Password for host "server_2" user "user_name" Re-enter Password
   forhost "server_2" user "user_name":
   Re-enter Password for host "server_2" user "user_name":
   Inserted host "server_2"
   ```

   where:

   Password = the same password set for the NDMP-Host Data Mover. Assign a username and a password to a Data Mover on page 56 provides more information.

   Note: VNX accesses a Data Mover by its internal name (movername), such as server_2, but this field can be anything the customer assigns.

2. If you are performing a Data Mover-to-Data Mover three-way backup, repeat step 1 for the Data Mover without an attached tape drive. After this step is completed, the Data Mover is listed in the GUI under NDMP hosts.

Using the NetBackup Administration Console

1. Start the **NetBackup Administration Console** for the NDMP server:
   - In a Windows environment: From the **Start** menu, select **Programs ➤ Symantec NetBackup ➤ NetBackup Administration Console**.
   - In a UNIX environment, type:

     ```
     /usr/openv/netbackup/bin/jnbSA &
     ```

2. From **Media and Device Management** in the left pane, select **Credentials** and click **NDMP Hosts**.

3. From the **Actions** menu, select **New** and click **NDMP Host**. The **NDMP host name** dialog box appears.

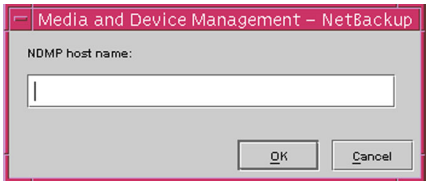

4. Type the name of the NDMP server to back up and click **OK**. The **New NDMP Host** dialog box appears.

Note: The NDMP host name is case-sensitive.

---

5. Under **NDMP Host Credentials**, select the appropriate option from the list given below and type the username and password that NetBackup uses to access the NDMP host:
   - For the **Use Global NDMP credentials for this NDMP host** option:
     a. Click **Host Properties** and select **Master Server**.
     b. Double-click the **Master Server** option from the right pane. The **Properties** dialog box appears.
     c. From **Properties**, select **NDMP**. The **NDMP Global Credentials** dialog box appears.
     d. In the **NDMP Global Credentials** dialog box, type the username and password in the **Username**, **Password**, and **Confirm password** fields, respectively. Use the same username and password that NetBackup uses to access the NDMP server.
     e. Click **Apply** and then click **OK**.

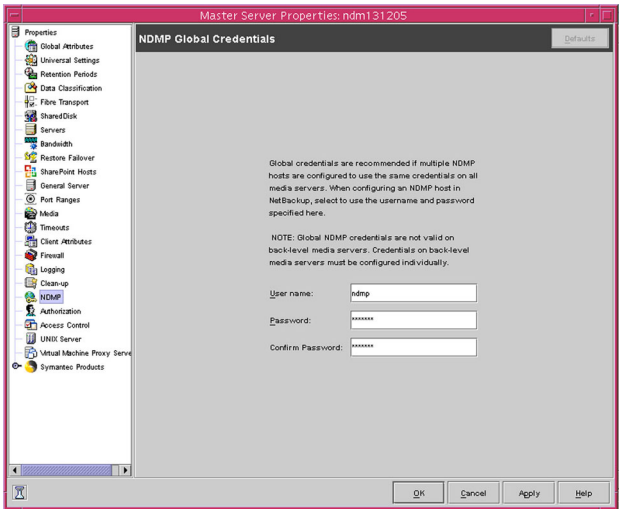

- For the **Use the following credentials for this NDMP host on all media servers** option:

  a. Type the username and password in the **Username**, **Password**, and **Confirm pass word** fields, respectively. Use the same username and password that NetBackup uses to access the NDMP server.

  Note: The user must have permission to run NDMP commands.

  b. Click **OK**.

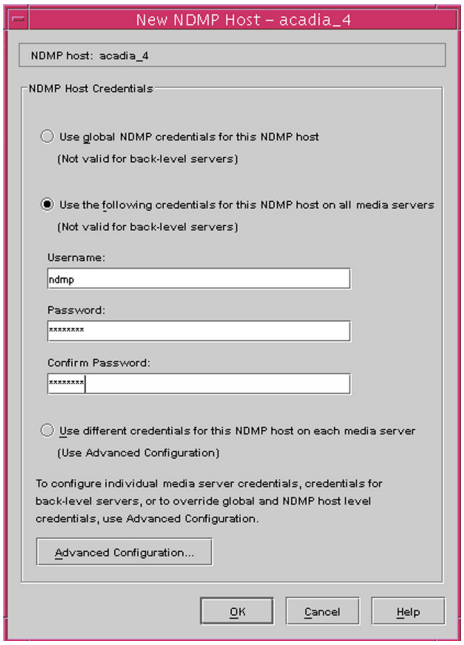

- For the **Use the following credentials for this NDMP host on all media servers** option:

  a. Click **Advanced Configuration**. The **Advanced NDMP Credentials** dialog box appears.

  b. Click **Add** to open the **Add Credentials** dialog box.

  c. Type the username and password in the **Username**, **Password**, and **Confirm pass word** fields, respectively. Use the same username and password that NetBackup uses to access the NDMP server.

  Note: The user must have permission to run NDMP commands.

  d. Click **OK**.

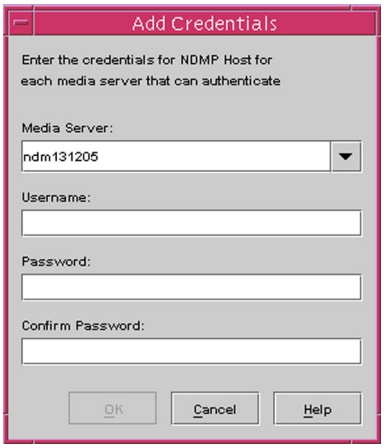

6. Verify access between the NDMP host and the NDMP client by using this command syntax:

   tpautoconf -verify *<movername>*

   where:

   <movername> = name of the Data Mover

   Example:

   Windows: **<install_path> \Volmgr\bin\tpautoconf -verify server_2**

   UNIX: **/usr/openv/volmgr/bin/ tpautoconf -verify server_2**

   Output:

```
# ./tpautoconf -verify server_2
Connecting to host "server_2" as user "ndmp"...
Waiting for connect notification message...
Opening session--attempting with NDMP protocol version 4...
Opening session--successful with NDMP protocol version 4
  host supports TEXT authentication
  host supports MD5 authentication
Getting MD5 challenge from host...
Logging in using MD5 method...
Host info is:
  host name "server_2"
  os type "DartOS"
  os version "EMC Celerra File Server.T.5.6.45.4"
  host id "abc1997"
Login was successful
Host supports LOCAL backup/restore
Host supports 3-way backup/restore
```

Repeat this procedure for each NDMP host that NetBackup backs up, and also for a Data Mover-to-Data Mover three-way backup, for the Data Mover without an attached tape drive. After this step is completed, the Data Mover is listed in the GUI under NDMP hosts.

Configure a robot

CAUTION Do not perform this step for versions 6.0 and later.

Specify the data for the robotic control that is attached to the NDMP-Host Data Mover. For NDMP client-attached robotic control information, consult the NetBackup documentation.

When configuring a robot that attaches directly to the NDMP-Host Data Mover, you must specify these information:

- Device name of the robot
- SCSI controller number for the robot
- SCSI ID for the robot
- SCSI LUN for the robot

Perform the following steps to configure a robot:

1. To specify the robotic control, use this command syntax:

   ```
   $ set_ndmp_attr -insert -robot <movername> <device_name> <SCSI_
   controller> <SCSI_ID> <SCSI_LUN>
   ```

   where:

   <movername> = name of the NDMP-Host Data Mover

   <device_name> = robotic device name (Set up the TLU on page 53 explains how to obtain a device name)

   <SCSI_controller> = controller number

<SCSI_ID> = SCSI ID for the robotic control

<SCSI_LUN> = LUN for the robotic control

Windows example:

```
<install_path> \Volmgr\bin\set_ndmp_attr -insert -robot server_2 c1t0l0 1 0 0
```

UNIX example:

```
/usr/openv/volmgr/bin/set_ndmp_attr -insert -robot server_2 c1t0l0 1 0 0
```

Output:

```
Inserted robot host "server_2:c1t010" device "1"
```

Note: If you make an error configuring the robot, delete the configuration by using this command syntax:

```
$ set_ndmp_attr -delete -robot <movername> <devicename> <SCSI
_controller> <SCSI_Id> <SCSI_Lun>
```

Example:

```
$ set_ndmp_attr -delete -robot server_2 c1t010
```

Output:

```
Deleted robot host "server_2" device "c1t0l0"
```

---

2. To verify that the authentication record is created and view the record for accuracy, type:

   ```
   set_ndmp_attr -list
   ```

   Windows example:

   ```
   <install_path> \Volmgr\bin\set_ndmp_attr -list
   ```

   UNIX example:

   ```
   /usr/openv/volmgr/bin/set_ndmp_attr -list
   ```

   Output:

   ```
   Record Type: Authentication
   NDMP server: server_2
   Username: user_name
   Password: <registered>
   Record Type: Robotic Information
   NDMP server: server_2
   Device: c1t0l0
   Controller: 1
   SCSI Id: 0
   SCSI Lun: 0
   ```

3. To verify access between the NDMP host and the NDMP client, use this command syntax:

   set_ndmp_attr -verify *<movername>*

   where:

   <movername> = name of the Data Mover

Windows example:

```
<install_path> \Volmgr\bin\set_ndmp_attr -verify server_2
```

UNIX example:

```
# /usr/openv/volmgr/bin/set_ndmp_attr -verify server_2
```

Output:

```
Connecting to host "server_2" as user "user_name"...
Waiting for connect notification message...
Opening session with NDMP protocol version 2..
Host info is:
host name "Server"
os type "SNAS"
os version "5.4"
host id "abc1997"
Logging in using TEXT method...
Login was successful
```

Add a robot

Use these steps to add a robot from the **NetBackup Administration Console**, for NDMP master server:

Note: For NetBackup versions prior to 6.0, follow the steps in Configure a robot on page 148 to add a robot.

1. From **Media and Service Management** in the left pane, click **Devices**.
2. From the **Actions** menu, select **New** and click **Robot**. The **Add a New Robot** dialog box appears.
3. In the **Add Robot** dialog box, type the appropriate information in the fields listed below (use the online help or Symantec documentation for assistance):
   - **Media Manager host**: Specify the host that manages the EMM database (by default, this is the NetBackup master server).
   - **Device host**: Use the drop-down list to select the NetBackup media server.
   - **Robot type**: Specify the robot type.
   - **Robot name**: Specify the robot name.
   - **Robot control**: Select the robot control that is attached to an NDMP host option.
   - **Robot device path**: Type the device name of the robot by using this command syntax:

     <SCSI_controller> <SCSI_Id> <SCSI_Lun>

     where:

     <SCSI_controller> = chain number of the controller

<SCSI_Id> = SCSI ID of the controller, also called target

<SCSI_Lun> = LUN of the controller

Example:

```
c1t0l0
```

Note: Do not include the NDMP host name as part of the device path. Set up the TLU on page 53 provides more information on determining the robot device name.

- **NDMP host name**: Type the name of the NDMP host to which the robot is attached.
- **Bus, Target, and LUN**: Use the values that you used specified for the robot device path. For example, use 1, 0, and 0 for bus, target, and LUN, respectively, from the previous example. The bus, target, and LUN values are 0 by default.

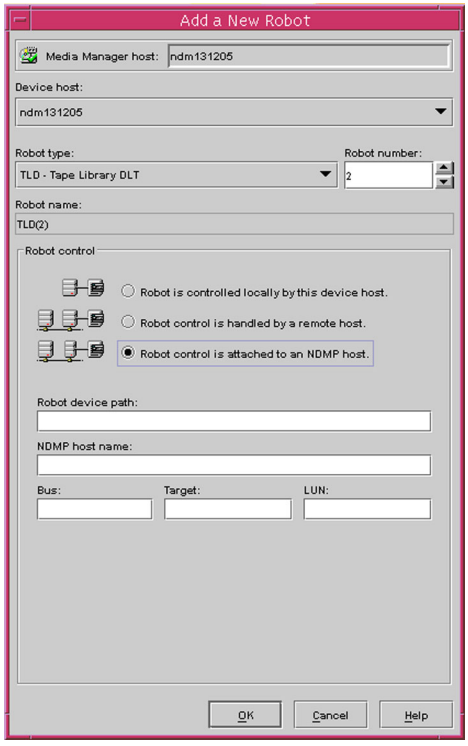

4. Click **OK** in the **Add New Robot** dialog box. A pop-up window prompts you to restart the device manager service.

5. Click **No** to continue configuring other devices before restarting the device manager.

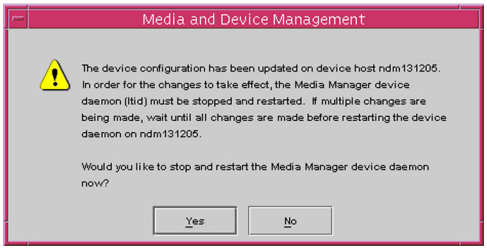

CAUTION Restarting this service can abort backups, archives, or restores in progress.

Update the volume database

1. Insert the volumes into the robotic library.
2. Under **Media and Device Management** in the **NetBackup Administration Console**, click **Media ➤ Robots**.
3. Select the robotic library where you inserted the volumes.
4. Click **Actions ➤ Inventory Robot**.
5. In the **Inventory Type** section, select **Update Volume Configuration**.
6. Click **Start** to start the inventory and update the database.

Add a drive

1. Under **Media and Device Management** in the **NetBackup Administration Console**, click **Devices**.
2. From the **Actions** menu, select **New ➤ New Drive**. The **Add Drive** dialog box appears.
3. Complete the dialog box fields. Consult online help for more information.

   For the **Device name** field, use this syntax for a drive directly attached to the NDMP-Host Data Mover:

   <movername>:<device_file>

   where:

   <movername> = name of the Data Mover

   <device_file> = name of the file for the drive

   Example:

```
server_2:c1t1l0
```

4. When you complete the dialog box fields, click **OK**. A message appears asking whether you want to restart the Media Manager.

5. Click **Yes.**

Add a storage unit

1. In **NetBackup Management** in the **NetBackup Administration Console**, click **Storage Units**.

2. From the **Action** menu, select **New ➤ New Storage Unit**. The **Add New Storage Unit** dialog box appears.

3. In **Storage Unit Type**, click the arrow and select **NDMP** from the list.

4. Type the appropriate information in the dialog box fields.

   Note: Consult online help for more information.

   In NDMP host, type the name of the Data Mover.

   Example:

   ```
   server_2
   ```

5. When you complete, click **OK**.

Create a backup policy

1. Specify these attributes:

   - **Policy Type**: NDMP
   - **Policy Storage Unit**: Name of the storage unit if the NDMP host has more than one
   - **For a three-way backup**: Storage unit that you defined for the NDMP host with an attached tape drive

2. Configure the client or clients as follows:

   - **Hostname**: Name of the NDMP-Host Data Mover
   - **Hardware and operating system**: `NDMP NDMP`

3. Configure the files so the file list contains directories that reflect the directories on the NDMP Host. NDMP does not support these file list capabilities:

- Wildcards in pathnames
- Individual filenames
- Exclude list

4. Verify the configuration by running a backup policy.

Specify a backup type

By default, Symantec uses the dump backup type. You can change the backup type from dump to tar or vbb by specifying a variable in the backup policy's file list as follows:

1. In the **NetBackup Administration Console**, expand **Master Server ➤ NetBackup Management ➤ Policies**.
2. In the middle pane, double-click the policy you want to edit. The **Change Policy** dialog box appears.
3. Click the **Backup Selections** tab, and select the path or directory where you want to use the vbb backup type.
4. Click **Insert**. The **Insert File** dialog box appears.
5. In Pathname or directive, define the variable by typing:

```
SET type = vbb
```

6. Press **Enter** and click **OK**. The new variable should now appear in the middle pane of the **NetBackup Administration Console**, above the directories or files to which it applies.

Note: In the file list, the variable definition must appear before the directories or files to which it applies.

NetBackup for NDMP System Administrator's Guide for UNIX and Windows provides more information on using variables.

Configure Symantec Backup Exec

Before configuring the NDMP server with the Backup Exec server, ensure that the following has been installed on the Backup Exec server:

- Windows Servers or Windows XP
- Backup Exec for Windows Servers

Note: While installing Backup Exec, make sure that you select the NDMP option, which is an add-on component.

To configure Symantec Backup Exec:

1. Add the NDMP server to the Backup Exec on page 155
2. Inventory the newly added TLU on page 156
3. Create a backup job on page 158
4. Create a restore job on page 160

Add the NDMP server to the Backup Exec

To configure Backup Exec in order to use the NDMP option, add the NDMP server to Backup Exec. If any storage devices are attached to the NDMP server, the Backup Exec locates them automatically and adds them to the list of devices.

Note: Before starting the installation process, edit the host file in the Windows machine that hosts the Backup Exec application to include the IP address and name of the NDMP servers.

To add an NDMP server to Backup Exec:

1. On the navigation bar, click the **Devices** tab.
2. In the task pane, under **NDMP Tasks**, click **Add NDMP Server**. The **Add NDMP Server** dialog box appears.
3. Add the following information in the **Add NDMP Server** dialog box:
   - **NDMP Server**: Type the name of the NDMP server.
   - **Port**: Type the appropriate port number (10000) that is used for communications between the Backup Exec server and the NDMP server.
   - **Logon Account**: Select the appropriate logon account for the NDMP server. To create a new account, click the button next to the **Logon Account** field and then click **New** in the **logon Account Selection** dialog box. Type the same username and password that you used for the NDMP-Host Data Mover.
   - **NDMP server supports ICMP ping operations**: Select this option if the NDMP server supports ICMP ping operations. It ensures that the Backup Exec can use ping to locate the NDMP server ping operations, by default.

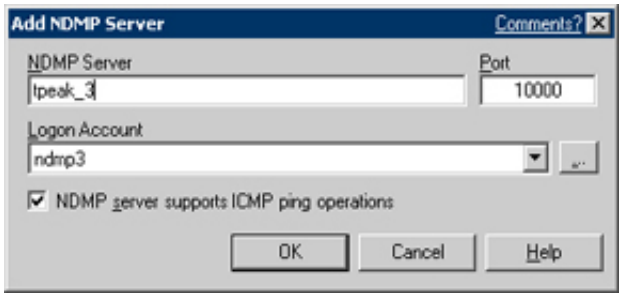

4. Click **OK** to add the NDMP server. A warning message box appears. Click **Yes** to restart all the Backup Exec services. The **Backup Exec Services Manager** dialog box appears.

5. Click **Restart all services** to restart the Backup Exec services.

6. Click **OK** when the message **Processing services completed** appears in the **Restarting Backups Exec Services** window.

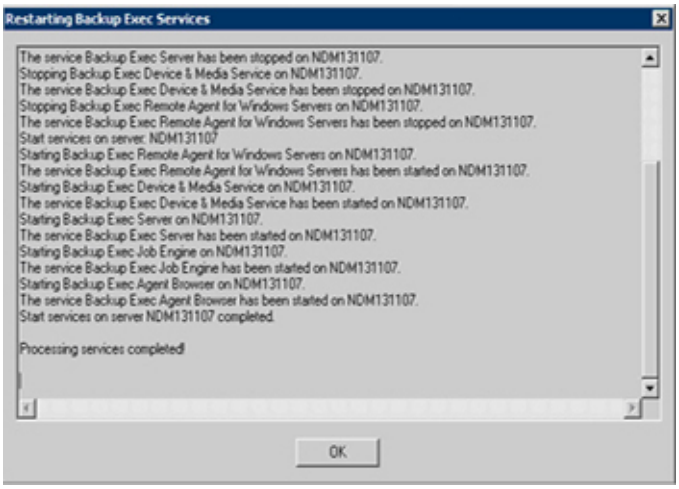

Note: The NDMP server and the storage devices that are directly attached to the Backup Exec appear in the Devices pane. Depending on the number of storage devices, it can take a while for all the devices to be displayed. You need to refresh the GUI to view all the services and devices that are added.

Inventory the newly added TLU

Following are steps:

1. From the navigation bar, click the **Devices** tab.

2. Select **Robotic Libraries** and right-click the required robot.

3. Select **Inventory**. The **Inventory Job Properties** dialog box appears.

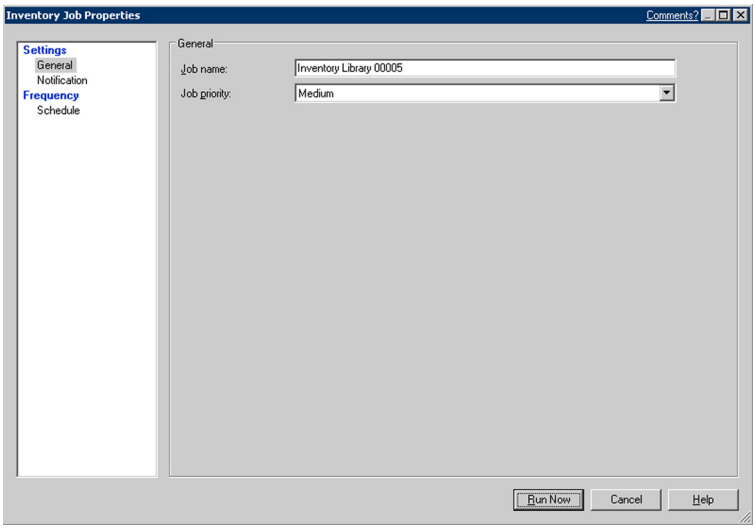

4. Click **Run Now** to start the inventory. To view and monitor the inventory status, click the **Job Monitor** tab from the navigation bar. The **Job Monitor** window appears.

Note: When the inventory job is successfully completed, it is moved to the Job History section in the Job Monitor window.

---

5. It is recommended to erase and label the tapes. To erase tapes, select **Robotic Library** and click **Slots**. A list of tapes is displayed.

6. Select all the tapes and then right-click to select **Quick Erase**. A warning box appears. Click **OK** to start the Quick Erase process.

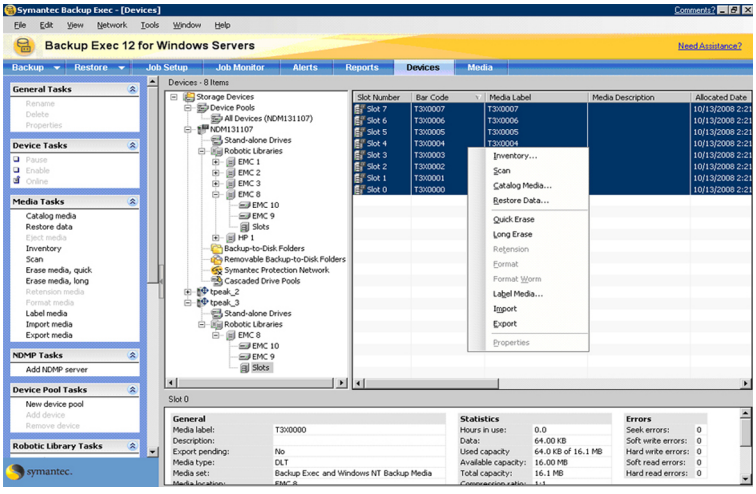

Note: To erase all the tapes together, click Yes to All in the Symantec Backup Exec dialog box.

---

7. In the **Erase Job Properties** dialog box, click **Run Now**. To view and monitor the status, click the **Job Monitor** tab on the navigation bar. The **Job Monitor** window appears.

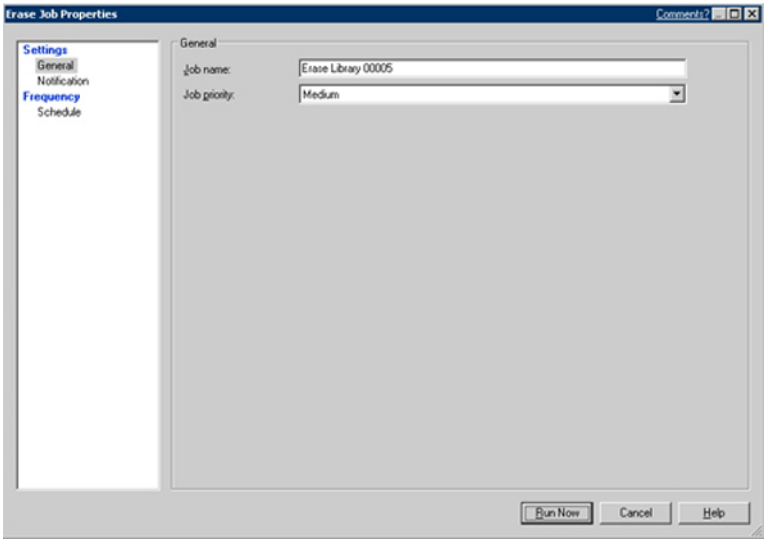

8. When the job is completed, the tapes are listed under **Scratch Media Set** and are ready for use. To view the tapes, click the **Media** tab on the navigation bar. From the left panel, select **Media Sets** and then click **Scratch Media set**.

9. To label the tapes, select **Robotic Library** and click **Slots**. A list of tapes is displayed.

10. Select all the tapes and right-click to select **Label Media**. A warning box appears. Click **OK** to start the labeling process for the first media. Click **OK** to label each media individually. To view and monitor the status of the labeling job, click the **Job Monitor** tab on the navigation bar. The **Job Monitor** window appears.

Create a backup job

Perform the following steps to create a backup job:

1. From the navigation bar, click the **Backup** tab.

2. Select **New Backup Job**. The **Backup Job Properties** window appears.

3. From the left task panel, select **Source** and click **Selections**.

4. To back up a file system, select **All Resources ➤ NDMP Servers** and then select the appropriate NDMP server (where the file system is located). All the file systems to be backed up (except the file systems on which NDMP2D TLUs are built) are displayed. If the **Logon Account Selection** dialog box appears when the NDMP server is selected, select or create the correct account name.

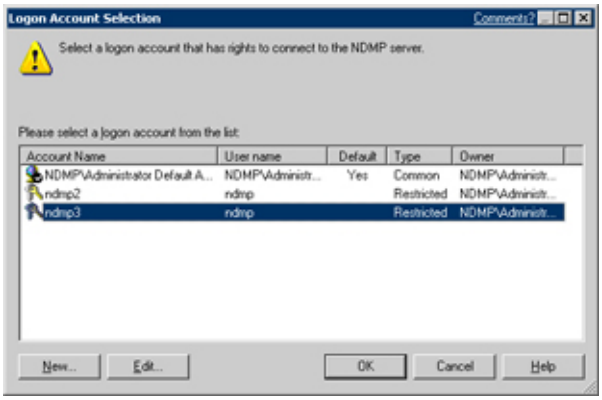

5. To select the required file system for backup, select the checkbox next to it.

6. To select the library to backup the data, select **Devices and Media** from the **Destination** tab.

7. From the **Device** list, select a storage device (for example, the library or one of its drives) where the data needs to be backed up. If you select a storage device that is attached to the Data Mover on which the source file system is mounted it would be a two-way backup, whereas selecting a storage device attached to a remote Data Mover would be a three-way backup.

8. Under **Settings**, click **NDMP** and select the following options:

   a. From the **Backup method** drop-down list, select the backup method. Level 0 is for full backup, whereas levels 1 through 9 are intended for incremental backups.

   b. From the **Backup type** drop-down list, select **Dump** or **VBB**.

   c. Select the checkbox next to the **Backup with integrated checkpoints (VNX SnapSure)** option. This option automatically creates a checkpoint for the file system that is to be backed up.

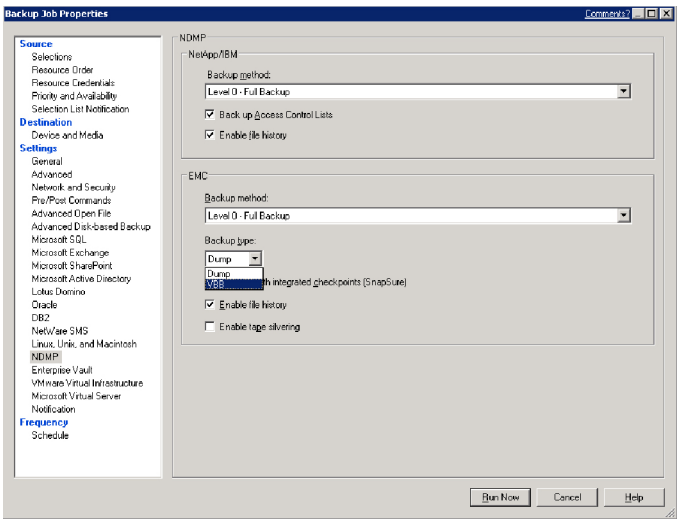

9. Click **Run Now**. A confirmation window appears. Click **OK** to start the backup operation.

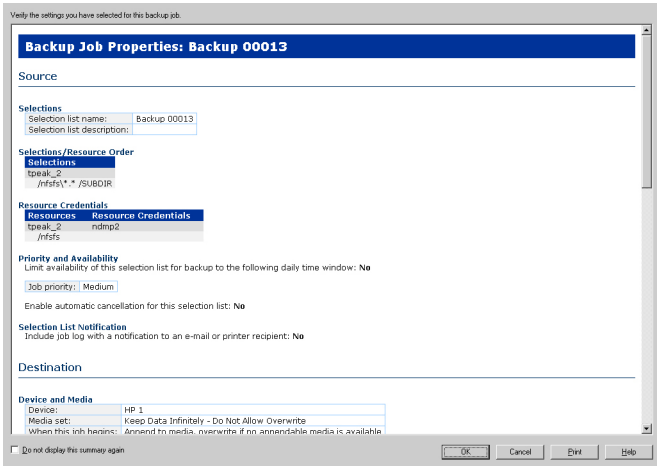

Create a restore job

Perform the following steps to create a restore job:

1. From the **Restore** tab, click **New Restore Job**. The **Restore Job Properties** window appears.

2. Select **Source** and click **Selections**. The files that are to be restored from their respective NDMP servers are now available.

3. Under **All Resources**, select the appropriate NDMP server to restore files from a backup image. A list of all datasets (file systems and directories) that were previously backed up is displayed.

4. Click the required dataset and then select the checkbox next to the required backup image to perform a full restore. You can expand the backup image to select the required files and directories in order to perform a partial restore.

5. Under **Source**, click **Device** and then select the storage device (the library or one of its drives) from where the data needs to be restored.

6. To enable an out-of-place restore, under **Destination**, select **File Redirection** and then select the checkbox next to **Redirect file sets**. In file redirection, the original data is not replaced by the restored data. If the **File Redirection** option is not selected, an In-place restore takes place.

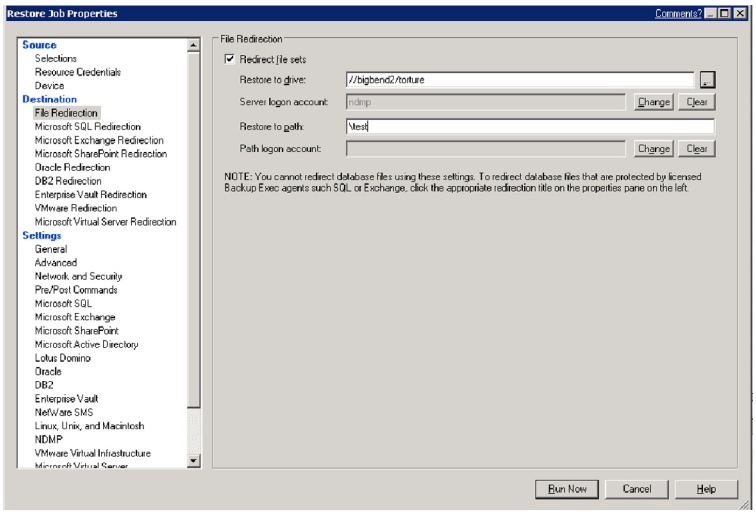

7. In the **Restore to drive** field, specify the NDMP server and the file system by clicking the button next to it. You can also type the drive name.

8. In the **Restore to path** field, specify the out-of-place location within the target file system.

9. Under **Settings**, click **NDMP** and select the checkbox next to the **Enable Direct Access Recovery** option. This option is selected by default to allow the direct access recovery of selected files and directories without restoring the backup image as a whole.

10. Click **Run Now**. A confirmation window appears. Click **OK** to start the restore operation.

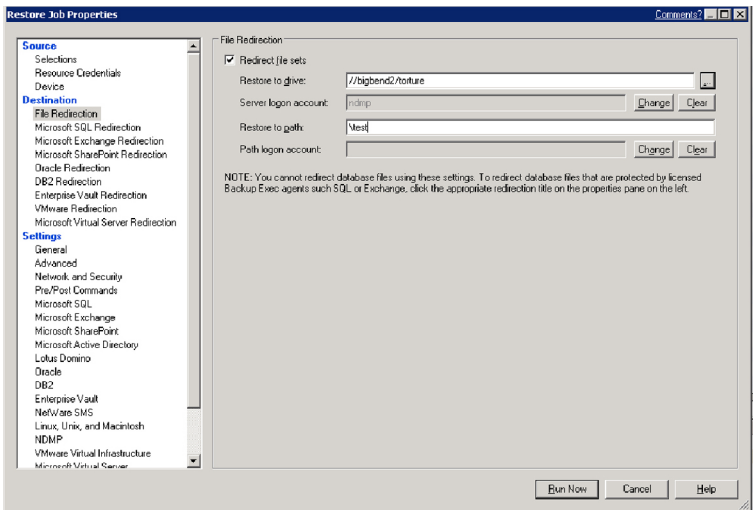
Restore Job Properties
Comments?
Source
Selections
Resource Credentials
Device
Destination
File Redirection
Microsoft SQL Redirection
Microsoft Exchange Redirection
Microsoft SharePoint Redirection
Oracle Redirection
DB2 Redirection
Enterprise Vault Redirection
VMware Redirection
Microsoft Virtual Server Redirection
Settings
General
Advanced
Network and Security
Pre/Post Commands
Microsoft SQL
Microsoft Exchange
Microsoft SharePoint
Microsoft Active Directory
Lotus Domino
Oracle
DB2
Enterprise Vault
NetWare SMS
Linux, Unix, and Macintosh
NDMP
VMware Virtual Infrastructure
File Redirection
Redirect file sets
Restore to drive:
//bigbend2/torture
Server logon account:
ndmp
Change
Clear
Restore to path:
\test
Path logon account:
Change
Clear
NOTE: You cannot redirect database files using these settings. To redirect database files that are protected by licensed Backup Exec agents such SQL or Exchange, click the appropriate redirection title on the properties pane on the left.
Run Now
Cancel
Help

Appendix B

Configure CNS and NS for NDMP

This appendix illustrates the VNX fibre HBA configuration option as well as Network Server (NS) fibre HBA configuration option for NDMP. Figure 12 on page 163 shows the supported CNS fibre HBA configuration option with 510/514 Data Movers. Figure 13 on page 164 shows the supported NS fibre HBA configuration option.

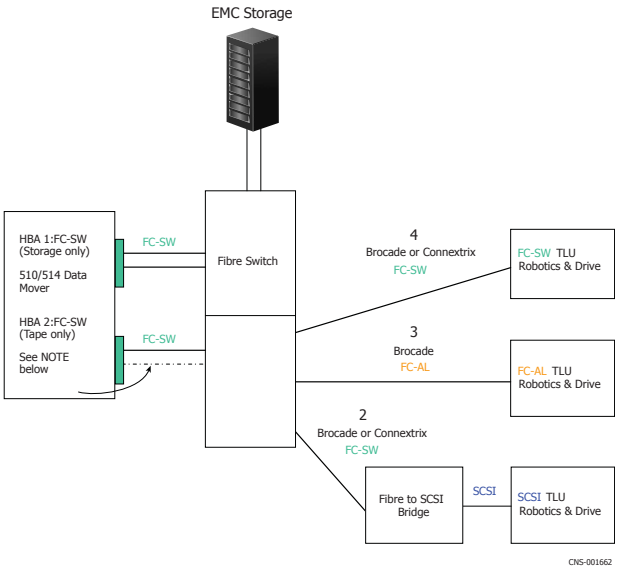

Figure 12. CNS FC-SW HBA configuration with 510/514 Data Movers

In Figure 13 on page 164, the two fibre HBAs are used for tape drive connectivity.

Note:

- The LP8000 HBA is a dual channel Performance module. It is used for both storage and tape drive connectivity.
- The LP9000 HBA is available in both dual and single channel Performance module, but only the single channel module is used for tape drive connectivity.
- The single channel Performance module, LP9002L-F2, is used to connect TLUs and tape drives.
- The dual channel Performance module, LP9002DC, is used to connect storage.

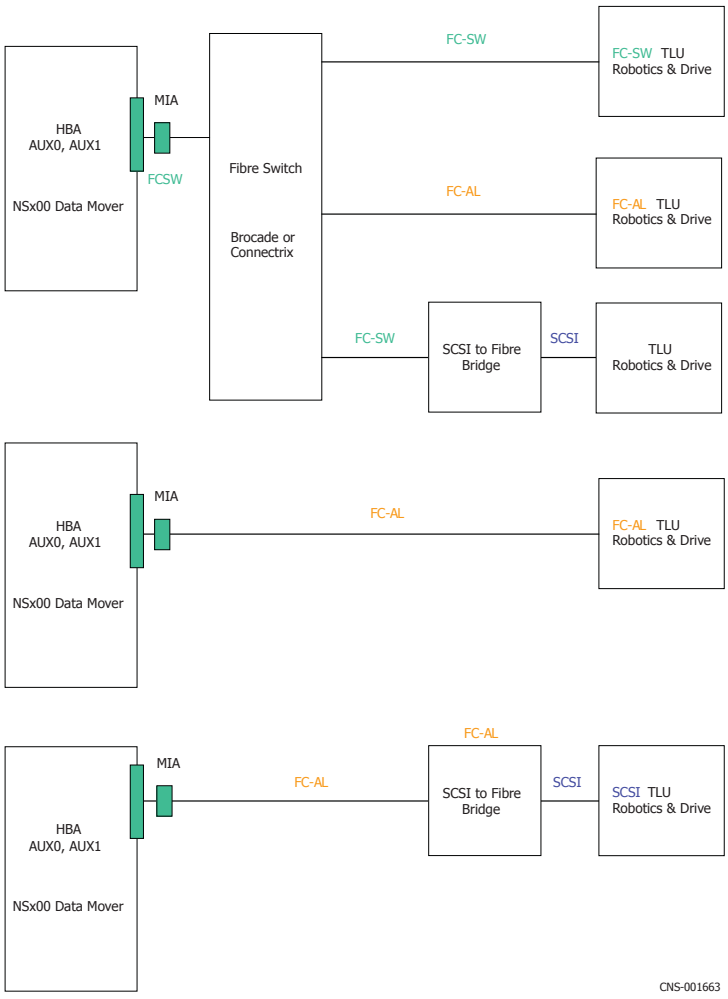

Figure 13. NS Performance-SW and FC-AL HBA configurations

Glossary

A

ASCII (American Standard Code for Information Interchange)
7-bit (one byte) character-encoding scheme for representing text and controlling printers and computer communication based on the ordering of the English alphabet. US-ASCII was the most commonly used character encoding on the Internet until it was recently surpassed by UTF-8. ASCII was incorporated into the Unicode character set as the first 128 symbols, so the ASCII characters have the same numeric codes in both sets.

automated backup and restore
Procedure in which a remote backup server running NDMP-compliant backup software backs up data to a tape library unit (TLU).

See also *local backup* and *remote backup*.

B

backup and restore
Technique for ensuring file system integrity and security by backing up to tape or disk a file system or incremental changes to a file system. This data can be restored later.

C

character set
Set of characters used by a language (ANSI/ASCII or Unicode). The character set affects file format (how the software stores and transmits data) and string processing (the logic with which text is manipulated).

checkpoint
Point-in-time, logical image of a PFS. A checkpoint is a file system and is also referred to as a checkpoint file system or an EMC SnapSure™ file system.

See also *Production File System.*

cumulative incremental backup
Backup that copies all data modified since the last full backup. To restore data when cumulative incremental backups are in use, you need only the latest full and cumulative incremental backups.

See also *differential backup*, *differential incremental backup*, *full backup*, and *incremental backup*.

D

differential backup
Backup that copies all data modified since the last full backup. To restore data when differential backups are in use, you need only the latest full backup and the latest differential backups.

See also *cumulative incremental backup*, *differential incremental backup*, *full backup*, and *incremental backup*.

differential incremental backup
Backup that copies data objects modified since the last full backup or incremental backup. To restore data when differential incremental backups are in use, the newest full backup and all subsequent differential backups are required.

See also *cumulative incremental backup*, *full backup*, and *incremental backup*.

direct access restore (DAR)
Data recovery operation that allows the NDMP client to directly access backed-up data anywhere in a tape set without having to traverse the tape set sequentially.

dump
Backup format in PAX that traverses a file tree in mixed width-first and depth-first order.

dynamic drive sharing (DDS)
Feature of NDMP-compliant backup software that allows NDMP servers to share tape drives.

F

filtering for NDMP
Backup feature that allows administrators to specify using environmental variables in the qualified backup vendor's software, which directories and files within a file system to exclude from the backup.

full backup
Backup that copies a defined set of data, regardless of whether it has changed since the last backup. A full backup is the basis from which incremental backups are taken.

See also *cumulative incremental backup*, *differential incremental backup*, and *incremental backup*.

H

heterogeneous three-way backup
Three-way NDMP backup involving a Data Mover, which serves as the interface to the file system or primary storage system, and a qualified third-party NDMP server, which serves as the interface to the tape devices or other secondary storage.

See also three-way backup.

I

I18N
An abbreviation for internationalization (I + 18 letters + N).

See internationalization.

incremental backup
Term that encompasses both cumulative incremental backups and differential incremental backups. Any backup that copies only data objects modified since the time of some previous backup.

See also *cumulative incremental backup*, *differential incremental backup*, and *full backup*.

internationalization (I18N)
The process of generalizing software such that it can be made to handle various languages and locales.

L

L10N
An abbreviation for localization (L+10 letters+N).

See localization.

local backup
Procedure that backs up a file system, or incremental changes to a file system from a Data Mover to a locally attached TLU.

See also *automated backup and restore*.

localization (L10N)
The process of adapting software for a specific culture or language by adding locale-specific components and translating text.

N

NDMP Backup-to-Disk (NDMP2D)
Software-based virtual tape library configured on a Data Mover that emulates a real TLU, and allows NDMP backup-to-disk functionality. This feature is sometimes referred to as a virtual tape library unit (VTLU). However, it does not provide true VTLU functionality as found in EMC Disk Library and similar storage products.

NDMP client
Application that controls the NDMP session. The NDMP client runs an NDMP-compliant backup application, such as EMC NetWorker®.

NDMP host
Host system (Data Mover) that executes the NDMP server application. Data is backed up from the NDMP host to either a local tape drive or to a backup device on a remote NDMP host.

NDMP server
NDMP process that runs on an NDMP host, which is a Data Mover in VNX for file environment.

NDMP Volume Backup (NVB)
EMC-specific type of NDMP backup mechanism that backs up data blocks at a volume level, rather than at a file level. NVB reads a set of disk data blocks in an efficient manner compared to the method used for traditional, file-based backups. NVB works only with EMC-qualified vendor backup software. NVB is commonly known as Volume Based Backup (VBB).

Network Data Management Protocol (NDMP)
Open standard network protocol designed for enterprise-wide backup and recovery of heterogeneous network-attached storage.

P

Portable Archive Interchange (PAX)
VNX for file archive protocol that works with standard UNIX tape formats and provides file-level backup and recovery operations.

R

remote backup
Remotely initiated procedure that backs up data to a TLU which can be attached to a target server different from the source Data Mover. The configuration determines whether data traverses the network.

See also *automated backup and restore* and *three-way backup*.

T

tape library unit (TLU)
Physical device that contains and manages sets of tape drives and cartridges.

tar
Backup format in PAX that traverses a file tree in depth-first order.

three-way backup
NDMP backup that directs data from one Data Mover through a network connection to a remote Data Mover to its attached tape library backup device.

See also *heterogeneous three-way backup* and *remote backup*.

U

Unicode
Family of universal character encoding standards used for representation of text for computer processing.

UTF-8 (8-bit Unicode Transformation Format)
Multibyte encoding form that uses an algorithmic mapping scheme to convert every Unicode value to a unique 1- to 4-byte sequence with no embedded null characters.

See *Unicode* or *UCS Transformation Format-8*.

Index

A

attributes
    for Windows NT 58

B

backup 24, 27, 70, 72, 114, 117, 124, 131, 135
    levels with VNX 24
    software
        CA BrightStor ARCserve Backup 114
        CommVault Galaxy 117
        EMC NetWorker 124
        Fujitsu-Siemens NetWorker 131
        IBM Tivoli Storage Manager 135
    statistics
        output explanation 72
        viewing 70
BakBone NetVault
    automatic device configuration 114
bufsz parameter
    configuring 46, 47

C

cabling the TLU 35
checkpoint
    integrating into backups 28
commands
    server_devconfig 55
    server_mount 74
    server_param 46, 47, 49, 66, 67
    server_pax 70
    server_user 57
CommVault Galaxy 117, 118, 120
    adding a client 118
    testing the TLU 120
configurations for NDMP
configurations for NDMP *(continued)*
    dual-TLU 22
    single TLU 22
configuring
    CA BrightStor ARCserve Backup 114
    CommVault Galaxy 117
    Data Movers 46, 47
    Fujitsu-Siemens NetWorker 131
    IBM Tivoli Storage Manager 135
    NS Series 47
    SyncSort Backup Express 141
    tape buffer size 46
    Tivoli Storage Manager 135
configuring Data Movers 47
connecting
    TLU robot 35
creating
    NDMP user 56

D

Data Movers
    configuring 45, 46, 47
    username and password creation 56
DDS
    overview 26
    restriction 12
device names
    sample 56
direct access restore
    DAR 26
directory filtering 32
dynamic drive sharing
    overview 26

dynamic drive sharing *(continued)*
restriction 12

E

EMC Data Manager
configuration 123
local backups 121
overview 120
EMC E-Lab Navigator 102
EMC Networker
restrictions 124
EMC NetWorker 46, 125, 129
bufsz parameter 46
jbconfig script 125
tape buffer size 129
environment variables
setting 58
with FileMover 74, 75
error messages 104
exclusion
file and directory 32

F

file filtering 32
file-level restore 29
FileMover
stub file 27
with NDMP 27
Fujitsu-Siemens NetWorker 131
full destructive restore 29

H

HP OpenView Storage Data Protector
configuring 132, 133
creating a media pool 133

I

IBM Tivoli Storage Manager 135
international character sets
implementation 36
support overview 36
vendor support for 36
with NDMP 37

M

maxProtocolVersion parameter 49
MD5 password encryption 57
media interface adapter 13
messages, error 104
multithreaded backup design 42

N

NASA thread 42
NASS thread 42
NASW thread 42
NDMP
and file filtering 32
and integrated checkpoints using VNX SnapSure 28
and international character sets 37
and vbb 29
and VNX NDMP Volume Backup 29
backups to disk
concepts 27
concepts 19
configurations 19
environment variables 58
user account 56
version
overriding 49
Network Data Management Protocol 19
NIS
user account 57
NS series configuration 47

O

overriding
NDMP version 49

P

parameters
bufsz 46
configuring 46, 47
maxProtocolVersion 49
password
for Data Mover access 56
Performance
TLU configuration 35

R

rawfs volume ID 81
related information 16
requirements
software and hardware 12
restrictions
configuration 12
general 12

restrictions *(continued)*
NDMP 13
TLU 13
robot
connection with Data Mover 35

S

SCSI TLU cabling 35
server_mount command 74
server_param command 47, 49, 66, 67
server_pararm command 46
server_pax command 70
setting
:environment variables on VNX 58
username and password 56
software
SyncSort Backup Express 141
statistics
backup performance 70
stripe size
optimal 12
SyncSort Backup Express 141
Configur the Data Mover NDMP 141
system requirements 12

T

tape buffer size
configuring 46
three-way backups
heterogeneous 23
overview 23
with EMC Data Manager 121
Tivoli Storage Manager 135
TLU
cabling 35
configurations
Fibre Channel 35
single TLU 22
three-way backup 23
robot connectivity 35
SCSI port
TLU *(continued)*
SCSI port *(continued)*
sample 56
virtual 27
to disk
with NDMP 27

U

Unicode
with NDMP 36
username
for NDMP 56

V

vbb 29
VDM 74
version
overriding NDMP 49
viewing backup
statistics 70
Data Mover
with NDMP 74
tape library unit
concepts 27
VNX
backup levels 24
Data Mover configuration 45, 46
environment variables 58
parameter configuration 45, 46
VNX SnapSure checkpoints
integrating 28
volume ID
rawfs 81
volume-based backup 29
VTLU
concepts 27

W

Windows NT attributes
backing up 58