

Data ONTAP® 7.1

Data Protection

Tape Backup and Recovery Guide

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number 210-01080_B0
December 2005

Copyright and trademark information

Copyright information

Copyright © 1994–2005 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Portions of this product are derived from the Berkeley Net2 release and the 4.4-Lite-2 release, which are copyrighted and publicly distributed by The Regents of the University of California.

Copyright © 1980–1995 The Regents of the University of California. All rights reserved.

Portions of this product are derived from NetBSD, which is copyrighted by Carnegie Mellon University.

Copyright © 1994, 1995 Carnegie Mellon University. All rights reserved. Author Chris G. Demetriou.

Permission to use, copy, modify, and distribute this software and its documentation is hereby granted, provided that both the copyright notice and its permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation.

CARNEGIE MELLON ALLOWS FREE USE OF THIS SOFTWARE IN ITS “AS IS” CONDITION. CARNEGIE MELLON DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING FROM THE USE OF THIS SOFTWARE.

Software derived from copyrighted material of The Regents of the University of California and Carnegie Mellon University is subject to the following license and disclaimer:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notices, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER

IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software contains materials from third parties licensed to Network Appliance Inc. which is sublicensed, and not sold, and title to such material is not passed to the end user. All rights reserved by the licensors. You shall not sublicense or permit timesharing, rental, facility management or service bureau usage of the Software.

Portions developed by the Apache Software Foundation (<http://www.apache.org/>). Copyright © 1999 The Apache Software Foundation.

Portions Copyright © 1995–1998, Jean-loup Gailly and Mark Adler

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

Portions Copyright © 1995 University of Southern California. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of Southern California, Information Sciences Institute. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

Portions of this product are derived from version 2.4.11 of the libxml2 library, which is copyrighted by the World Wide Web Consortium.

Network Appliance modified the libxml2 software on December 6, 2001, to enable it to compile cleanly on Windows, Solaris, and Linux. The changes have been sent to the maintainers of libxml2. The unmodified libxml2 software can be downloaded from <http://www.xmlsoft.org/>.

Copyright © 1994–2002 World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>

Software derived from copyrighted material of the World Wide Web Consortium is subject to the following license and disclaimer:

Permission to use, copy, modify, and distribute this software and its documentation, with or without modification, for any purpose and without fee or royalty is hereby granted, provided that you include the following on ALL copies of the software and documentation or portions thereof, including modifications, that you make:

The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.

Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, a short notice of the following form (hypertext is preferred, text is permitted) should be used within the body of any redistributed or derivative code: "Copyright © [\$date-of-software] World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>.

Notice of any changes or modifications to the W3C files, including the date changes were made.

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS

FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders.

Software derived from copyrighted material of Network Appliance, Inc. is subject to the following license and disclaimer:

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, Data ONTAP, FAServer, FilerView, MultiStore, NearStore, NetCache, SecureShare, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapVault, SyncMirror, and WAFL are registered trademarks of Network Appliance, Inc. in the United States, and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The Evolution of Storage are trademarks of Network Appliance, Inc. in the United States and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexClone, FlexVol, FPolicy, HyperSAN, InfoFabric, LockVault, Manage ONTAP, NOW, NOW NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapMigrator, SnapSuite, SnapValidator, SohoFiler, vFiler, VFM, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the United States and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the United States. Spinnaker Networks, the Spinnaker Networks logo, SpinAccess, SpinCluster, SpinFS, SpinHA, SpinMove, and SpinServer are registered trademarks of Spinnaker Networks, LLC in the United States and/or other countries. SpinAV, SpinManager, SpinMirror, SpinRestore, SpinShot, and SpinStor are trademarks of Spinnaker Networks, LLC in the United States and/or other countries.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	.xi
Chapter 1	Introduction to Data Protection	1
	Data protection options	2
	Tape backup and recovery	4
Chapter 2	Tape Drive Management	7
	Understanding tape devices.	8
	Displaying tape device statistics	12
	Displaying supported tape drives.	14
	Assigning tape aliases	16
	Displaying tape device and tape drive information.	24
	Understanding tape files	28
	Controlling tape drives	29
	Qualifying tape drives dynamically	34
	Using nonqualified tape drives	38
	Enabling and disabling persistent reservations	42
Chapter 3	Data Backup to Tape	45
	Introduction to data backup to tape.	46
	Types and methods of backup	47
	What the dump command backs up	49
	How the dump command works	50
	Preparing to use the dump command.	52
	Determining the amount of backup data	53
	Estimating the number of tapes for the backup	54
	Deciding how to enter the dump command	55
	Using the dump command	56
	Dump command syntax	57
	Recommendations for performing a backup	60
	Specifying the backup level	63

Understanding increment chains	64
Specifying the tape device	66
Specifying the dump path	68
Specifying a list of files	69
Excluding items from a backup	70
Specifying a name for a backup	72
Specifying no Windows NT ACLs.	73
Specifying the blocking factor	74
Specifying the tape file size	76
Specifying updates to the /etc/dumpdates file	77
Appending backups to a tape	79
Examples of the dump command.	80
Level-0 backups to a local tape drive	81
Incremental backups to a local tape drive	83
Backups to a remote tape drive.	84
Backups to multiple tapes	86
Backup of a directory from a Snapshot copy	88
Backup of a list of files or directories	89
Backup data without ACLs.	90
Backup specifying a blocking factor.	91
Backup excluding specified files.	92
Verifying the contents of a dump command backup	93
Checking the status of a dump command backup	94
Restarting a dump command backup.	96
Deleting restartable dump command backups	98

Chapter 4

Data Restoration from Tape	99
About data restoration	100
What the restore command recovers	101
How the restore command works.	103
Preparing to use the restore command	104
Getting required materials	105
Deciding where to enter the restore command.	106
Preparing the destination	107
Using the restore command.	108
Restore command principles	109
Specifying an individual file restore	113
Specifying a full restore	114
Specifying a table-of-contents restore	115

Specifying a resumed restore.	117
Specifying tape devices	119
Specifying a single tape file on a multiframe tape	121
Specifying the restore destination	122
Specifying the blocking factor	123
Specifying detailed status output.	124
Specifying ignoring inode limitations	125
Specifying automatic confirmations	127
Specifying no Windows NT ACLs.	128
Specifying no qtree information	129
Specifying a test restore	130
Restoring incremental backups	131
Restoring the entire storage system.	132
Restore examples	133
Restoring an entire backup.	134
Restoring individual files or directories	136
Restoring using a remote tape drive	137
Multiple tape restores	138
 Chapter 5	
Using NDMP Services	141
Learning about NDMP services on the storage system.	142
Using tape devices with NDMP	146
Preparing a storage system for basic NDMP backup application management	148
Enabling or disabling NDMP service.	150
Managing NDMP security features.	151
Specifying the NDMP version	157
Specifying a preferred network interface.	159
Enabling or disabling enhanced DAR functionality	161
Specifying environment variables for backup applications	162
Displaying NDMP session status.	171
Displaying NDMP debug messages	174
Terminating an NDMP session	177
Copying data with ndmpcopy	178

Chapter 6

Event Logging for Backup and Restore	183
About event logging for backup and restore	184
Event log messages	185
Enabling or disabling event logging	190
 Glossary	 191
 Index	 197

Preface

About this guide	This guide describes how to protect, back up, restore, and copy data between Network Appliance® filers that run Data ONTAP® software and that are Serving Data by Design™. This guide covers all filer models.
Audience	This guide is for system administrators who are familiar with operating systems that run on the filer's clients, such as UNIX®, Windows 95™, Windows NT®, and Windows® 2000. It also assumes that you are familiar with how to configure the filer and how the NFS, CIFS, and HTTP protocols are used for file sharing or transfers. This guide does not cover basic system or network administration topics, such as IP addressing, routing, and network topology; it emphasizes the characteristics of the NetApp® filer.
Terminology	Network Appliance's storage products (filer, FAS appliances, and NearStore® systems) are all <i>storage systems</i>. All these products are also sometimes called <i>filers</i> or <i>storage appliances</i>. This guide uses the term “type” to mean pressing one or more keys on the keyboard. It uses the term “enter” to mean pressing one or more keys and then pressing the Enter key.
Command conventions	You can enter filer commands on the system console or from any client that can obtain access to the filer using a Telnet session. In examples that illustrate commands executed on a UNIX workstation, the command syntax and output might differ, depending on your version of UNIX.
Formatting conventions	The following table lists different character formats used in this guide to set off special information.

Formatting convention	Type of information
<i>Italic type</i>	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fcstest adaptername</code> command, you enter the characters “fcstest” followed by the actual name of the adapter. ◆ Book titles in cross-references.
Monospaced font	◆ Command and daemon names. ◆ Information displayed on the system console or other computer monitors. ◆ The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

Keyboard conventions

This guide uses capitalization and some abbreviations to refer to the keys on the keyboard. The keys on your keyboard might not be labeled exactly as they are in this guide.

What is in this guide...	What it means...
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
<i>Enter</i>	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards.
<i>type</i>	Used to mean pressing one or more keys on the keyboard.
<i>enter</i>	Used to mean pressing one or more keys and then pressing the Enter key.

Special messages

This guide contains special messages that are described as follows:

Note

A note contains important information that helps you install or operate the system efficiently.

Caution

A caution contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

WARNING

A warning contains instructions that you must follow to avoid personal injury.

About this chapter

This chapter introduces the data backup, restore, and protection capabilities that are available in this release of Data ONTAP.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Data protection options”](#) on page 2
- ◆ [“Tape backup and recovery”](#) on page 4

Data protection options

About data protection

Data protection means backing up data and being able to recover it. You protect the data by making copies of it so that it is available for restoration even if the original is no longer available.

Why data needs protection

- Reasons that businesses need data backup and protection systems include the following:
- ◆ To protect data from accidentally deleted files, application crashes, data corruption, and viruses
 - ◆ To archive data for future use
 - ◆ To recover from a disaster

Methods of protecting data

Depending on your data protection and backup needs, Data ONTAP offers a variety of features and methods to ensure against accidental, malicious, or disaster-induced loss of storage system data. This guide describes how to use Data ONTAP to back up and recover data from tape. See the *Online Backup and Recovery Guide* for information about how to use Data ONTAP features like SnapMirror® software to protect data.

Data Protection Feature	Description
Tape backup dump and restore commands	Tape backup and restore. The <code>dump</code> and <code>restore</code> commands allow you to back up storage system Snapshot copies to tape. The <code>dump</code> command takes a Snapshot copy of the volume and then copies that data to tape. Because the Snapshot copy, not the active file system, is backed up to tape, the storage system can continue its normal functions while the tape backup takes place. See Chapter 2, “Tape Drive Management” on page 7, Chapter 3, “Data Backup to Tape” on page 45, and Chapter 4, “Data Restoration from Tape” on page 99.

Data Protection Feature	Description
NDMP support	Support for tape backup management of the storage system by third-party applications that use the Network Data Management Protocol (NDMP). Data ONTAP also provides an NDMP copy command that you can use to carry out NDMP-compliant backups and restores. Security login restricts access to NDMP operations. See Chapter 5, “Using NDMP Services” on page 141.

Tape backup and recovery

When to use tape backup and recovery

You use tape backup and recovery to create tape archives or retrieve data from tape archives. Tape archives complement online backup methods, and you should plan to do tape backups regardless of the online backup and recovery methods you use.

Advantages of tape archives

Making tape archives gives you the following advantages over relying solely on online storage:

- ◆ Once made, they require few resources to maintain.
- ◆ You can put the archives in a more secure place than you can put a storage system.
- ◆ You can recover data on Sun Solaris systems with the `ufsrestore` command.
- ◆ Data archived on tape is on a physically stable medium and is stored in a standard format. You can recover data from any release of Data ONTAP.

Disadvantages of tape archives

Tape archives have the following disadvantages compared to online storage:

- ◆ It takes a long time to restore data from tape.
- ◆ It is time-consuming to find a particular file or directory on tape.
- ◆ Some systems are too large to be backed up to tape efficiently.
Large file systems might need to be broken into multiple volumes and backed up on different nights. Volumes over 500 GB or total data set sizes greater than 1.4 TB can take a long time to restore. These volumes are good candidates for backup using the SnapMirror® or SnapVault® features.

Where to find backup and restore information

If you want to perform backups to tape and restores from tape, you should be familiar with the topics listed in the following table.

Topic	Chapter
Understanding tapes and controlling tape drives	Chapter 2, “ Tape Drive Management ” on page 7
Backing up data to tape	Chapter 3, “ Data Backup to Tape ” on page 45
Restoring data from tape	Chapter 4, “ Data Restoration from Tape ” on page 99
Using the NDMP protocol, which third-party backup software vendors use	Chapter 5, “ Using NDMP Services ” on page 141

About this chapter

This chapter describes methods for using tape *devices* (software functionality designations), tape *drives* (the physical hardware modules that house and control the tapes), and tape aliases (methods of addressing tape drives). You need to manage tape drives when you back up data from the storage system to tape or when you recover data from tape to the storage system. In particular, you need to understand the relationship between tape devices and tape drives so you can specify devices in performing backups and restores. Tape aliases are names you assign to tape drives to make it easier to keep track of tape drives used for multiple storage systems.

In this and subsequent chapters, the term *tape device* means a specific functionality of a physical tape drive that you create by specifying information in the tape device name when you install a tape drive, which handles a single tape, or a tape stacker or tape library (medium changer), which handles multiple tapes.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Understanding tape devices](#)” on page 8
- ◆ “[Displaying tape device statistics](#)” on page 12
- ◆ “[Displaying supported tape drives](#)” on page 14
- ◆ “[Assigning tape aliases](#)” on page 16
- ◆ “[Displaying tape device and tape drive information](#)” on page 24
- ◆ “[Understanding tape files](#)” on page 28
- ◆ “[Controlling tape drives](#)” on page 29
- ◆ “[Qualifying tape drives dynamically](#)” on page 34
- ◆ “[Using nonqualified tape drives](#)” on page 38
- ◆ “[Enabling and disabling persistent reservations](#)” on page 42

Understanding tape devices

About tape devices

When you install a tape drive or tape stacker on a storage system according to the tape drive manufacturer's instructions, Data ONTAP creates tape devices associated with the tape drive or stacker. A device is created for each combination of rewind and compression capability. Therefore, a tape drive or tape stacker can have several tape devices associated with it. You must specify a tape device to move, write, or read tapes.

Data ONTAP detects tape drives and tape stackers attached before or after bootup and assigns logical numbers and tape devices to them.

Kinds of tape devices

The following kinds of tape devices are available:

- ◆ A local tape device on the storage system that is performing a tape operation
- ◆ A remote tape device on a NetApp or Solaris machine that
 - ❖ Is not the machine that is performing a tape operation, but is connected through the network to a host that is performing the tape operation
 - ❖ Is running the `rmt` protocol
 - ❖ Has a trust relationship with the storage system that is performing the tape operation

Note

You cannot use tape devices associated with tape stackers or tape libraries (medium changers) on a remote Solaris system.

Tape device name format

The format of a tape device name is as follows:

```
[remote_host:]rewind_type st drive_number compression_type
```

remote_host is optional. You specify a remote host storage system if you want to use a tape drive that is attached to that host. You must follow the remote host name with a colon (:).

rewind_type is the rewind type, as shown in the following table.

Rewind type	Description
r	Rewind. Data ONTAP rewinds the tape after it finishes writing the tape file.
nr	No rewind. Data ONTAP does not rewind the tape after it finishes writing the tape file. If you want to write multiple tape files on the same tape, specify no-rewind tape device names.
ur	Unload/reload. You can use this only if you are performing operations on tapes in tape stackers. When you use these devices, the stacker unloads the tape when it reaches the end of a tape file, and then loads the next tape if there is one. Use this rewind type only under the following circumstances: ◆ The tape drive associated with this device is housed in a tape stacker or a medium changer that is set to stacker mode. ◆ The tape drive associated with this device is attached to a storage system. ◆ Sufficient tapes for the operation that you are performing are available in the stacker tape sequence defined for this tape drive.

Note

If you record a tape using a no-rewind device, you must rewind the tape before you read it.

st is the standard designation for a tape drive.

drive_number is the drive number. It is incremented for each tape drive; if there is only one tape drive attached to the storage system or remote host, the number is 0. The assigned drive number persists through reboots.

compression_type is a drive-specific code for the density of data on the tape and the type of compression. The values for *compression_type* are shown in the following table.

Compression type	Description
a	Highest compression
h	High compression
l	Low compression
m	Medium compression

Examples: `nrst0a` specifies a no-rewind device on tape drive 0 using the highest compression.

`remfiler:nrst0a` specifies a no-rewind device on tape drive 0 on the remote host remfiler that uses the highest compression.

Caution

When using the `urst` device with the `dump` or `restore` command, ensure that the tape drives it is used with are stackers and that there are sufficient tapes in the stacker sequence. Otherwise, the tape drives involved will terminate the command sequence or overwrite the same tape multiple times.

Sample listing of tape devices

The following example shows the tape devices associated with a Digital™ DLT7000:

```
rst0l - rewind device, format is: 81633 bpi 40 GB (w/comp)
nrst0l - no rewind device, format is: 81633 bpi 40 GB (w/comp)
urst0l - unload/reload device, format is: 81633 bpi 40 GB (w/comp)
rst0m - rewind device, format is: 85937 bpi 35 GB
nrst0m - no rewind device, format is: 85937 bpi 35 GB
urst0m - unload/reload device, format is: 85937 bpi 35 GB
rst0h - rewind device, format is: 85937 bpi 50 GB (w/comp)
nrst0h - no rewind device, format is: 85937 bpi 50 GB (w/comp)
urst0h - unload/reload device, format is: 85937 bpi 50 GB (w/comp)
rst0a - rewind device, format is: 85937 bpi 70 GB (w/comp)
nrst0a - no rewind device, format is: 85937 bpi 70 GB (w/comp)
urst0a - unload/reload device, format is: 85937 bpi 70 GB (w/comp)
```

The abbreviations in the example are defined as follows:

- ◆ bpi—Bits per inch; this is the density of the tape.
- ◆ GB—Gigabytes; this is the capacity of the tape.
- ◆ w/comp—With compression; this shows the tape capacity with compression.

Displaying tape device statistics

Why to display tape drive statistics

It can be helpful to display tape device statistics to understand tape performance and check usage patterns. You can zero out the statistics reading and restart the process of displaying the statistics whenever you want.

How to display tape drive statistics

To display the statistics for a specified tape device, complete the following step.

Step	Action
1	Enter the following command: storage stats tape <i>tape</i> <i>tape</i> is the name of a tape device.

Example: storage stats tape nrst01

Example of the output of the storage stats tape command

The following example shows the statistics for the tape device nrst01.

```
filerA> storage stats tape nrst01

Bytes Read: 71471104
Bytes Written: 382147584
Command      Num issued  Max (ms)  Min (ms)  Avg (ms)
-----
WRITE - Total      2518      1927        2        24  6269 KB/s
  44-48KB           897       372         2         6  6531 KB/s
  60-64KB           421      1927         3        13  4796 KB/s
 128-132KB          800       131         8        19  6761 KB/s
 508KB+             400       481        32        83  6242 KB/s
READ - Total      1092      1570         5        14  4582 KB/s
  60-64KB           92      1390         5        25  2493 KB/s
  64-68KB          1000     1570         5        13  4958 KB/s
WEOF              5       2827     2787     2810
FSF                1    13055    13055    13055
BS                 0         0         0         0
FSR                2     1390         5        697
BSR                1        23        23        23
REWIND             9    67606        94    22260
```

ERASE	0	0	0	0
EOM	1	63223	63223	63223

How to zero out tape drive statistics

To zero out the statistics so that only new tape device statistics beginning from the time this command is entered will be displayed the next time you enter the command, complete the following step.

Step	Action
1	Enter the following command: storage stats tape zero <i>tape</i> <i>tape</i> is the name of a tape drive.

Example: `storage stats tape zero nrst01`

Displaying supported tape drives

Why to display supported tape drives

Before connecting a tape drive to a storage system, you should make sure the tape drive is supported by the storage system.

How to display supported tape drives

To display a list of the tape drives supported by the storage system, complete the following step.

Step	Action
1	Enter the following command: storage show tape supported [-v] -v is an option that displays more complete information about each tape drive.

Sample output of the storage show tape supported command

The following example shows the output of the storage show tape supported command.

```
storage show tape supported
Supported Tapes
-----
Exabyte 8500C 8mm
Exabyte 8505 8mm
Digital DLT2000
Quantum DLT2000
```

Sample output of the storage show tape supported -v command

The following example shows the output of the storage show tape supported command with the -v option.

```
storage show tape supported -v
Supported Tapes
-----
Exabyte 8500C 8mm
Density Compression
```

Setting	Setting		
-----	-----		
0x14	0x00	EXB-8200	2.5GB
0x90	0x10	EXB-8200C	(w/compression)
0x15	0x00	EXB-8500	5.0GB
0x8c	0x10	EXB-8500C	(w/compression)

Exabyte 8505 8mm

Density Compression

Setting	Setting		
-----	-----		
0x14	0x00	EXB-8200	2.5GB
0x90	0x10	EXB-8200C	(w/compression)
0x15	0x00	EXB-8500	5.0GB
0x8c	0x10	EXB-8500C	(w/compression)

Digital DLT2000

Density Compression

Setting	Setting		
-----	-----		
0x00	0x00	default density	
0x17	0x00	42500 bpi	2.6GB
0x18	0x00	42500 bpi	6.0GB
0x19	0x10	62500 bpi	20 GB (w/comp)

Quantum DLT2000

Density Compression

Setting	Setting		
-----	-----		
0x00	0x00	default density	
0x17	0x00	42500 bpi	2.6GB
0x18	0x00	42500 bpi	6.0GB
0x19	0x10	62500 bpi	20 GB (w/comp)

Assigning tape aliases

Why to assign tape aliases

When you use more than one tape drive (or tape stacker or medium changer) on a storage system, Data ONTAP assigns the addresses on a first-come, first-served basis when the software is installed. However, if a storage system shares tape drives, or if tape drives are moved around, the software-assigned addresses do not move with the physical devices. After reboot, these addresses are not necessarily reassigned in the same order.

If you use large or multi-storage system configurations, or if your storage systems share backup tape devices, the storage systems may assign different names to the same tape drives. Tape aliasing enables you to make sure that the same name is always associated with the same tape drive.

What tape aliasing does

Tape aliasing provides a correspondence between the logical names of backup devices (for example, `st0` or `mc1`) and a name permanently assigned to a port or to a tape drive or medium changer.

Note

`st0` and `st00` are considered different logical names.

You can use tape aliases as arguments to the `dump`, `restore`, and `snapmirror store` commands.

Note

Logical and worldwide names are used only to access a device. Once accessed, the device returns all error messages using the physical path name.

Names available for aliasing

The two kinds of names available for aliasing are physical path names (PPNs) and worldwide names (WWNs).

What physical path names are

Physical path names (PPNs) are the numerical address sequences that Data ONTAP assigns to tape drives and tape libraries based on which SCSI-2/3 adapter or switch (specific location) they are connected to on the storage system. PPNs are also known as electrical names.

Format of the PPN

Direct-attached devices: PPNs of direct-attached devices use the following format:

host_adapter.device_id_lun

For example, the PPN 8.6 indicates that the host adapter number is 8, the device ID is 6, and the logical unit number (LUN) is 0.

Note

On disk drives, the LUN is always 0 and is therefore never shown. On tape drives and medium changers, the LUN is also not shown if it is 0.

Switch-attached devices: PPNs of switch-attached devices use the following format:

switch:port_id.device_id_lun

For example, the PPN MY_SWITCH:5.3L2 indicates that the tape drive connected to port 5 of a switch called MY_SWITCH is set with SCSI ID 3 and has the LUN 2.

The LUN is determined by the drive itself. Both Fibre Channel and SCSI tape drives and libraries, as well as disks, have PPNs.

Examples using PPNs

In the following example, the `dump` command is using the tape device name of a tape drive:

```
dump 0f /dev/nrst0a /vol/vol0
```

In the following example, the same `dump` command is using the PPN of the tape drive:

```
dump 0f /dev/nr.MY_SWITCH:5.6.a /vol/vol0
```

PPNs of tape drives and libraries do not change unless the name of the switch changes, the tape drive or library moves, or the tape drive or library is reconfigured. PPNs remain unchanged after reboot.

For example, if a tape drive named MY_SWITCH:5.3L2 is removed and a new tape drive with the same SCSI ID and LUN is connected to port 5 of switch MY_SWITCH, the new tape drive would be accessible via the name MY_SWITCH:5.3L2.

What worldwide names are

Tape drives and libraries are assigned worldwide names (WWNs) at the time of manufacture. WWNs are similar to the media access control (MAC) addresses on Ethernet cards. All Fibre Channel devices have WWNs, but SCSI-attached devices do not have WWNs and cannot be accessed with a WWN.

Accessing a tape drive or library via the WWN allows multiple storage systems to track the same device. Depending on whether a tape drive is connected to a Fibre Channel switch or hub or is directly attached to a Fibre Channel adapter, different storage systems may have different PPNs for the same device. Using the WWN in these cases eliminates any confusion.

Also, if a switch is renamed or a tape drive is moved on the storage system, the WWN of the tape drive or library does not change, and scripts or backup programs do not need to change the name of the tape drive or library to which they are backing up.

Format of the WWN

The WWN of a tape device uses the following format:

WWN[#:###:#####:#####]L##

is a hexadecimal character and L is the LUN of the device. If the LUN is 0, it is not shown.

Each WWN consists of eight bytes, and the format for the WWN is not case-sensitive.

Examples using WWNs

In the following example, the `dump` command uses the logical name of a tape drive:

```
dump 0f /dev/nrst0a /vol/vol0
```

In the following example, the same `dump` command uses the worldwide name of the tape drive:

```
dump 0f /dev/nr.WWN[2:000:00e08b:01523e].a /vol/vol0
```

Displaying existing aliases of tape drives

To display existing aliases of tape drives, complete the following step:

Step	Action												
1	Enter the following command: storage alias Result: A display like the following appears: <table><tr><td>Alias</td><td>Mapping</td></tr><tr><td colspan="2">-----</td></tr><tr><td>st0</td><td>MY_SWITCH:5.3L3</td></tr><tr><td>st2</td><td>MY_SWITCH:5.4L6</td></tr><tr><td>mc1</td><td>2:4e3:38fe3f:758eab</td></tr><tr><td>mc348</td><td>MY_SWITCH:5.3L0</td></tr></table> In this example, the display shows that there are two tape drives and two medium changers attached to the storage system. Tape drives st0 and st2 and medium changer mc348 are attached to port 5 of the switch MY_SWITCH. Medium changer mc1 has the worldwide name 2:4e3:38fe3f:758eab.	Alias	Mapping	-----		st0	MY_SWITCH:5.3L3	st2	MY_SWITCH:5.4L6	mc1	2:4e3:38fe3f:758eab	mc348	MY_SWITCH:5.3L0
Alias	Mapping												

st0	MY_SWITCH:5.3L3												
st2	MY_SWITCH:5.4L6												
mc1	2:4e3:38fe3f:758eab												
mc348	MY_SWITCH:5.3L0												

Displaying information about tape drives or libraries

To display information about tape drives and tape libraries (medium changers) to help you assign tape aliases, complete the following step.

Step	Action
1	Enter the following command: storage show {tape me} [{alias PPN WWN}] <i>alias</i> is the logical name of the tape drive or medium changer. <i>PPN</i> is the physical path name. <i>WWN</i> is the worldwide name.

Examples:

```
storage show tape

Tape Drive:      MY_SWITCH:5.3L4
Description:     Quantum DLT7000
Serial Number:   12345679
World Wide Name: WWN[2:333:444444:555555]L4
```

```

Alias Name(s):      st0  st1  st2  st3

Tape Drive:         MY_SWITCH:5.3L5
Description:        Quantum DLT7000
Serial Number:      12345678
World Wide Name:    WWN[2:777:888888:999999]L5
Alias Name(s):      st10 st11 st12 st13

```

```

storage show tape st0
Tape Drive:         MY_SWITCH:5.3L4
Description:        Quantum DLT7000
Serial Number:      12345679
World Wide Name:    WWN[2:333:444444:555555]L4
Alias Name(s):      st0  st1  st2  st3

```

```

storage show tape MY_SWITCH:5.3L4
Tape Drive:         MY_SWITCH:5.3L4
Description:        Quantum DLT7000
Serial Number:      12345679
World Wide Name:    WWN[2:333:444444:555555]L4
Alias Name(s):      st0  st1  st2  st3

```

Assigning tape aliases

To assign an alias to a tape drive or medium changer, complete the following step.

Step	Action
1	Enter the following command: <pre>storage alias [alias {PPN WWN}]</pre> <i>alias</i> is the logical name of the tape drive or medium changer to which you want to add the alias. <i>PPN</i> is the physical path name to which you want to assign the tape drive or medium changer. <i>WWN</i> is the worldwide name to which you want to assign the tape drive or medium changer.

Example: `storage alias st0 MY_SWITCH:5.3L3`
`storage alias mc80 WWN[2:4e3:38fe3f:758eab]`

Result: The tape device st0 is assigned the physical path name MY_SWITCH:5.3L3. The medium changer mc80 is assigned the worldwide name WWN[2:4e3:38fe3f:758eab].

Removing tape aliases

To remove an alias from a tape drive or medium changer, complete the following step.

Step	Action
1	Enter the following command: <pre>storage unalias {alias -a -m -t}</pre> <i>alias</i> is the logical name of the tape drive or medium changer from which you want to remove the alias. -a removes all aliases. -m removes the aliases from all medium changers. -t removes the aliases from all tape drives.

Example:

```
storage unalias st0
storage unalias mc80
```

Result: The aliases are removed from the tape device st0 and from the medium changer mc80.

Propagating tape aliases to multiple storage systems

If you need to use the same set of tape drives to back up more than one storage system, you can save the tape alias information in a file.

To propagate tape aliases to multiple storage systems, complete the following steps.

Step	Action
1	Create a file named <i>tape_alias</i> containing the tape alias information. Example: <pre>storage unalias -a storage alias st0 8.6 storage alias st1 8.7 storage alias mc0 8.1</pre>

Step	Action
2	Copy the file to the root volume of each storage system.
3	Execute the following command on each storage system: <pre>source /vol/root_volume_name/tape_alias</pre> <i>root_volume_name</i> specifies the root volume. Result: The storage systems all contain the same configuration information.

To ensure that multiple storage systems assign the same alias to a tape drive or medium changer, you can type the same set of storage alias commands on each storage system.

Alternatively, UNIX users can use a shell script like the following to propagate the source script information to the storage systems.

```
#!/bin/sh
# Check for the source file
if [ "$#" != "1" ]
then
    echo "Usage: $0 <source_file>"
    exit 1
fi

if [ ! -r $1 ]
then
    echo "Cannot open source file \"$1\""
    exit 1
fi

while [ 1 ]
do

    echo Hit ctrl-c to terminate program when all filers have been
    entered.

# Gather up filer and passwd from user
    printf "File Server: "
    read FILER
    printf "Password: "
    read PASSWD
    stty -echo
    stty echo
    printf "\n"
```

```

# Now issue the commands in the source script to
# the remote filer.

while read cmd
do
    echo Issuing command \"$cmd\" to filer $FILER
    rsh $FILER -l root:$PASSWD \"$cmd\" < /dev/null
done < $1
printf "\n"
done

```

Adding Fibre-Channel-attached tape drives and libraries

You can add Fibre-Channel-attached tape drives and libraries to storage systems dynamically (without taking the storage systems offline). When the new medium changer is added, the storage system detects its presence and adds it to the configuration. If the medium changer is already referenced in the alias information, no new logical names are created. If the library is not referenced, the storage system creates a new alias for the medium changer.

Displaying tape device and tape drive information

About displaying tape device information

Before you back up to a tape device attached to the storage system, you might want to verify that the storage system detects the tape drive associated with the device and verify the available tape device names associated with the tape drive.

Qualified versus nonqualified tape drives

Qualified tape drives are tape drives that Network Appliance has tested and found to work properly with storage systems. A list of qualified tape drives is available on the NOW™ (NetApp on the Web) site at

http://now.netapp.com/NOW/knowledge/docs/olio/cert_matrices/

Nonqualified tape drives are not supported by Network Appliance. You can use a nonqualified tape drive by qualifying the tape drive dynamically, as described in “[Qualifying tape drives dynamically](#)” on page 34, or by making it emulate a qualified tape drive, as described in “[Using nonqualified tape drives](#)” on page 38.

Information about qualified tape drives in tape libraries

Qualified tape drives that are in tape stackers or libraries are listed by drive. For example, a Breece-Hill library is not listed in all displays, whereas the DLT7000 tape drives that it contains are listed.

Types of information available

You can display information about

- ◆ Qualified tape drives
- ◆ Nonqualified tape drives
- ◆ Tape stackers
- ◆ Tape drive connections to the storage system

Displaying information about tape drives

To display information about all tape drives on a storage system, complete the following step.

Step	Action
1	Enter the following command: <pre>sysconfig -t</pre> Result: A display like the following appears. The numbers following “Tape drive” show the slot on the storage system that the drive is attached to, followed by the drive’s SCSI ID. In the following example, the Exabyte 8900 has SCSI ID 1 and is attached to a controller in slot 0b. Information about the Digital DLT7000 includes bpi specifications provided by the manufacturer. Note _____ Compression capacity in the display is estimated; actual capacity depends on how much the data being written to the tape can be compressed. <pre>Tape drive (0b.1) Exabyte 8900 8mm rst0l - rewind device, format is: EXB-8500 5.0GB (readonly) nrst0l - no rewind device, format is: EXB-8500 5.0GB (readonly) urst0l - unload/reload device, format is: EXB-8500 5.0GB (readonly) rst0m - rewind device, format is: EXB-8500C (w/compression) nrst0m - no rewind device, format is: EXB-8500C (w/compression) urst0m - unload/reload device, format is: EXB-8500C (w/compression) rst0h - rewind device, format is: EXB-8900 10.0GB nrst0h - no rewind device, format is: EXB-8900 10.0GB urst0h - unload/reload device, format is: EXB-8900 10.0GB rst0a - rewind device, format is: EXB-8900C (w/compression) nrst0a - no rewind device, format is: EXB-8900C (w/compression) urst0a - unload/reload device, format is: EXB-8900C (w/compression) Tape drive (5.3) Digital DLT7000 rst1l - rewind device, format is: 81633 bpi 40 GB (w/comp) nrst1l - no rewind device, format is: 81633 bpi 40 GB (w/comp) urst1l - unload/reload device, format is: 81633 bpi 40 GB (w/comp) rst1m - rewind device, format is: 85937 bpi 35 GB nrst1m - no rewind device, format is: 85937 bpi 35 GB urst1m - unload/reload device, format is: 85937 bpi 35 GB rst1h - rewind device, format is: 85937 bpi 50 GB (w/comp) nrst1h - no rewind device, format is: 85937 bpi 50 GB (w/comp) urst1h - unload/reload device, format is: 85937 bpi 50 GB (w/comp) rst1a - rewind device, format is: 85937 bpi 70 GB (w/comp) nrst1a - no rewind device, format is: 85937 bpi 70 GB (w/comp) urst1a - unload/reload device, format is: 85937 bpi 70 GB (w/comp)</pre>

Displaying information about tape medium changers

To display information about tape medium changers, complete the following step.

Step	Action
1	Enter the following command: <pre>sysconfig -m</pre> Result: A display like the following appears. In this example, a tape stacker with SCSI ID 4 is attached to slot 6 of the storage system. <pre>Medium changer (6.4) BHTi Quad 7 mc0 - medium changer device</pre> Note_____ If autoloader on the medium changer is set to On, the medium changer information might not appear. _____

Displaying information about tape drive connections to the storage system

To display information about tape drive connections to the storage system, complete the following steps.

Step	Action
1	Enter the following command: <pre>sysconfig -v</pre>
2	Look for the information about the SCSI adapter to which the tape devices are attached.

Example: This example shows a tape medium changer with SCSI ID 6 and a tape drive with SCSI ID 4 attached to slot 6 of the storage system. The SCSI firmware is 2.26, and the SCSI adapter clock rate is 60 MHz. For each drive, the following information appears:

- ◆ SCSI ID
- ◆ Vendor ID
- ◆ Product ID
- ◆ Firmware version

```
slot 6: SCSI Host Adapter 6 (QLogic ISP 1040B)
      Firmware Version 2.26      Clock Rate 60MHz.
6: BHTi      Quad 7      1.41
4: QUANTUM DLT7000      1B41
```

Understanding tape files

What a tape file is

All data on a tape is stored in tape files. Tape files are separated by file marks on the tape and have no names. You specify a tape file by its position on the tape.

Reading and writing tape files

You write a tape file using a tape device that you specify. When you read the tape file, you must specify a device that has the same compression type that you used to write it. The compression type is indicated by the last character in the tape device name, for example, the `a` in `nrst0a`.

For example, if you have two tape drives of the same model with tape drive numbers 1 and 2 and you write a file on `nrst0a`, you can read it safely on the other tape drive using `rst1a`, because devices of the same model have the same compression type (type `a`).

Controlling tape drives

About controlling tape drives

You can control tape drives with the `mt` command to move and position the tape in several ways. This section describes how to perform the following tasks:

- ◆ Move a tape to the end of data. You do this to append a backup.
- ◆ Skip forward over files. You do this to access a particular tape file.
- ◆ Skip backward over files. You do this to access a particular tape file.
- ◆ Append a backup. You do this to save tape if you have small backups.
- ◆ Rewind a tape. You do this to get to the beginning of the tape after using a no-rewind device.
- ◆ Take a tape drive offline. You do this to service a tape drive.
- ◆ Display status information. You do this to find out whether you can read the tape with a device or to verify that a tape drive is not in use.

The `mt` command syntax

The syntax of the `mt` command is as follows:

```
mt {-f|-t} device command [ count ]
```

The `-f` and `-t` options are interchangeable; they indicate that the next argument is a device.

device is a tape device.

command is an `mt` command. The following table discusses the most common `mt` commands.

count specifies the number of times to execute a command that supports multiple operations.

This section discusses the `mt` commands listed in the following table. For a description of all the `mt` commands, see the `mt(1)` man page.

Command	Task
<code>eom</code>	Position the tape to the end of the data or the end of the medium if the tape is full.
<code>fsf</code>	Move the tape forward, skipping a specified number of files.

Command	Task
<code>bsf</code>	Move the tape backward, skipping a specified number of files.
<code>rewind</code>	Rewind the tape.
<code>offline</code>	Rewind the tape and unload the tape medium if possible.
<code>status</code>	Display information about a device and the drive associated with it.

Note

Use `nrst` devices for all tape status and movement operations. Using other rewind types might produce unwanted side effects.

Caution

When you use the `urst` device with the `mt` command, the tape drives it is used with must be stackers and there must be sufficient tapes in the stacker sequence. Otherwise, the tape drives involved will terminate the command sequence or overwrite the same tape multiple times.

Moving a tape to the end of data

You move a tape to the end of data if you want to append data on a tape.

To move a tape to the end of data, or end of medium if the tape is full, complete the following step.

Step	Action
1	Enter the following command: <code>mt -f <i>device</i> eom</code> <i>device</i> is the name of a no-rewind tape device.

Example: `mt -f nrst0a eom`

Note

If you use a rewind or unload/reload tape device, this command rewinds the device, moves the tape to the beginning of data, and unloads it if possible.

Moving forward to a file

You move forward to access a particular tape file further along the tape. You can skip over a specified number of file marks and stop at the end-of-tape side of a file mark. This puts the tape drive head at the beginning of a file.

To move forward to the beginning of a tape file, complete the following step.

Step	Action
1	Enter the following command: <code>mt -f device fsf n</code> <i>device</i> is the name of a tape device used on the tape. <i>n</i> is the number of tape file marks you want to skip over going forward. The tape moves to the beginning of the <i>n</i> th file after the file where it is when the command is issued.

Example: If you enter the following command in the middle of the third file on the tape, it moves the tape to the beginning of the eighth file on the tape:

```
mt -f nrst0a fsf 5
```

Moving backward to the beginning of a file

You move backward to access a particular tape file further back on the tape.

To move backward to the beginning of a tape file, complete the following step.

Step	Action
1	Enter the following command: <code>mt -f device bsf n</code> <i>device</i> is the name of a tape device used on the tape. <i>n</i> is the number of tape file marks you want to skip over going backward. The tape moves to the beginning of the <i>n</i> th file before the file where it is when the command is issued.

Example: If you enter the following command in the middle of the fifth file on the tape, it moves the tape to the beginning of the second file on the tape:

```
mt -f nrst0a bsf 3
```

Rewinding a tape

You rewind a tape to get to the beginning of the tape when using a no-rewind device. You do this before performing a restore of a backup that was done using a no-rewind device, provided that there is no other data before the backed-up data on the tape and it is not unloaded.

To rewind a tape, complete the following step.

Step	Action
1	Enter the following command: mt -f <i>device</i> rewind <i>device</i> is the name of a tape device used on the tape.

Example: `mt -f nrst0a rewind`

Taking a tape drive offline

You take a tape drive offline to service it.

To rewind the tape and take the tape drive offline by unloading the tape, complete the following step.

Step	Action
1	Enter the following command: mt -f <i>device</i> offline <i>device</i> is the name of a tape device used on the tape.

Example: `mt -f nrst0a offline`

Displaying status information

You display status information to find out whether you can read with a device or to verify that a tape drive is not in use.

To display status information about a tape device and the drive associated with it, complete the following step.

Step	Action
1	Enter the following command: mt -f <i>device</i> status <i>device</i> is the name of a tape device.

Example: `mt -f nrst0a status`

Result: A display like the following appears:

```
Tape drive: Quantum DLT7000
Status: ready, write enabled
Format: 85937 bpi 70 GB (w/comp)
fileno = 0  blockno = 95603  resid = 0
```

The fields in the display are as follows:

- ◆ Tape drive—The model of the tape drive
- ◆ Status—Whether the tape drive is ready and write-enabled
- ◆ Format—The number of bytes per inch, total capacity in gigabytes, and whether data compression is used
- ◆ fileno—The current tape file number; numbering starts at 0
- ◆ blockno—The current block number
- ◆ resid—The number of bytes that the drive attempted to write or read but could not because the end of tape was reached

Qualifying tape drives dynamically

What a qualified tape drive is

A qualified tape drive is a tape drive that has been tested and found to work properly on Network Appliance storage systems. A qualified tape drive is in the Data ONTAP kernel's internal tape qualification list.

How to qualify a tape drive dynamically

You can qualify a nonqualified tape drive by entering a valid tape qualification file in the storage system's `/etc/tape_config` directory. You get the information you need for this file from the tape vendor's drive documentation. You can use a text editor to create the tape qualification file. There can be multiple files in the directory, one for each model being qualified.

How to format the tape qualification file

The `/etc/tape_config` directory contains a sample tape qualification file, which includes an explanation of the requirements for a tape qualification file, a list of the default SCSI command timeout values used by the tape drive, and an example of a tape qualification file.

When you create a new tape qualification file, you can use any valid file name, except a file name that begins with a period (.).

The format of the tape qualification file is as follows:

vendor_id (string, limit of 8 bytes)—The vendor ID as reported by the SCSI Inquiry command.

product_id (string, limit of 16 bytes)—The product ID as reported by the SCSI Inquiry command.

id_match_size (number)—The number of bytes of the product ID to be used for matching to detect the tape drive to be identified, beginning with the first character of the product ID in the Inquiry data.

vendor_pretty (string, limit 16 bytes)—If this parameter is present, it is specified by the string displayed by `sysconfig -v` or `sysconfig -t`; otherwise, the `INQ_VENDOR_ID` is displayed.

product_pretty (string, limit 16 bytes)—If this parameter is present, it is specified by the string displayed by `sysconfig -v` or `sysconfig -t`; otherwise, the `INQ_VENDOR_ID` is displayed.

Note

The *vendor_pretty* and *product_pretty* fields are optional, but if you specify a value for one of these fields, you must also specify a value for the other.

You must specify a description, density code, and compression algorithm for *l*, *m*, *h*, and *a*, as follows:

- ◆ `{l|m|h|a}_description=` (string, limit 16 bytes)—The string to print for `sysconfig -t` that describes characteristics of the particular density setting
- ◆ `{l|m|h|a}_density=` (hex codes)—The density code to be set in the SCSI mode page block descriptor corresponding to the desired density code for *l*, *m*, *h*, or *a*
- ◆ `{l|m|h|a}_algorithm=` (hex codes)—The compression algorithm to be set in the SCSI Compression Mode Page corresponding to the density code and the desired density characteristic

The following fields are optional:

- ◆ *autoload*= (Boolean yes/no)—Set this field to *yes* if the tape drive has an automatic loading feature; that is, after you insert a tape cartridge, the tape drive becomes ready without the need to execute a SCSI load (start/stop unit) command. The default for this field is *no*.
- ◆ *cmd_timeout_0x**—Individual timeout value. Use this field only if you want to specify a different timeout value from the one being used as a default by the tape driver. The sample file lists the default SCSI command timeout values used by the tape driver. The timeout value can be expressed in minutes (m), seconds (s), or milliseconds (ms).

Example of a tape configuration file

The following example shows a tape configuration file for the tape drive Tape Drive 2001.

```
vendor_id="TAPECO"
product_id="TAPEDRIVE2001"
id_match_size=11

vendor_pretty="Tape Company"
product_pretty="Tape Drive 2001"

l_description="2 GB"
l_density=0x01
l_algorithm=0x81
```

```

m_description="4 GB (compression)"
m_density=0x01
m_algorithm=0x82

h_description="6 GB"
h_density=0x10
h_algorithm=0x88

a_description="12 GB (compression)"
a_density=0x10
a_algorithm=0x89

autoload="yes"
# no entry for autoload is the same as autoload="no"

cmd_timeout_0x01=5m
cmd_timeout_0x12=700ms
cmd_timeout_0x1b=60s

```

How the storage system qualifies a new tape drive dynamically

The storage system's `/etc/tape_config` directory is created automatically when the storage system boots. When you add a tape qualification file to this directory, the storage system checks the file's format at the next boot time or the next time any tape is accessed. If the format is valid, the information is entered into the internal tape qualification table and a success message similar to the following is printed to the console and system log. Information about the tape persists as long as the file is in the directory or until the file is altered.

```
Dynamic Tape Qualification file:/etc/tape_config/filename format
check was successful. Internal tape qualification table was updated
with Tape Drive product information.
```

If the format is incorrect, an error message similar to one of the following is printed to the console and system log.

```
Dynamic Tape Qualification file: /etc/tape_config/filename has
missing or badly formatted required key(s).
Dynamic Tape Qualification file: /etc/tape_config/filename has a
format error in the information appended to the required key.
```

When you connect a tape drive to the storage system, the storage system looks for a *product_id* match between information obtained during the tape discovery process and information contained in the internal tape qualification table. If a match is discovered, the drive is marked as qualified and the storage system can access the tape drive. If a match is not found, the tape drive remains in the unqualified state and is not accessed.

Qualified SCSI drives must be attached to the storage system before reboot to be recognized; qualified Fibre Channel drives are recognized whenever they are attached.

Using nonqualified tape drives

How you can use a nonqualified tape drive

You can use a nonqualified tape drive (one that is not on the list of qualified tape drives) on a storage system if it can emulate a qualified tape drive. It is then treated as though it were a qualified tape drive.

Displaying information about nonqualified tape drives

To display information about nonqualified tape drives, complete the following steps.

Step	Action	
1	If the storage system...	Then...
	Has accessed the tape drive through the <code>dump</code> or <code>mt</code> command	Go to Step 3.
	Has not accessed the tape drive through the <code>dump</code> or <code>mt</code> command	Go to Step 2, then proceed to Step 3.
2	Enter the following command to access the tape drive: mt -f device status <i>device</i> is any device that contains the tape drive number that you think is assigned to the tape drive. Example: <code>mt -f nrst1a status</code>	
3	Enter the following command: sysconfig -t Result: If the storage system has registered a tape drive as emulating a qualified tape drive, it displays a message like the following: Tape drive (6.5) DLT9000 emulates Digital DLT7000. If the storage system has not registered a tape drive as emulating a qualified tape drive, it displays a message like the following: Tape drive (6.5) DLTXXXX (Non-qualified tape drive)	

Getting tape drive information for emulation

You must supply the following information about the nonqualified tape drive from your tape drive documentation to emulate a qualified tape drive:

- ◆ What qualified tape drive the nonqualified tape drive can emulate.
- ◆ The vendor ID string. This is a SCSI string and should be in the SCSI section of your tape drive manual.
- ◆ The product ID string. This is a SCSI string and should be in the SCSI section of your tape drive manual.

Note

Your nonqualified tape drive must be able to emulate a qualified tape drive for the emulation to succeed.

Emulating a qualified tape drive

To emulate a qualified tape drive, complete the following steps.

Step	Action
1	Make sure that you have a tape adapter available on the storage system.
2	Halt the storage system.
3	Connect the tape drive to the storage system according the tape drive manufacturer's instructions.
4	Turn on the tape drive.
5	Turn on the storage system. Result: When the storage system boots, an error message tells you that the tape drive is unsupported. It also contains the vendor ID and the product ID for the tape drive.
6	Record the vendor ID and product ID from Step 5 for later use.
7	Enter the following command to determine the highest tape drive number on your storage system: sysconfig -t

Step	Action
8	Enter the following command to create the storage system /etc/cloned_tapes file: <pre>mt -f nrstna status</pre> <i>n</i> is one plus the highest tape drive number from Step 7. Example: <code>mt -f nrst2a status</code>
9	Open the storage system's /etc/cloned_tapes file in a text editor on a client that can access it.
10	For each nonqualified tape drive, create a line with either of following formats: <pre><i>clone_vendor_ID clone_product_ID EMULATES vendor_ID product_ID</i></pre> <pre><i>clone_product_ID EMULATES product_ID</i></pre> <i>clone_vendor_ID</i> is the vendor of the nonqualified tape drive. <i>clone_product_ID</i> is the model number of the nonqualified tape drive. <i>vendor_ID</i> is the vendor of a qualified tape drive that you want the nonqualified tape drive to emulate. <i>product_ID</i> is the model number of a qualified tape drive that you want the nonqualified tape drive to emulate. Example: The following entry in the /etc/cloned_tapes file enables the storage system to treat the nonqualified Quantum DLT9000 tape drive as a clone of the qualified Quantum DLT7000 tape drive: <pre>QUANTUM DLT9000 EMULATES QUANTUM DLT7000</pre>
11	Enter the following command to put the emulation into effect: <pre>mt -f rstna status</pre> <i>n</i> is one plus the highest tape drive number from Step 7. Example: <code>mt -f rst2a status</code>

Step	Action
12	Verify that the drive appears, along with all your previously installed devices, by entering the following command: sysconfig -t

Enabling and disabling persistent reservations

What tape reservations are

Because multiple storage systems can share access to one tape drive or medium changer, it is advantageous to prevent more than one storage system from accessing a tape drive or a medium changer at a time. Tape reservation solves this problem by enabling SCSI-persistent reservations for all Fibre-Channel-attached tape drives and medium changers. These reservations provide a method for tape drives and medium changers to accept commands from only one storage system at a time.

These reservations cannot be enabled or disabled on an individual device basis.

Why tape reservations are not automatically enabled

Because not all tape drives and medium changers implement SCSI-persistent reservations correctly, this option is disabled by default to prevent problems from occurring with these devices.

How to enable persistent reservations

To turn on persistent reservations, complete the following step.

Step	Action
1	Enter the following command: <code>options tape.persistent_reservations on</code>

When to disable persistent reservations

Enabling this option may cause problems if tape drives or medium changers do not work properly. If tape commands report the device is reserved when no other storage systems are using the device, this option should be disabled.

How to disable persistent reservations

To turn off persistent reservations, complete the following step.

Step	Action
1	Enter the following command: <code>options tape.persistent_reservations off</code>

About this chapter

This chapter describes how to use the `dump` command to back up data on the storage system to tape.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Introduction to data backup to tape](#)” on page 46
- ◆ “[How the dump command works](#)” on page 50
- ◆ “[Preparing to use the dump command](#)” on page 52
- ◆ “[Using the dump command](#)” on page 56
- ◆ “[Appending backups to a tape](#)” on page 79
- ◆ “[Examples of the dump command](#)” on page 80
- ◆ “[Verifying the contents of a dump command backup](#)” on page 93
- ◆ “[Checking the status of a dump command backup](#)” on page 94
- ◆ “[Restarting a dump command backup](#)” on page 96
- ◆ “[Deleting restartable dump command backups](#)” on page 98

Introduction to data backup to tape

About data backup

Although data backup can also mean copying data to standard output, for this chapter, data backup means copying data from disk to tape.

Reasons to back up data from disk to tape

You back up data from disk to tape for the following reasons:

- ◆ You can store the backup tapes at an off-site archive to protect the data against natural disasters.
- ◆ You can restore data from tape if an application or a user inadvertently corrupts or deletes files that cannot be recovered using the Snapshot™ feature.
- ◆ You can restore data from tape after you reinstall the file system on the storage system (for example, when migrating to larger disks or converting a single-volume storage system to a multivolume storage system).

For detailed information

The following sections present information about the types of backups you can perform, how to perform them, and what gets backed up:

- ◆ [“Types and methods of backup”](#) on page 47
- ◆ [“What the dump command backs up”](#) on page 49

Introduction to data backup to tape

Types and methods of backup

- Types of backups
- You can do both full and incremental backups:
- ◆ In a full backup, everything is backed up.
 - ◆ In an incremental backup, only files changed since a previous backup are written to the backup media.

- Backup methods
- Data ONTAP supports three backup methods:
- ◆ Using the Data ONTAP `dump` command.
 - ◆ Using an NDMP-compliant backup application. For more information about NDMP, see Chapter 5, “[Using NDMP Services](#),” on page 141.
 - ◆ Using commands or a non-NDMP-compliant application from an NFS client or a CIFS client, such as the `tar` or `cpio` command on an NFS client.

Caution_____

If you use the `tar` or `cpio` command on an NFS client to back up the storage system, be aware that some versions of these commands fail with file systems that contain long path names, unusual file names, or hard links.

Comparison of backup methods

The following table compares the backup methods.

	storage system dump command	NDMP applications	Commands or non-NDMP applications on NFS or CIFS client
Where backup is initiated	Initiates the backup session from the storage system console or from a trusted host through the <code>rsh</code> command.	Initiate the backup session from the computer running the NDMP application.	Initiate the backup session from an NFS client or a CIFS client.

	storage system dump command	NDMP applications	Commands or non-NDMP applications on NFS or CIFS client
What is backed up	Backs up any files or directories.	Back up any files or directories.	Back up the files or directories that are mounted or shared by the client.
Where data is backed up	Backs up data to tape drives attached to the storage system being backed up, to another storage system, or to another computer that supports the rmt protocol.	Back up data to tape drives attached to the storage system being backed up, to the computer running the NDMP application, or to another storage system, as well as to NDMP- compliant tape libraries.	Back up data to tape drives attached to the client or, if the client supports the rmt protocol, to tape drives attached to the storage system. The details are application- dependent.
What attributes are preserved	Preserves both NFS and CIFS attributes.	Preserve both NFS and CIFS attributes.	Preserved attributes are application- dependent.

What the dump command backs up

What the dump command can back up

The `dump` command backs up data so that it can look the same when restored as it did when it was backed up.

The `dump` command can back up the following objects:

- ◆ A file
- ◆ A directory
- ◆ A qtree
- ◆ An entire volume

In addition to backing up data in files, the `dump` command can back up the following information about each file, as applicable:

- ◆ UNIX GID (Group ID), owner UID (User ID), and file permissions
- ◆ UNIX access, creation, and modification time
- ◆ File type
- ◆ File size
- ◆ DOS name, DOS attributes, and creation time
- ◆ Windows NT Access Control Lists (ACLs)
- ◆ Qtree information

Dump command data format

The data is written to tape in a format compatible with the standard BSD-dump format. This means that you can restore Network Appliance dump tapes using the Solaris `ufsrestore` command. The only data that you cannot restore with this command is the CIFS (Windows NT) attributes, because Solaris does not support them. However, the data itself is accessible.

Dump and restore permissions limitations

If you dump a file that has only Windows NT permissions and restore it into a UNIX-style qtree or volume, the file gets the default UNIX permissions for that qtree or volume.

If you dump a file that has only UNIX permissions and restore it into an NTFS-style qtree or volume, the file gets the default Windows permissions for that qtree or volume. Other dumps and restores preserve permissions.

How the dump command works

What the dump command does

The `dump` command writes file system data from disk to tape. It is optimized for data restoration to a storage system using the Data ONTAP `restore` command.

Kinds of backups

There are three kinds of backups:

- ◆ Full volume backup, where you back up an entire volume
- ◆ Full qtree backup, where you back up an entire qtree
- ◆ Miscellaneous backup, where you back up a subtree that is neither an entire volume nor an entire qtree

The backup process

Data ONTAP uses the process described in the following table to back up the object indicated by the dump path.

Stage	Action
1	For less than full volume or full qtree backups, Data ONTAP traverses directories to identify files to be backed up. If you are backing up an entire volume or qtree, Data ONTAP combines this step with Step 2.
2	For a full volume or full qtree backup, Data ONTAP identifies the directories in the volumes or qtrees to be backed up.
3	Data ONTAP writes the directories to tape.
4	Data ONTAP writes the files to tape.
5	Data ONTAP writes the access control list (ACL) information (if applicable) to tape.

How the dump command uses Snapshot copies in backing up data

The `dump` command uses a Snapshot copy of your data for the backup, so you do not need to take the storage system or volume offline before initiating the backup.

The `dump` command names each Snapshot copy it creates `snapshot_for_backup.n`.

- ◆ *n* is an integer starting at 0.

- ◆ Each time the `dump` command creates a Snapshot copy, it increments the integer by 1.
- ◆ The storage system resets the integer to 0 when it is rebooted.

When Data ONTAP executes multiple `dump` commands simultaneously, the `dump` commands create multiple Snapshot copies. For example, if Data ONTAP is running two `dump` commands simultaneously, you find these Snapshot copies in the volumes from which data is being backed up: `snapshot_for_backup.0` and `snapshot_for_backup.1`.

Note

When you are backing up from a Snapshot copy, `dump` does not create an additional Snapshot copy.

Preparing to use the dump command

About dump command preparations

Before backing up data with the `dump` command, you need to make preparations. You must have sufficient tapes to store the data and know how you will enter the `dump` command.

For detailed information

The following topics discuss what you need to do to prepare for using the `dump` command:

- ◆ [“Determining the amount of backup data”](#) on page 53
- ◆ [“Estimating the number of tapes for the backup”](#) on page 54
- ◆ [“Deciding how to enter the dump command”](#) on page 55

Determining the amount of backup data

About determining the amount of backup data

Before you enter the `dump` command, it is helpful to estimate the amount of backup data so that you can determine the number of tape files and the number of tapes required for the backup.

Note _____
This method of estimation works only for full backups of volumes. It is impossible to determine how much data is being backed up in an incremental backup.

Estimating the amount of backup data in a volume

To estimate the amount of data in a full backup, complete the following steps.

Step	Action
1	For each item that you want to back up, enter the following command: <code>df path_name</code> <i>path_name</i> is the name of the item.
2	For multiple items, such as multiple volumes, add up the data for each item to determine the total amount of data to be backed up.

Estimating the number of tapes for the backup

Reasons to determine the number of tapes

You should estimate the number of tapes required for the backup before entering the `dump` command for these reasons:

- ◆ You can ensure that the `dump` command does not fail because it runs out of tapes.
- ◆ You can load the required number of tapes in the tape drives or stackers in advance for an unattended backup. If you do not load enough tapes before entering the `dump` command and you start the `dump` command from the console, Data ONTAP prompts you to load additional tapes. But if you start the `dump` command through a Remote Shell connection, you do not see the prompts from Data ONTAP and the `dump` command terminates if there are not enough tapes.

Estimating the number of tapes

To estimate the number of tapes required for the backup, complete the following steps.

Step	Action
1	Determine the capacity of the tape device you are using for the backup by entering the following command: <code>sysconfig -t</code>
2	Determine the amount of data that needs to be backed up, as described in “ Determining the amount of backup data ” on page 53.
3	Divide the amount of data by the capacity of a tape.
4	If your estimate indicates that your data will nearly fill the last tape, add a tape to the estimate in case the backup exceeds your estimate. This is especially important when using compression, because compression rates vary based on the data.

Deciding how to enter the dump command

About entering the dump command

You can enter the `dump` command through a Remote Shell connection, such as through the `rsh` command, through a Telnet session accessing the storage system console, or through the storage system console directly.

Other than potential problems associated with any remote connection, console access through a Telnet session and direct console connection to the storage system behave the same way.

Benefits of entering the dump command through a Remote Shell connection

Entering the `dump` command through a Remote Shell connection gives you these benefits:

- ◆ When the `dump` command is in progress, you can still use the console to manage the storage system. If the `dump` command entered on the console is backing up a large number of files, you cannot use the console for a long time.
- ◆ You can start multiple `dump` commands through the `rsh` command.
- ◆ Data ONTAP is less likely to inadvertently terminate the `dump` command, especially if it is run in the background from a Solaris system. If you enter a `dump` command on the storage system console, it could be terminated by Ctrl-C entered on a host connected to the storage system using a Telnet session.
- ◆ You can automate storage system backups through shell scripts and crontab entries.

Benefits of entering the dump command at the console

If you enter the `dump` command at the console, you can read and respond to screen messages and prompts displayed by the command. For example, the command might prompt you for another tape to complete the backup, whereas a `dump` command entered through a Remote Shell connection does not generate any messages when the command needs user intervention, and terminates instead.

Using the dump command

About the dump command

You can enter the `dump` command any time the tape devices you want to use are free to back up data in a specified path. After the `dump` command is finished, the data in the path is written to tape.

You can run up to eight `dump` commands (depending on the hardware you are using) in parallel on up to eight tape drives, one command per drive. Parallel backups increase throughput.

For detailed information

The `dump` command is complex and has many options. The following sections discuss how to construct a `dump` command:

- ◆ [“Dump command syntax”](#) on page 57
- ◆ [“Recommendations for performing a backup”](#) on page 60
- ◆ [“Specifying the backup level”](#) on page 63
- ◆ [“Understanding increment chains”](#) on page 64
- ◆ [“Specifying the tape device”](#) on page 66
- ◆ [“Specifying the dump path”](#) on page 68
- ◆ [“Specifying a list of files”](#) on page 69
- ◆ [“Excluding items from a backup”](#) on page 70
- ◆ [“Specifying a name for a backup”](#) on page 72
- ◆ [“Specifying no Windows NT ACLs”](#) on page 73
- ◆ [“Specifying the blocking factor”](#) on page 74
- ◆ [“Specifying the tape file size”](#) on page 76
- ◆ [“Specifying updates to the `/etc/dumpdates` file”](#) on page 77

Dump command syntax

Syntax of the dump command

The `dump` command syntax is as follows:

`dump options arguments dump_path`

Note

Not all options are mandatory, and some do not have any arguments.


Rules for entering the dump command

The following list describes the rules for entering the `dump` command:

- ◆ You can list one or more options. You must list all options together; do not separate the options by commas or spaces.
- ◆ You can list options in any order.
- ◆ You must include a backup level and a tape file in the options.
- ◆ *arguments* can be one argument or a list of arguments, each of which is associated with an option.
- ◆ List all arguments in the same order as their corresponding options.
- ◆ Separate each argument with one or more spaces.
- ◆ If the argument is a list, use commas to separate the items in the list.
- ◆ *dump_path* is the complete path name of the volume, directory, or qtree batch file to be backed up by the `dump` command.
- ◆ Always precede the volume name by `/vol/` even if the volume is a root volume, because between different levels of backups, you might have changed the root volume.

Example of a simple dump command

The following figure illustrates the dump command syntax:



Dump options

The following table summarizes the meanings of the dump command options. For general recommendations on using them, read [“Recommendations for performing a backup”](#) on page 60.

Option	Meaning	For more information, see...
backup level	Level 0 is a full backup; levels 1 through 9 are for incremental backups.	“Specifying the backup level” on page 63 and “Understanding increment chains” on page 64
A	Do not back up Windows NT ACLs.	“Specifying no Windows NT ACLs” on page 73
b	The blocking factor. Argument: The number of 1-KB blocks in each write. For a storage system, the range is 4 through 64, and the default is 63.	“Specifying the blocking factor” on page 74
B	Write the specified number of tape blocks to a tape file before starting a new tape file. Argument: The number of tape blocks in a tape file.	“Specifying the tape file size” on page 76

Option	Meaning	For more information, see...
f	Specify the tape device for the backup. It is mandatory. Argument: At least one tape device name as an argument. Separate additional tape device names with commas.	“Specifying the tape device” on page 66
l	Back up only specific files and directories in the dump path. You must use the n option when using the l option.	“Specifying a list of files” on page 69
n	Provide a name for the backup to be recorded in the /etc/dumpdates file. It takes a string as an argument. It is required if you use the l option.	“Specifying a name for a backup” on page 72
Q	Back up all data in the specified volume that does not reside in a qtree.	“Excluding all data in qtrees” on page 70
u	Update the /etc/dumpdates file. You must use this option if you plan to perform incremental backups in the future.	“Specifying updates to the /etc/dumpdates file” on page 77
x	Exclude specified files from the backup. It takes a string specifying exclusion prefixes or suffixes as an argument.	“Excluding specified files” on page 71

Recommendations for performing a backup

General recommendations

Avoid backing up too much data in a single `dump` command to reduce the risk of encountering an error that requires restarting a `dump` command. If the `dump` command encounters an error, however, you might be able to correct the error and proceed from the point where the command terminated. (See “[Restarting a dump command backup](#)” on page 96.)

If the storage system console appears to be hung during a backup, it is because a backup can take a long time. The console becomes responsive and the prompt returns when the backup is completed.

Recommendations for minimizing backup time and data loss

You can minimize both the time required to perform a backup and the possibility of data loss. The shorter the time for the `dump` command to finish, the more incremental backups you can perform.

- ◆ Perform frequent incremental backups to minimize the amount of unrecoverable data in case of errors.

Note

There is a disadvantage to having a large number of incremental backups. When you restore data, you must restore from all the incremental backup tapes, which requires running multiple restores and manipulating multiple tape sets.

- ◆ Use local tape drives. The storage system can write faster to a local tape drive than to a tape drive attached to a remote system.
- ◆ Organize data to be backed up.

The `dump` command runs faster if the dump path specifies one of the following:

- ❖ A full volume
- ❖ Full qtrees
- ❖ A full volume excluding qtrees

Recommendations for decreasing tape backup time

Follow these guidelines to decrease tape backup time:

- ◆ Divide large volumes into smaller volumes or qtrees.
For example, if you divide a 500-GB volume into three qtrees, you can back up each qtree to a separate tape drive or run separate full backups on three different nights.
- ◆ Limit the amount of data in a volume or qtree to be backed up to 200 GB.
- ◆ Schedule the backups in appropriate rotations.
- ◆ Schedule backups when the load on the storage system is light.
- ◆ Do not divide a backup into more than 15 qtrees.

Recommendations for minimizing the number of tape drives required

Because Data ONTAP supports the `rmt` protocol, several storage systems can share the same tape drive. Attach the tape drive to the storage system with the most data to back up. Follow these guidelines if multiple storage systems back up to the same tape drive:

- ◆ Use a private network for the backup so that the traffic load on the network does not slow down the backup process.
- ◆ Schedule the `dump` command on each storage system so that it starts only when no other storage systems are using the tape drive.

Recommendations for recording information for restores

For ease of use during a restore, you should label the backup tapes with the following information:

- ◆ The dump path of each backup on a tape.
- ◆ The level of each backup on a tape.
- ◆ The date of each backup.
- ◆ The blocking factor. This must match for backups and restores.
- ◆ Tape file contents of a multifile tape. A brief description of the contents of each tape file on a multifile tape helps you locate a desired tape file for restoring.
- ◆ The sequence of tape files on a multifile tape. This enables you to specify which file to restore. To specify a tape file, you need to know where in the sequence of tape files it is.
- ◆ The Data ONTAP version of each backup.

Recommendation for verifying backups

If you suspect that a backup might have errors, verify the backup by performing a test restore, as described in “[Specifying a test restore](#)” on page 130. For information about restoring information from tape backups, see “[Data Restoration from Tape](#)” on page 99.

Using the dump command

Specifying the backup level

- What a backup level is for**

You use a backup level to specify whether all files or only the most recently changed files are backed up.
- Kinds of backup levels**

A level-0 backup is a full backup. A full backup backs up all the data in the dump path.

Backups at levels from 1 through 9 are incremental backups. An incremental backup backs up only the items in the dump path that have been created or changed since the most recent backup of a lower level.

Specifying a backup level To specify the backup level, complete the following step.

Step	Action
1	Include the level number as an option. The range is 0 through 9.

Example: The following command performs a level-0 backup of the /vol/vol0 volume:

```
dump 0f rst0a /vol/vol0
```

↑
0 specifies a level 0 backup.

Note Incremental updates do not run unless the baseline transfer has updated the dumpdates file.

Understanding increment chains

Definition of an increment chain

An increment chain consists of a series of incremental backups of the same path. Because you can specify any level of backup at any time, you must understand increment chains to be able to perform backups that result in the files that you want being restored.

Types of increment chains

There are two types of increment chain:

- ◆ A consecutive increment chain is a sequence of incremental backups that starts with level 0 and is incremented by one. If this is how you do incremental backups, you do not need to read the rest of this section.
- ◆ A nonconsecutive increment chain is one in which incremental backups skip levels or have levels that are out of sequence. An example of such a chain is one where the levels have been specified in the following order: 0, 2, 3, 1, 4, or more commonly, 0,1,1,1 or 0,1,2,1,2.

Incremental backup sources

Incremental backups base themselves on the most recent lower-level backup. For example, the sequence of backup levels 0, 2, 3, 1, 4 gives two increment chains: 0, 2, 3 and 0, 1, 4. The bases of the incremental backups are shown in the following table.

Back-up order	Increment level	Increment chain	Base	Files backed up
1	0	Both	Files on the storage system	All files in the backup path
2	2	0, 2, 3	The level-0 backup	Files in the backup path created since the level-0 backup
3	3	0, 2, 3	The level-2 backup	Files in the backup path created since the level-2 backup

Back-up order	Increment level	Increment chain	Base	Files backed up
4	1	0, 1, 4	The level-0 backup, because that is the most recent level that is lower than the level-1 backup	Files in the backup path created since the level-0 backup, including files that are in the level-2 and level-3 backups
5	4	0, 1, 4	The level-1 backup, because it is both of a lower level and more recent than the level-0, level-2, or level-3 backups	Files created since the level-1 backup

Using the dump command

Specifying the tape device

About specifying tape devices

You must specify at least one tape device to do the backup. If you specify more than one tape device, each tape device in the list is used in the order listed to write a tape file. If the backup requires more tape files than the number of tape devices specified, the last tape device is used for all remaining tape files.

Caution_____

If you specify more than one rewind device on the same tape drive, the storage system displays a warning and terminates the `dump` command.

Note_____

storage system device names might not be valid on remote tape drive hosts. For tape drives attached to remote hosts, use tape device names that follow the host naming conventions.


You can specify two kinds of tape devices: local and remote.

Specifying a local tape device

To specify local tape devices for a backup, complete the following step.

Step	Action
1	Use the <code>f</code> option with a list of one or more tape devices, separated by commas, as an argument. Note_____ You cannot combine local and remote tape devices in one command, and you can write to only one remote machine in a command.

Example: The following command specifies that one tape file be written with one device:




Specifying a tape device on a remote storage system

To specify tape devices for a backup on a tape drive attached to a remote storage system, complete the following step.

Step	Action
1	Use the <code>f</code> option with a list of one or more tape devices, separated by commas, as an argument. Do not repeat the remote machine name. Note _____ You cannot combine local and remote tape devices in one command, and you can write to only one remote machine in a command. _____

Example: The following command specifies that two tape files, one using `nrst0a` and the other using `rst0a`, be written on the remote storage system named `remfil`:



Specifying the dump path

What the dump path does The dump path specifies one volume, qtree, or subtree to back up. A subtree is a directory in a volume or qtree.

Specifying a single path name You can specify a dump path by specifying a volume, qtree, or subtree to back up all the data in it. The volume, qtree, or subtree can be in either of the following locations:

- ◆ The active file system, for example, /vol/volname/home
- ◆ A Snapshot copy, for example, /vol/volname/.snapshot/weekly.0/home

To specify a single dump path, complete the following step.

Step	Action
1	Put the path name of the volume, qtree, or subtree that you want to back up at the end of the <code>dump</code> command.

Example: The following command contains the dump path /vol/vol0:

```
dump 0f rst0a /vol/vol0
```

↑
The dump path is /vol/vol0.

Specifying a list of files

About specifying a list of files

You can specify for backup a list of one or more files. The files must all be in the same dump path.

You use this method to back up some, but not all, subdirectories or files in the dump path with only one `dump` command. Using a list is easier than using a `dump` command for each subdirectory or file, and it is quicker because it avoids the overhead of starting up multiple `dump` commands.

Note Because the `dump` command reads the files from standard input, you must use a script to specify a list of files if you use a Remote Shell connection.

Specifying the file list

To specify a list of files, complete the following steps.

Step	Action
1	Use the <code>n</code> and <code>l</code> options in the command line.
2	Include a name for the backup as an argument to the <code>n</code> option.
3	Put the path name of the volume, <code>qtree</code> , or subtree that you want to back up at the end of the <code>dump</code> command.
4	Enter the <code>dump</code> command line.
5	In response to prompts, enter each name as a path name relative to the dump path in the <code>dump</code> command. Note Do not specify a parent directory (<code>..</code>) or a directory that is a symbolic link.
6	To end the list, press the Enter key.

Using the dump command

Excluding items from a backup

What you can exclude from a backup

You can exclude the following items from a backup:

- ◆ Data in qtrees
- ◆ Specified files

About excluding all data in qtrees

You can back up all data in a specified volume that is not in a qtree. The specified volume is the dump path. You use this method if you back up on a qtree basis and want to back up the remaining data in a volume. Usually, the data in qtrees changes frequently, while the remaining data, such as configuration files, changes rarely.


Note _____
You cannot do incremental backups using this method.

Excluding all data in qtrees

To back up all non-qtree data in a specified volume, complete the following step.

Step	Action
1	Use the <code>Q</code> option in the command line.

Example: The following command backs up all items in `/vol/vol0` that are not in a qtree:



About excluding specified files

You can exclude a list of files and all files that match a specified pattern. You specify a pattern to avoid backing up classes of files that you don't want, such as files that end with .core.

Rules for excluding specified files

The rules for constructing a string for excluding files are as follows:


- ◆ A string can be a file name.
- ◆ You can use the asterisk (*) as a wildcard character.
- ◆ The wildcard character must be the first or last character of the string. Each string can contain up to two wildcard characters. For example, you can specify *.core, core.*, or *core.*, but not core*.1.
- ◆ If you have more than one string, you must separate the strings with a comma.
- ◆ You cannot have a comma in the file name or pattern.
- ◆ You can specify up to 32 strings.

Excluding specified files

To exclude files from a backup, complete the following steps.

Step	Action
1	Use the x option in the command line.
2	Include a string or comma-separated list of strings as an argument to the x option.

Example: The following command backs up all items in /vol/vol0 except for the file dontbackup and all files that end in .core:



Specifying a name for a backup

When to specify a name for a backup

You can name a backup and, if you use the `u` option, have the name of the backup recorded in the `/etc/dumpdates` file. You do this in two situations:

- ◆ You are specifying a list of directories or files in the backup with the `l` option.
- ◆ You want to monitor the backup history of a backup.

Naming a backup

To specify a name for a backup, complete the following steps.

Step	Action
1	Include the <code>n</code> option in the <code>dump</code> command line.
2	Include a name for the backup as an argument to the <code>n</code> option.

Example: The following command gives the name `thisbackup` to a backup:

```
dump 0fn rst0a thisbackup /vol/vol0
```



`n` specifies to name this backup.



The name of the backup is `thisbackup`.

Result: A line like the following appears in the `/etc/dumpdates` file:

```
thisbackup 0 Tue Jul 24 20:40:09 2001
```

Using the dump command

Specifying no Windows NT ACLs

Reasons to omit Windows NT ACLs

You can omit Windows NT ACLs (access control lists) from a backup. This provides a slight performance enhancement. You omit Windows NT ACLs in two situations:

- ◆ You plan to restore to a volume that will be used in a UNIX environment or to a pre-5.1 version of Data ONTAP.
- ◆ You are backing up no files or directories that contain Windows NT ACLs.

Omitting Windows NT ACLs

To omit Windows NT ACLs from a backup, complete the following step.

Step	Action
1	Include the A option in the dump command line. Note_____ This option does not take an argument. _____

Example: The following command backs up the /vol/vol0 volume without backing up Windows NT ACLs:

```
dump 0fA rst0a /vol/vol0
```

↑
A specifies that no Windows NT ACLs will be backed up.

Specifying the blocking factor

What the blocking factor is

A tape block is 1,024 bytes of data. In the `dump` command, you can specify the number of tape blocks that are transferred in each write operation. This number is called the blocking factor.

Blocking factor range on a storage system

If you use a storage system for both the backup and the restore, you do not need to specify a blocking factor. Network Appliance recommends that you use the default blocking factor of 63. On a storage system, you can specify a blocking factor from 4 through 64.

Note

Dump backups do not allow tape block sizes bigger than 64KB. An attempt to use block sizes of greater than 64KB will result in failed backups.

Blocking factor choice on other machines

On a remote host that is not a storage system, you can use a blocking factor from 4 through 256, provided that the host supports the blocking factor that you choose.

Note

SunOS 4.x and earlier versions do not support a blocking factor greater than 63.

If you plan to restore to a system other than the system that did the backup, the restore system must support the blocking factor that you used for the backup. For example, if you use a blocking factor of 128, the system on which you restore that backup must support a blocking factor of 128. A blocking factor of 128 would mean that you could not restore the backup using a tape drive that is connected to a storage system.

Specifying a blocking factor

To specify a blocking factor for a backup, complete the following step.

Step	Action
1	Include the <code>b</code> option in the <code>dump</code> command line. Note_____ This option does not take an argument.

Example: The following command backs up the `/vol/vol0` volume using a blocking factor of 32, so that 32 KB of data are transferred at a time:

`dump 0fb rst0a 32 /vol/vol0`

↑
b specifies that
a blocking factor
is in the command.

↑
The blocking
factor is 32.

Using the dump command

Specifying the tape file size

Maximum tape file sizes

In the `dump` command, you can specify the maximum size of the tape file in terms of tape blocks. For example, if you want the maximum tape file to be 2 GB, specify 2,097,151. This specifies that the largest tape file can contain 2,097,151 tape blocks, which are 1 kilobyte each. The tape file size must be equal to or greater than the blocking factor; otherwise, the dump process terminates.

When to specify the tape file size

If you do a remote backup or plan to restore the backup on a system other than the storage system, you might need to specify a tape file size. Some systems support only tape files of a limited size; for example, some Solaris systems do not support tape files larger than 2 GB.

Specifying the tape file size

To specify a tape file size, complete the following steps.

Step	Action
1	Include the <code>B</code> option in the <code>dump</code> command line.
2	Include the tape file size, in kilobytes, in the <code>dump</code> command as an argument to the <code>B</code> option. The size applies to all tape files in the backup.

Example: The following command backs up the `/vol/vol0` volume using a tape file size of 2,097,151, so that a tape file is no larger than 2 GB:

`dump 0fB rst0a 2097151 /vol/vol0`

↑
B specifies that
the file size is
in the command.

↑
The file size is
2,097,15KB

Specifying updates to the */etc/dumpdates* file

What the */etc/dumpdates* file is for

The */etc/dumpdates* file enables you to keep track of backups. It records the following information:

- ◆ The name of the backup. The name is one of the following:
 - ❖ If you use the *n* option, the name you supply
 - ❖ If you use the *Q* option, the volume you are backing up followed by the notation */all_non_quota_files*
 - ❖ If you use neither, the dump path
- ◆ The level of the backup.
- ◆ The time of the Snapshot copy used for the backup.

Why to update the */etc/dumpdates* file

You update the */etc/dumpdates* file for the following reasons:

- ◆ You plan to perform incremental backups. The storage system uses the data in the */etc/dumpdates* file to determine what to include in incremental backups.
- ◆ You want to keep a history of a backup.

Principles applying to the */etc/dumpdates* file

The following principles apply to the */etc/dumpdates* file:

- ◆ If the */etc/dumpdates* file does not exist when you try to update it, the storage system creates it.
- ◆ You can edit the */etc/dumpdates* file manually, if needed.
- ◆ A new backup of the same path and level overwrites the old entry.

Example of an */etc/dumpdates* file

An */etc/dumpdates* file lists one backup per line. Each line contains the name of the backup, followed by the level of the backup, then the date of the backup. An */etc/dumpdates* file looks like the following:

```
/vol/vol1/ 0 Tue Jul 24 22:07:48 2001
/vol/vol0/ 0 Tue Jul 24 21:06:53 2001
/vol/vol0/etc 0 Tue Jul 24 19:06:15 2001
my_named_dump 0 Tue Jul 24 20:40:09 2001
/vol/vol0/all_non_quota_files 0 Tue Jul 24 20:54:06 2001
```

```
/vol/vol0/home 0 Tue Jul 24 21:06:39 2001
/vol/vol1/ 1 Tue Jul 24 22:08:09 2001
/vol/vol1/ 2 Tue Jul 24 22:08:20 2001
my_named_dump 1 Tue Jul 24 22:12:26 2001
/vol/vol0/home 5 Tue Jul 24 22:12:45 2001
```

**Updating the
/etc/dumpdates file**

To update the /etc/dumpdates file, complete the following step.

Step	Action
1	Include the <code>u</code> option in the <code>dump</code> command line. Note _____ This option does not take an argument.

Example: The following command backs up the /vol/vol0 volume and adds the backup information to the /etc/dumpdates file:

`dump 0fu rst0a /vol/vol0`



`u` updates the */etc/dumpdates* file.

If the command is issued on Monday, April 16, 2001, at 45 seconds after 1:12 p.m., the following line is added to the /etc/dumpdates file:

```
/vol/vol0/ 0 Mon Apr 16 13:12:45 2001
```

Appending backups to a tape

Why to append backups

If you are backing up small volumes, qtrees, or files, you might want to put several backups on one tape to conserve tapes. Also, adding each backup to the tape after the last backup ensures that backups are sequential.

Appending a backup to a tape

To append a backup to a tape, complete the following steps.

Step	Action
1	Move the tape to the desired location using the <code>mt</code> command. See “ Controlling tape drives ” on page 29.
2	Execute a <code>dump</code> command. Caution _____ Use no-rewind device names to make sure that the tape is not rewound and that previous backups are not overwritten. _____

Examples of the dump command

About the examples Examples of how to perform the following tasks can help you use the `dump` command.

For detailed information You can use the following examples of the `dump` command to understand the `dump` command and as boilerplate for performing backups.

- ◆ [“Level-0 backups to a local tape drive”](#) on page 81
- ◆ [“Incremental backups to a local tape drive”](#) on page 83
- ◆ [“Backups to a remote tape drive”](#) on page 84
- ◆ [“Backups to multiple tapes”](#) on page 86
- ◆ [“Backup of a directory from a Snapshot copy”](#) on page 88
- ◆ [“Backup of a list of files or directories”](#) on page 89
- ◆ [“Backup data without ACLs”](#) on page 90
- ◆ [“Backup specifying a blocking factor”](#) on page 91
- ◆ [“Backup excluding specified files”](#) on page 92

Examples of the dump command

Level-0 backups to a local tape drive

- Examples of level-0 backups**
- This section provides examples of the following kinds of level-0 backups:
- ◆ A full backup
 - ◆ A full backup with a record
 - ◆ A full backup minus qtrees

Full backup

The following command performs a full backup of the /vol/vol1/users/tom/specs directory. After the dump command is finished, the tape drive rewinds the tape.

```
dump 0f rst0a /vol/vol1/users/tom/specs
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
f	A tape device is supplied in the command line.
rst0a	The tape drive rewinds.
/vol/vol1/users/tom/specs	This is the directory to be backed up.

Full backup with a record

The following command performs a full backup of the /vol/vol1/users/tom directory and records the backup in the /etc/dumpdates file. The tape drive rewinds the tape.

```
dump 0uf rst0a /vol/vol1/users/tom
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.

Element	Meaning
u	Record the backup in the /etc/dumpdates file.
f	A tape device is supplied in the command line.
rst0a	The tape drive rewinds.
/vol/vol1/users/tom	This is the directory to be backed up.

Full backup minus qtrees

The following command performs a full backup of all non-qtree data in the /vol/vol1 volume. The tape drive rewinds the tape.

```
dump 0fQ rst0a /vol/vol1
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
f	A tape device is supplied in the command line.
Q	Do not back up data that is in a qtree.
rst0a	The tape drive rewinds.
/vol/vol1	This is the volume to be backed up.

Incremental backups to a local tape drive

Level-1 backup

The following command performs a level-1 backup of the /vol/vol1 volume to nrst0a. This means that only files that have changed since the most recent level-0 backup of the same path are written to tape. After the command is finished, it does not rewind the tape.

```
dump 1uf nrst0a /vol/vol1
```

The elements of this command line are described in the following table.

Element	Meaning
1	Do a level-1 backup.
u	Record the backup in the /etc/dumpdates file.
f	A tape device is supplied in the command line.
nrst0a	The tape drive does not rewind.
/vol/vol1	This is the volume to be backed up.

Examples of the dump command

Backups to a remote tape drive

Examples of backups to a remote tape drive

This section provides examples of backups to the following kinds of remote tape drives:

- ◆ A tape drive attached to a remote storage system
- ◆ A tape drive attached to a Solaris system
- ◆ A tape drive attached to a Solaris system with a 2-GB tape file limit

Tape drive attached to a remote storage system

The following command performs a backup to a tape drive attached to a remote storage system named netapp1. The tape drive does not rewind the tape.

```
dump 0f netapp1:nrst0a /vol/vol1
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
f	A tape device is supplied in the command line.
netapp1	This is the name of the storage system that the tape drive is attached to.
nrst0a	The tape drive does not rewind.
/vol/vol1	This is the volume to be backed up.

Tape drive attached to a Solaris system

The following command performs a backup to a tape drive on a Solaris system. The tape drive rewinds the tape.

```
dump 0f ritchie:/dev/rmt/0 /vol/vol1
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
f	A tape device is supplied in the command line.
ritchie	This is the name of the Solaris machine to which the tape drive is connected.
/dev/rmt/0	This is the name of the tape drive. Tape drive names vary according to the type of Solaris system you use.
/vol/vol1	This is the volume to be backed up.

Tape drive attached to a Solaris system with a 2-GB tape file limit

The following command performs a backup to a tape drive on a Solaris system. The Thompson system has a 2-GB limit on tape file sizes. The size of the backup is greater than 2 GB but less than 4 GB, so the backup must be broken up into two tape files.

```
dump 0fB thompson:/dev/rmt/0n,/dev/rmt/0n 2097151 /vol/vol1
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
f	A tape device is supplied in the command line.
B	A tape file size maximum is supplied in the command line.
thompson	This is the name of the Solaris machine to which the tape drive is connected.
/dev/rmt/0n	This is the name of the remote tape drive.
rst0a	The tape drive rewinds after the second file.
2097151	The maximum tape file size is 2 GB.
/vol/vol1	This is the volume to be backed up.

Examples of the dump command

Backups to multiple tapes

Examples of backups to multiple tapes

This section provides examples of the following kinds of multiple-tape backups:

- ◆ A backup using multiple tapes on single-tape drives
- ◆ A backup using multiple tapes on a stacker

Multiple tapes on single-tape drives

The following command backs up the /vol/vol1 volume to the tape on tape drive 0 using rst0a. If the backup exceeds the capacity of the tape, the command writes the second tape file on tape drive 1 using rst1a. Otherwise, it does not use the tape on rst1a.

```
dump 0f rst0a,rst1a /vol/vol1
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
f	A tape device is supplied in the command line.
rst0a	The first tape drive rewinds.
rst1a	The second tape device rewinds.
/vol/vol1	This is the directory to be backed up.

Multiple tapes on a stacker

The following command backs up the /vol/vol1 volume to the tape on urst0a. After the command finishes writing to the first tape drive, it rewinds and unloads the tape. If the volume exceeds the capacity of the tape on urst0a, the command writes the second tape file to the same tape drive on urst0a. If the command needs to write a third tape file, it prompts you to load a new tape.

```
dump 0f urst0a,urst0a /vol/vol1
```

Caution

If you do not use a stacker when you specify an unload/reload device, the tape drive will unload and reload the same tape, overwriting the tape.

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
f	A tape device is supplied in the command line.
urst0a,urst0a	The tape stacker rewinds and unloads the first tape. If a second tape is needed, it loads the second tape.
/vol/vol1	This is the directory to be backed up.

Backup of a directory from a Snapshot copy

Directory backup from a Snapshot copy

From the weekly Snapshot copy, the following command performs a level-0 backup of the /vol/vol1/home/users directory:

```
dump 0f rst0a /vol/vol1/.snapshot/weekly.0/home/users
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
f	A tape device is supplied in the command line.
rst0a	The tape drive rewinds.
/vol/vol1/.snapshot/weekly.0/home/users	This is the directory to be backed up.

Examples of the dump command

Backup of a list of files or directories

List of files or directories

The following example shows the prompts and path name entry when you back up a list of files or directories. The example ends the list of path names with a blank line.

```
dump 0ufnl rst0a user.1.3.5 /vol/vol1/home

DUMP: creating "snapshot_for_backup.0" snapshot.
creating.....
DUMP: Date of this level 0 dump: Tue Jun  4 12:47:14 2001
DUMP: Date of last level 0 dump: Tue May 28 4 12:45:51 2001
DUMP: Dumping /vol/vol0/home to nrst0a
DUMP: mapping (Pass I) [regular files]
DUMP: Reading file names from standard input
user1
user3/jdoe
user5/rroe/src
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
u	Record the backup in the /etc/dumpdates file.
f	A tape device is supplied in the command line.
n	A name for the backup is supplied.
l	The names of individual files and directories to be backed up are entered interactively.
rst0a	The tape drive rewinds.
user.1.3.5	This is the name of the backup.
/vol/vol1/home	This is the directory that contains the files to be backed up.

Examples of the dump command

Backup data without ACLs

Backup not including ACLs

The following command performs a level-0 backup of the /vol/vol1 volume. The A option means that the backup does not include any Windows NT ACL information. Use this option if the files in the volume are for a UNIX-only environment.

```
dump 0uAf rst0a /vol/vol1
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
u	Record the backup in the /etc/dumpdates file.
A	Do not back up Windows NT ACLs.
f	A tape device is supplied in the command line.
rst0a	The tape drive rewinds.
/vol/vol1	This is the volume to be backed up.

Examples of the dump command

Backup specifying a blocking factor

Blocking factor specification

The following command performs a level-0 backup of the /vol/vol1 volume. This command writes 32 KB of data at a time, enabling you to restore the data from systems that limit each write to 32 KB.

```
dump 0ufb rst0a 32 /vol/vol1
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
u	Record the backup in the /etc/dumpdates file.
f	A tape device is supplied in the command line.
b	A blocking factor is specified.
rst0a	The tape drive rewinds.
32	The blocking factor is 32, so write 32 KB of data at a time.
/vol/vol1	This is the volume to be backed up.

Examples of the dump command

Backup excluding specified files

**File exclusion
backup**

The following command performs a level-0 backup of the /vol/vol1 volume, but excludes the files that meet certain requirements:

```
dump 0ufX rst0a tmp,*.o,core*,*backup*, /vol/vol1
```

The elements of this command line are described in the following table.

Element	Meaning
0	Do a full backup.
u	Record the backup in the /etc/dumpdates file.
f	A tape device is supplied in the command line.
X	An exclude list is specified.
rst0a	The tape drive rewinds.
tmp,*.o,core*,*backup*	The exclude list specifies files as follows: ◆ tmp specifies that the file name is tmp. ◆ *.o specifies that the file name ends in .o (for example, program.o). ◆ core* specifies that the file name begins with the core string (for example, core.small). ◆ *backup* specifies that the file name contains the backup string (for example, spec.backup.1).
/vol/vol1	This is the volume to be backed up.

Verifying the contents of a dump command backup

Reason to verify a backup

You might want to verify a backup initiated by the `dump` command to ensure that all the files you wanted to back up are on the tape.

Verifying a backup

To verify a backup, complete the following steps.

Step	Action
1	From your client, preserve output to the console by using a utility such as <code>script</code> .
2	List all the files in a backup by entering the following command: <code>restore tf rst0a</code>
3	Compare the list to what you intended to back up.
4	For more detailed verification, use the <code>N</code> option of the <code>restore</code> command, as described in “ Specifying a test restore ” on page 130.

Checking the status of a dump command backup

How to find out whether a backup needs to be restarted

To find out whether a backup initiated by the `dump` command is proceeding as expected or has terminated, you can run the `backup status` command. The command provides output in the following format.

ID	State	Type	Device	Start Date	Level	Path
1	ACTIVE	dump	nrst0a	Nov 28 00:22	0	/vol/vol0
2	ACTIVE	dump	nrst0a	Nov 28 00:22	3	/vol/vol1
4	ACTIVE	NDMP	urst1a	Nov 28 00:22	1	/vol/vol0
6	RESTARTABLE	dump		Nov 27 00:22	3	/vol/vol1

The elements of this output (the “dump table”) are described in the following table.

Element	Meaning
ID	The unique ID assigned to the dump and the index in the software’s internal dump table. As soon as a dump completes, its ID number is deallocated and returned to the pool of available slots. The total number of entries in the dump table is limited to 32.
State	The state of the dump: ACTIVE or RESTARTABLE
Type	The type of dump: standard format (dump) or NDMP protocol (NDMP)
Device	The current device to which the dump is writing
Start Date	Date on which the backup began
Level	The level of the dump (0 through 9)
Path	The dump path

How to get details about a specific backup

To get more detailed information about a specific backup initiated by the `dump` command, supply the dump ID at the end of the `backup status` command, for example:

```
filer> backup status 2
State:      ACTIVE                                Type:      dump
Path:       /vol/vol0/src                          Level:     0
Options:    b=63, u
Devices:    rst1a,rst2a,rst3a
Completed:  3 tape files
Last Update: Mon Nov 26 00:14:35 2001
```

The `backup status` output for an individual dump ID supplies the following additional information.

Element	Meaning
Options	All the options specified for the backup and their respective arguments
Completed	The number of tape files that have already been copied
Last Update	The time and date of the last completed update

Restarting a dump command backup

Why you might need to restart a backup

You might need to restart a backup initiated by the `dump` command. The `dump` command sometimes does not finish because of such internal or external errors as tape write errors, power outages, accidental user interruptions, or internal inconsistency on the storage system.

You might also choose to interrupt and restart a backup to avoid periods of heavy traffic on the storage system or to avoid competition for other limited resources on the storage system, such as a tape drive. A long backup can be interrupted and restarted later if a more urgent restore (or backup) requires the same tape drive. Restartable backups persist across reboots.

What qualifies as a restartable backup

A failed backup to tape can be restarted if it has completed one tape that has ended in phase IV (file data and ACL data).

How to find out whether a backup needs to be restarted

To find out whether a backup is proceeding as expected or has terminated, you can run the `backup status` command. For more information, see “[Checking the status of a dump command backup](#)” on page 94.

How to restart a dump process

To restart a dump process that has been shown to be restartable, complete the following step.

Step	Action
1	From the storage system, enter the following command: <pre>dump R[f comma-separated_device-list] {path ID}</pre> <i>f</i> is an option that enables you to supply a device list. <i>comma-separated_device-list</i> lets you direct the dump stream to output devices other than those originally designated in the failed dump. A restarted dump process uses this device list the same way a regular dump would. Any device list that is valid to a regular dump will be valid in this case. If a device list is not specified, the command defaults to the remainder of the devices listed but not yet consumed by the failed dump. For example, suppose the following device list was supplied to the previous dump, which failed while writing to <code>rst2a</code>: <pre>rst0a,rst1a,rst2a,rst3a,rst4a</pre> The command would use <code>rst3a,rst4a</code> as the default. However, if the original device list contained any non-rewinding (<code>nrst</code>) devices or any devices not supported by NetApp, users are required to supply a new device list at restart of the dump. <i>path</i> is the path that is listed in the dump table (the output of the <code>backup status</code> command). If there are multiple entries, that is, entries with exactly the same path, the command prompts you to use the ID to restart the dump. <i>ID</i> is the unique ID displayed by the <code>backup status</code> command. You can use either <i>path</i> or <i>ID</i> in most cases. Result: The command starts rewriting the dump stream from the beginning of the tape file in which the previous dump was interrupted.

Deleting restartable dump command backups

Deleting restartable backups

To delete a restartable backup, complete the following step.

Step	Action
1	Enter the following command: backup terminate <i>ID</i> <i>ID</i> is the unique ID displayed in the dump table displayed by the <code>backup status</code> command. Result: The command deletes the specified restartable backup.

How the dump command deletes restartable backups automatically

To prevent restartable backups from accumulating on a storage system and taking up unreasonable amounts of disk space, the `dump` command automatically checks the the snap reserve every 10 minutes. If the snap reserve is over 100%, the oldest restartable backups are deleted until snap reserve usage drops below 100% or there are no more restartable backups to delete.

About this chapter

Data restoration from tape is the process of copying data files back to the storage system from backup tapes that were created using the storage system `dump` command.

This chapter explains how to use the storage system `restore` command for data recovery.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“About data restoration”](#) on page 100
- ◆ [“What the restore command recovers”](#) on page 101
- ◆ [“How the restore command works”](#) on page 103
- ◆ [“Preparing to use the restore command”](#) on page 104
- ◆ [“Using the restore command”](#) on page 108
- ◆ [“Restoring incremental backups”](#) on page 131
- ◆ [“Restoring the entire storage system”](#) on page 132
- ◆ [“Restore examples”](#) on page 133

About data restoration

When to restore data from tape

You restore data from tape under the following circumstances:

- ◆ Files were deleted from disk but backed up to tape.
For example, if you inadvertently delete a file and want to restore the file, you can recover the file from tape.
- ◆ Files are corrupted. If some files are corrupted, you can restore the subtree containing the files.
- ◆ No disk slots are available for expansion. If the storage system runs out of storage space, you can do the following tasks:
 - ❖ Back up the entire storage system.
 - ❖ Replace the current disks with disks of greater capacity.
 - ❖ Set up volumes on new disks.
 - ❖ Restore the storage system from tapes.
- ◆ The entire storage system is damaged and unusable. Call Network Appliance Technical Support to determine whether you can repair the storage system and restore data from tape.

When not to restore data from tape

You do not need to restore a deleted file from tape if you can recover the deleted file from a Snapshot copy on the storage system. If the file is in a Snapshot copy, copying the file to the active file system is faster than recovering the file from tape.

Note

If you attempt an incremental restore to a storage system running Data ONTAP 6.2 or later from a storage system running a version earlier than 6.2, the restore will fail because of a formatting code change between the two code releases. You need to run the full backup again after you have upgraded to 6.2.

What the restore command recovers

What data can be recovered

The `restore` command enables you to recover all the information that you backed up with the `dump` command so that it is the same as when you did the backup.

The `restore` command can recover the following data:

- ◆ Contents of files and directories
- ◆ UNIX file permissions
- ◆ Windows NT access control lists (ACLs)

Note

If you restore a file that has only UNIX file permissions into an NT file system (NTFS) qtree or volume, the file has no Windows NT ACLs. The storage system uses only the UNIX file permissions on this file until you create a Windows NT ACL on it.

- ◆ Qtree information

Note

Qtree information is used only if a qtree is restored to the root of a volume. If a qtree is restored to a lower directory, such as `/vol/vol0/subdir/lowerdir`, it ceases to be a qtree and qtree information is not used.

- ◆ All other file and directory attributes
- ◆ Windows NT streams information

Compatibility

You can restore data from previous versions of Data ONTAP using the `restore` command. You can perform an incremental restore to a storage system running data ONTAP 6.2 or later using data backed up from a storage system running a version of Data ONTAP earlier than 6.2 by doing one of the following:

- ◆ Perform a level 0 restore and incremental restores before you upgrade to Data ONTAP 6.2 or later.
- ◆ Perform a level 0 restore and incremental restores after you upgrade to Data ONTAP 6.2 or later.

Performing a level 0 restore, upgrading Data ONTAP software, then performing incremental restores will not restore the incremental backups because they are in a different format from the level 0 restore. If you do this, you have to repeat the level 0 restore before you can restore incremental backups.

How the restore command works

The restore process The process in the following table shows how the `restore` command works.

Stage	Action
1	Data ONTAP catalogs the files that need to be extracted from the tape.
2	Data ONTAP creates directories and empty files.
3	Data ONTAP reads a file from tape, writes it to disk, and sets the permissions—including ACLs—on it.
4	Data ONTAP repeats steps 2 and 3 until all the specified files are copied from the tape.

Preparing to use the restore command

General preparations

Before restoring data with the `restore` command, you should make sure that you have the required materials, decide where to enter the `restore` command, and prepare the destination for the restore.

For detailed information

The following topics describe in detail how to prepare to restore files using the `restore` command:

- ◆ [“Getting required materials”](#) on page 105
- ◆ [“Deciding where to enter the restore command”](#) on page 106
- ◆ [“Preparing the destination”](#) on page 107

Getting required materials

Required information

Before restoring data with the `restore` command, you must have the following information:

- ◆ The level of the restore
- ◆ The tape device you used for each tape file in the backup that you are restoring
- ◆ The path into which you are restoring the material
- ◆ The blocking factor used in the backup

Required tape drives and tapes

You must meet these prerequisites for the `restore` command to run successfully:

- ◆ If you are doing an incremental restore, you need all the tapes in the backup chain.
- ◆ You need a tape drive that is available and compatible with the tape to be restored from.

Required space on destination storage system

You need about 100 MB more space on the destination storage system than the amount of data to be restored.

Caution

The `restore` command terminates if it runs out of space.

Deciding where to enter the restore command

Where to enter the command

You can enter the `restore` command through a Remote Shell connection, such as `rsh`, or on the console.

Benefits of entering the restore command through a Remote Shell

Entering the `restore` command through a Remote Shell connection gives you these benefits:

- ◆ When the `restore` command is in progress, you can still use the console to manage the storage system.
- ◆ You can start multiple `restore` commands through a Remote Shell connection if other tape drives are available.
- ◆ It is less likely that someone will inadvertently terminate the `restore` command, especially if it is run in the background from a UNIX system. If you enter a `restore` command on the console, it could be terminated by `Ctrl-C` entered on a host connected to the storage system using `Telnet`.

Benefit of entering the restore command on the console

The benefit of entering the `restore` command on the console is that you can read and respond to screen messages displayed by the command. For example, the command might prompt you for another tape to complete the recovery.

Preparing the destination

About preparing the destination

If you are restoring the backup to its original path, you do not need to prepare the target volume, qtree, or subtree. If you are restoring the backup to a different destination, you must prepare the location.

Preparing the destination

If you are restoring to a location other than the original location, complete the following step.

Step	Action	
1	If you are restoring...	Then...
	A volume	Create a new volume.
	A qtree or a directory	Rename or move files that are likely to have the same names as files you are restoring. Caution ————— If a restored file has the same name as an existing file, the existing file is overwritten by the restored file. Directories are not overwritten. —————

Using the restore command

About the restore command

The `restore` command enables you to specify four types of restores and to specify several items in a restore. This section describes the principles of the `restore` command, then the kinds of restores you can perform and the additional items that you can specify.

For detailed information

The following topics describe how to use and construct a `restore` command:

- ◆ [“Restore command principles”](#) on page 109
- ◆ [“Specifying an individual file restore”](#) on page 113
- ◆ [“Specifying a full restore”](#) on page 114
- ◆ [“Specifying a table-of-contents restore”](#) on page 115
- ◆ [“Specifying a resumed restore”](#) on page 117
- ◆ [“Specifying tape devices”](#) on page 119
- ◆ [“Specifying a single tape file on a multiframe tape”](#) on page 121
- ◆ [“Specifying the restore destination”](#) on page 122
- ◆ [“Specifying the blocking factor”](#) on page 123
- ◆ [“Specifying detailed status output”](#) on page 124
- ◆ [“Specifying ignoring inode limitations”](#) on page 125
- ◆ [“Specifying automatic confirmations”](#) on page 127
- ◆ [“Specifying no Windows NT ACLs”](#) on page 128
- ◆ [“Specifying no qtree information”](#) on page 129
- ◆ [“Specifying a test restore”](#) on page 130

Restore command principles

Syntax of the restore command

The `restore` command syntax is as follows:

```
restore options [arguments] [files ...]
```

options can be one restore type with modifiers.


Rules for using the restore command

Follow these rules when you enter the `restore` command:

- ◆ Specify one and only one restore type.
- ◆ Specify multiple options without intervening spaces.
- ◆ Enter the arguments for each option in the order that you specify the options. Separate each argument from the next with a space.
- ◆ If the destination for each file is the same as the location from which it was backed up, you do not need to explicitly specify a destination.

Example of a simple restore command

The following command illustrates the `restore` command syntax. It restores all the files in the dump tape to the same destinations from which they were backed up.



Restore types

A restore type specifies the type of restore you are performing. For a restore from tape, you must specify one and only one restore type. The following table summarizes the restore types.

Restore type	Meaning	Option	For more information, see...
Restart	Restarts data recovery after an interruption.	R	“Specifying a resumed restore” on page 117
Qtree table of contents	Lists qtree names and qtree information in a restore.	T	“Specifying a table-of-contents restore” on page 115
Full	Rebuilds the file system or subtree. If you are applying incrementals, you must specify this option.	r	“Specifying a full restore” on page 114
File table of contents	Lists file names in a restore.	t	“Specifying a table-of-contents restore” on page 115
File	Extracts an individual file or subtree from the backup.	x	“Specifying an individual file restore” on page 113

Modifiers

Modifiers specify optional actions. The following table summarizes the modifiers.

Option	Meaning	For more information, see...
A	Specifies not to restore Windows NT ACLs.	“Specifying no Windows NT ACLs” on page 128

Option	Meaning	For more information, see...
D	Specifies the directory into which the files are restored. Argument: The directory into which you are restoring files. Without an argument, the files are restored to the directory from which they were backed up.	“Specifying the restore destination” on page 122
F	Forces <code>restore</code> to continue regardless of inode limitations.	“Specifying ignoring inode limitations” on page 125
N	Reads backup tapes without writing to the storage system.	“Specifying a test restore” on page 130
Q	Ignores qtree information.	“Specifying no qtree information” on page 129
b	Specifies the blocking factor. Argument: The blocking factor that you used in the backup that you are restoring	“Specifying the blocking factor” on page 123
f	Specifies the tape device for each tape file. Argument: The name of one or more tape devices, separated by commas	“Specifying tape devices” on page 119
s	Specifies the relative position of a tape file if multiple tape files exist on a tape. File numbering starts at 1 from the current tape position. Argument: The tape file number	“Specifying a single tape file on a multfile tape” on page 121

Option	Meaning	For more information, see...
v	Specifies that <code>restore</code> displays the inode number of each file restored.	“Specifying detailed status output” on page 124
y	Specifies that <code>restore</code> not ask whether it should terminate when it gets an error.	“Specifying automatic confirmations” on page 127

Executing a restore command

Follow these steps to execute a `restore` command.

Step	Action
1	Place the tape containing the first tape file of the backup in the tape drive that you specify.
2	Enter the <code>restore</code> command.
3	If prompted, insert the next tape of the backup that you are restoring into the appropriate tape drive.
4	Repeat Step 3 until the restore is complete.

Using the restore command

Specifying an individual file restore

About restoring an individual file

You can restore one or more directories or files from a backup.

Specifying individual file restores

To restore an individual directory or file, complete the following steps.

Step	Action
1	Use the <code>x</code> option in the <code>restore</code> command line.
2	At the end of the command line, include the path names relative to the dump path of the files or directories that you want to restore. Separate path names with a space. Caution————— If you do not have a path in the command line, the <code>restore</code> command restores everything on the tape.

Example: The following command restores the `/src` directory and puts it in the location from which it was backed up:

`restore xf rst0a /src`

↑
x restores a
specified file.

↑
The specified file is
the `/src` directory.

Specifying a full restore

About full restores A full restore rebuilds the file system, qtree, or subtree that was in the backup that a tape file contains.

Specifying a full restore To specify a full restore, complete the following step.

Step	Action
1	Use the <code>r</code> option in the <code>restore</code> command line.

Example: The following command restores the `/src` directory and puts it in the same directory that it was backed up from:

```
restore rf rst0a
```


`r` performs a full restore.

Specifying a table-of-contents restore

About displaying a table of contents of a tape file

You can display a table of contents of the files or qtrees in a tape file. This is useful in determining what files or qtrees are on a tape and their locations. For qtrees, the restore lists qtree properties.

A table-of-contents restore takes much less time than a full restore because only the list of files in the backup is read. However, it uses a lot of CPU time because of the extensive output produced.

Use Remote Shell for a table-of-contents restore

In general, you should run a table-of-contents restore from a Remote Shell connection because an enormous output is generated. Usually, you can control the output more easily when it is sent to a client console rather than to the storage system console. Client consoles are also more flexible and enable you to save output.

Furthermore, you rarely need to change tapes with a table-of-contents restore. The command needs to read only the directory information from the tape and none of the files or qtrees. Because directory information tends to constitute a small part of a backup, it is almost always located on one tape. Also, table-of-contents restores work with multiple tape files specified on the command line.

Types of table-of-contents restores

You can specify two kinds of tables of contents: file and qtree. You cannot combine the two types in a single command.

Type	Meaning	Option
File	Lists all the file names in a backup. If you specify path names, only the files in the path names are listed.	t
Qtree	Lists qtrees and their settings for security style and Windows NT oplocks for all qtrees. If you specify qtree names, the information for only those qtrees is listed if they are in the backup.	T

Specifying table-of-contents restores

To specify a table-of-contents restore, complete the following step.

Step	Action
1	Use the T or t option in the <code>restore</code> command line, with files as an argument. If there is no argument, the entire content of a backup is listed.

Example: The following command lists all files in a backup:

```
restore tf rst0a
```

↑
t lists all files.

Example: The following command lists all qtree names in a backup:

```
restore Tf rst0a
```

↑
T lists all qtree names.

Specifying a resumed restore

About resuming a restore

If an entire tape file restore is stopped, you can resume the restore and avoid restoring again what has already been restored.

Resume restrictions

Resuming a restore has the following restrictions:

- ◆ You can resume only restores that you started with the `r` or `R` options.
- ◆ You can resume a `restore` command only if the backup consists of multiple tape files.
- ◆ You can resume a `restore` command only if the command is for a full restore. If the `restore` command is for extracting an individual file or subtree from a backup (that is, if you use the `x` option), or for a table-of-contents restore, you cannot resume the restore.
- ◆ You can resume a restore only if you received a message similar to the following during the restore:

```
RESTORE: Fri Aug 31 22:22:35 2001: Writing data to files.
```

Resuming restores

To resume a restore, complete the following steps.

Step	Action
1	In the <code>restore</code> command line, use the <code>R</code> option instead of the <code>r</code> option as the first option. It does not take an argument.
2	Enter the rest of the same <code>restore</code> command that was interrupted, but include only the tape files that were not restored.
3	Follow the prompts.

Example: The following command resumes a restore.

```
restore Rf rst0a
```



R resumes a restore.

Using the restore command

Specifying tape devices

About backup tape devices

When you performed a backup, you specified one or more tape devices. The files written by these devices can be on one or more tapes.

Compression types must be the same

You must use the same compression type to restore a backup as you did to perform the backup; however, you can use a different rewind type and device number. For example, you can use `rst1a` and tape drive 1 to restore a backup done on `nrst0a`, provided that the two tape drives use the same kind of tape.

Specifying the tape devices

To specify the tape devices for restores, complete the following steps.

Step	Action
1	Use the <code>f</code> option in the <code>restore</code> command line.
2	List the tape devices as an argument to the <code>f</code> option in the same order that you used in the backup. Separate multiple tape devices with a comma. The <code>restore</code> command restores from tape files consecutively, using the tape devices in the order that they appear in the command line. Note_____ If you do not specify at least one tape device, the <code>restore</code> command terminates. _____

Example: The following command specifies the `rst0a` device for a backup:


The tape device is `rst0a`.

↓

```
restore rf rst0a
```

↑

`f` indicates that a tape device is specified.



Specifying a single tape file on a multifile tape

About multifile tapes


You can have more than one tape file on a tape. Tape files do not have names. You can restore a single tape file on a tape that contains more than one tape file by moving the tape to the beginning of the file.

Specifying the file

To specify a single tape file to be restored from a multifile tape, complete the following steps.

Step	Action
1	Use the <code>f</code> option in the <code>restore</code> command line.
2	Use the same tape compression type as an argument to the <code>f</code> option that you used in the backup.
3	Use the <code>s</code> option in the <code>restore</code> command line to select the appropriate backup.
4	Include the relative position of the tape file that you are restoring in the command line as an argument to the <code>s</code> option. Note _____ Count the relative position from the current tape position. It might be best to rewind the tape and start from the beginning of the tape.

Example: The tape has been rewound, and the following command restores the third tape file from the beginning of the tape, then rewinds the tape:



Specifying the restore destination

When to specify the restore destination

The destination acts as the root of the backup that you are restoring. You specify a different restore destination if you are moving items. For example, if you created a backup and then installed multiple volumes on the storage system, when you perform a restore, you might specify a different volume or directory.

If you do not specify a restore destination, the files are restored into the locations from which they were backed up.

Note


Network Appliance recommends that you specify a restore destination even if you are restoring to the same destination that you backed up from. This ensures the files are restored where you want them to go and are traceable to that location.

Specifying the destination

To specify the restore destination, complete the following steps.

Step	Action
1	Use the <code>D</code> option in the <code>restore</code> command line.
2	Include the absolute path name of the restore destination as an argument to the <code>D</code> option.

Example: The following command restores a backup and puts it in the `/vol/destination` volume:



Specifying the blocking factor

About the blocking factor


The blocking factor specifies the number of tape blocks that are transferred in each write operation. A tape block is 1 kilobyte of data. When you restore, you must use the same blocking factor that you used for the backup. The default blocking factor is 63.

Specifying the factor

To specify the blocking factor, complete the following steps.

Step	Action
1	Use the <code>b</code> option in the <code>restore</code> command line.
2	Include the blocking factor as an argument to that option.

Example: The following command restores a backup and puts it in the `/vol/destination` volume:



Using the restore command

Specifying detailed status output

About detailed status output

You can get information about the progress of a restore on a file-by-file basis. If you have a restore problem, this output can be useful for your own diagnostics, as well as for Network Appliance Technical Support.

Because of the volume of information that needs to be processed by a console, getting detailed output can slow down a restore considerably.

Specifying detailed output

To get status information about each file recovered, complete the following step.

Step	Action
1	Use the <code>v</code> option in the <code>restore</code> command line. Note _____ This option does not take an argument.

Example: The following command restores a backup and produces status information about each file recovered:

`restore rfv rst0a`



`v` produces status information about each file recovered.

Using the restore command

Specifying ignoring inode limitations

About inode limitations

Inodes are data structures that contain information about files. The number of files, and, therefore, the number of inodes per volume, is determined by the `maxfiles` command. For information about setting the maximum number of files per volume and displaying inode information, see the *Storage Management Guide*.

Reasons to ignore inode limitations

The `restore` command assumes that the files being restored are added to the number of files on the storage system, and, therefore, that the inodes are added to the storage system. When the total of inodes in the restore and on a storage system exceeds the number of inodes that are allowed on a storage system, the restore is terminated.

However, if a restore updates an existing file, the inode count remains the same. Therefore, if you are sure that the restore consists mostly of files to be updated rather than new files, you can safely tell the storage system to ignore the calculations of the `restore` command.

Note_____

If a restore causes the number of inodes to exceed the maximum allowed, the restore is terminated.

Ignoring inode limitations

To specify a restore that ignores inode limitations, complete the following step.

Step	Action
1	Use the <code>F</code> option in the <code>restore</code> command line.
	Note _____
	This option does not take an argument.

Example: The following command restores a backup and ignores inode limitations:

```
restore rfF rst0a
```



F specifies to ignore
inode limitations.

Using the restore command

Specifying automatic confirmations

About automatic confirmations

Automatic confirmation means answering *yes* to any restore questions. You usually use this mode on restores run using a Remote Shell connection. A Remote Shell connection does not let you interact with the `restore` command; therefore, if the `restore` command requires user input and is run using a Remote Shell connection, it usually terminates. Specifying confirmation mode enables such restores to be completed in most cases. Even with the `y` option, however, the `restore` command fails if it encounters hard media errors or unclean drives.

Caution_____

This option is not recommended for critical restores because it can cause silent failure.

Specifying automatic confirmations

To specify automatic confirmations, complete the following step.

Step	Action
1	Use the <code>y</code> option in the <code>restore</code> command line. Note _____
	This option does not take an argument. _____

Example: The following command restores a backup with automatic confirmations:

```
restore rfy rst0a
```

 ↑
 y specifies automatic confirmations.

Specifying no Windows NT ACLs

About omitting Windows NT ACLs

You can omit Windows NT ACLs from a restore. This provides a slight performance enhancement. You omit Windows NT ACLs in two situations:

- ◆ You plan to restore to a UNIX-only environment or to a version of Data ONTAP earlier than 5.1.
- ◆ Your backup has no files or directories that contain Windows NT ACLs.

Omitting Windows NT ACLs

To omit Windows NT ACLs from a restore, complete the following step.

Step	Action
1	Include the A option in the <code>restore</code> command line. Note_____ _____ _____ This option does not take an argument. _____

Example: The following command restores a backup but does not restore Windows NT ACLs:

`restore rfA rst0a`

↑
A specifies not to restore Windows NT ACLs.

Using the restore command

Specifying no qtree information

About omitting
qtree information

You can omit qtree information from a restore. Qtrees are then restored as ordinary directories. You specify no qtree information in a restore if you want to restore qtrees as regular directories.

Omitting qtree
information

To omit qtree information from a restore, complete the following step.

Step	Action
1	Include the Q option in the <code>restore</code> command line. Note_____ This option does not take an argument. _____

Example: The following command restores a backup but does not restore qtree information:

```
restore rfQ rst0a
```


Q specifies not to restore
qtree information.

Specifying a test restore

About test restores You can test a restore by performing a restore that reads the tape but does not write to the storage system.

Note _____
Because a test restore is dependent on the speed of reading from tape, it takes almost as much time as an actual restore.

When to test a restore You would do a test restore in the following situations:

- ◆ To verify a backup tape that is old and might have deteriorated
- ◆ To verify that the set of tapes you have is complete
- ◆ To verify a backup tape that you believe was not written properly
- ◆ To quickly make sure that a block size works if it the block size is unknown

Specifying a test restore To specify a test restore, complete the following step.

Step	Action
1	Include the <code>N</code> option in the <code>restore</code> command line. Note _____ This option does not take an argument.

Example: The following command performs a test restore of a backup:

```
restore rfN rst0a
```

↑

N specifies a test restore.

Restoring incremental backups

Required tapes

Incremental restores build on each other the way incremental backups build on the initial level-0 backup. Therefore, to restore an incremental backup, you need all the backup tapes from the level-0 backup through the last backup that you want to restore.

For additional information about incremental backups and backup chains, see [“Specifying the backup level”](#) on page 63 and [“Understanding increment chains”](#) on page 64.

Restoring an incremental backup

To restore an incremental backup, complete the following steps.

Step	Action
1	Restore the level-0 backup.
2	Follow the prompts. You might be asked to remove or insert tapes.
3	Restore each incremental backup in the increment chain that you want to restore, starting with the lowest-level backup and going to the last backup that you want to restore. Note _____ During an incremental restore operation, a temporary directory labeled <code>.restore_do_not_touch_xxxxxxx</code> will appear in the active file system. Do not edit or delete this directory. The system will delete this directory after the current incremental restore operation is completed.
4	From the client, remove the <code>restore_symboltable</code> file in the root of the destination directory.

Restoring the entire storage system

About restoring the entire storage system

You can restore a storage system only to disks that you have initialized.

How you restore an entire storage system depends on whether you backed up each volume in one backup or as individual directories and qtrees.

If there is one backup for each volume

To restore an entire storage system if you used one `dump` command to back up each volume, complete the following step.

Step	Action
1	Restore each backup.

If each volume was backed up as subtrees or qtrees

To restore an entire storage system if you used separate `dump` commands to back up files, directories, and qtrees that make up each volume, complete the following steps.

Step	Action
1	Create the desired volumes.
2	Restore each backup to the appropriate volume.

Restore examples

About the examples Examples of how to perform the following tasks can help you use the `restore` command.

For detailed information You can use the following examples of the `restore` command to understand the `restore` command and as boilerplate for performing restores:

- ◆ “[Restoring an entire backup](#)” on page 134
- ◆ “[Restoring individual files or directories](#)” on page 136
- ◆ “[Restoring using a remote tape drive](#)” on page 137
- ◆ “[Multiple tape restores](#)” on page 138

About prompts In the `restore` command prompts, the following terms are used interchangeably to mean tape file or tape cartridge:

- ◆ Tape name
- ◆ Tape volume
- ◆ Tape

Restoring an entire backup

Restoring a backup to its original location

Backup: `dump 0f rst0a /vol/vol1/`

Restore: `restore rfD rst0a /vol/vol1/`

This command restores the entire tape file to its original location, in this case /vol/vol1. The tape drive rewinds the tape.

The elements of this command line are described in the following table.

Element	Meaning
r	Do a full restore.
f	A tape device is supplied in the command line.
D	The target directory is specified.
rst0a	The restore is done using the rst0a tape device.

Restoring a backup to a different location

Backup: `dump 0f rst0a /vol/vol1/users/tom`

Restore: `restore rfD rst0a /vol/vol1/home/tom`

This command restores the /vol/vol1/users/tom directory to the /vol/vol1/home/tom directory. The tape drive rewinds the tape.

The elements of this command line are described in the following table.

Element	Meaning
r	Do a full restore.
f	A tape device is supplied in the command line.
D	The target directory is specified.
rst0a	The restore is done using the rst0a tape device.

Element	Meaning
/vol/vol1/home/tom	The restored data is put in the /vol/vol1/home/tom directory.

Restoring individual files or directories

About restoring a single file

The `restore` command can extract individual qtrees, subdirectories, or data files from a backup that contains many qtrees, ordinary directories, or data files.

Restoring individual files

Backup: `dump 0f rst0a /vol/vol1/users/tom`

Restore: `restore xFD rst0a /vol/vol1/u1/tom spec4 spec5`

This command extracts the `/vol/vol1/users/tom/spec4` and `/vol/vol1/users/tom/spec5` files from the `/vol/vol1/users/tom` directory and places them in the `/vol/vol1/u1/tom` directory. The tape drive rewinds the tape.

The elements of this command line are described in the following table.

Element	Meaning
x	Restore a single file or directory.
f	A tape device is supplied in the command line.
D	Restore to a different location from the original.
rst0a	The restore is done using the <code>rst0a</code> tape device.
/vol/vol1/u1/tom	The restored data is put in the <code>/vol/vol1/u1/tom</code> directory.
spec4 spec5	The <code>/vol/vol1/users/tom/spec4</code> and <code>/vol/vol1/users/tom/spec5</code> files are extracted.

Restoring using a remote tape drive

Storage system restore using a tape drive attached to a remote storage system

Backup: `dump 0f netappl:rst0a /vol/vol1`

Restore: `restore rf netappl:rst0a`

This command performs a restore from a tape drive attached to a remote storage system named `netappl`. The tape drive then rewinds the tape.

The elements of this command line are described in the following table.

Element	Meaning
<code>r</code>	Do a full restore.
<code>f</code>	A tape device is supplied in the command line.
<code>netappl</code>	This is the storage system name.
<code>rst0a</code>	The restore is done using the <code>rst0a</code> tape device.

Storage system restore using a tape drive attached to a Solaris system

Backup: `dump 0f ritchie:/dev/rmt/0 /vol/vol1`

Restore: `restore rf ritchie:/dev/rmt/0`

This command performs a restore from a tape drive on a Solaris system.

The elements of this command line are described in the following table.

Element	Meaning
<code>r</code>	Do a full restore.
<code>f</code>	A tape device is supplied in the command line.
<code>ritchie</code>	This is the name of the Solaris machine to which the tape drive is connected.
<code>/dev/rmt/0</code>	This is the name of the tape device.

Restore examples

Multiple tape restores

Examples of multiple tape restores

This section provides examples of the following kinds of multiple tape restores:

- ◆ Multiple tapes on one single-tape drive
- ◆ Multiple tapes on two single-tape drives
- ◆ Multiple tapes on a tape stacker

Multiple tapes on one single-tape drive

Backup: `dump 0f rst0a /vol/vol1`

Restore: `restore rf rst0a`

This command restores the /vol/vol1 volume from the two tapes it took to back it up. You are prompted for the next tape when the first tape is restored.

The elements of this command line are described in the following table.

Element	Meaning
r	Do a full restore.
f	A tape device is supplied in the command line.
rst0a	The restore is done using the rst0a tape device; the restore command prompts for the second tape.

Multiple tapes on two single-tape tape drives

Backup: `dump 0f rst0a,rst1a /vol/vol1`

Restore: `restore rf rst0a,rst1a`

The first tape is in tape drive 0 and the second tape is in tape drive 1. This command restores the /vol/vol1 volume from the two tapes it took to back it up. It uses the tape in the second tape drive when the first tape is restored.

The elements of this command line are described in the following table.

Element	Meaning
r	Do a full restore.
f	A tape device is supplied in the command line.
rst0a	The restore is done using the <code>rst0a</code> tape device for the first tape
rst1a	The restore is done using the <code>rst1a</code> tape device for the second tape.

Multiple tapes on a tape stacker

Backup: `dump 0f urst0a,urst0a /vol/vol1`

Restore: `restore rf urst0a,urst0a /vol/vol1`

This command restores the `/vol/vol1` volume from the two tapes it took to back it up. It unloads the first tape and loads the second tape.

The elements of this command line are described in the following table.

Element	Meaning
r	Do a full backup.
f	A tape device is supplied in the command line.
urst0a,urst0a	The tape drive unloads and loads each tape.

About this chapter

This chapter describes how you can use Network Data Management Protocol (NDMP) services on your storage system to enable network-based backup and recovery using a number of NDMP-enabled commercial backup applications. It also describes how you can monitor NDMP services running on the storage system and how you can use `ndmpcopy` to migrate data efficiently within or between storage systems.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Learning about NDMP services on the storage system”](#) on page 142
- ◆ [“Using tape devices with NDMP”](#) on page 146
- ◆ [“Preparing a storage system for basic NDMP backup application management”](#) on page 148
- ◆ [“Enabling or disabling NDMP service”](#) on page 150
- ◆ [“Managing NDMP security features”](#) on page 151
- ◆ [“Specifying the NDMP version”](#) on page 157
- ◆ [“Specifying a preferred network interface”](#) on page 159
- ◆ [“Enabling or disabling enhanced DAR functionality”](#) on page 161
- ◆ [“Specifying environment variables for backup applications”](#) on page 162
- ◆ [“Displaying NDMP session status”](#) on page 171
- ◆ [“Displaying NDMP debug messages”](#) on page 174
- ◆ [“Terminating an NDMP session”](#) on page 177
- ◆ [“Copying data with `ndmpcopy`”](#) on page 178

Learning about NDMP services on the storage system

What NDMP is

The Network Data Management Protocol (NDMP) is a standardized protocol for controlling backup, recovery, and other transfers of data between primary and secondary storage devices such as storage systems and tape libraries.

By enabling NDMP protocol support on a storage system, you enable that storage system to carry out communications with NDMP-enabled commercial network-attached backup applications (also called Data Management Applications or DMAs), data servers, and tape servers participating in backup or recovery operations.

NDMP also provides low-level control of tape devices and medium changers.

Note

You can use the following `snapmirror` commands when using NDMP:

- ◆ `abort`
- ◆ `help`
- ◆ `retrieve`
- ◆ `status`
- ◆ `store`
- ◆ `use`

See the *Online Backup and Recovery Guide* for information on the `snapmirror` command.

Advantages of NDMP

Data ONTAP NDMP backup and recovery operations use the same dump and restore services described in the chapters about data backup to tape and data restoration from tape. However, accessing these data protection services through backup applications that support NDMP offers a number of advantages.

- ◆ NDMP backup applications provide sophisticated scheduling of data protection operations across multiple storage systems. They also provide media management and tape inventory management services to eliminate or minimize manual tape handling during data protection operations.
- ◆ NDMP backup applications support data cataloging services that simplify the process of locating specific recovery data. Direct Access Recovery (DAR) optimizes the access of specific data from large backup tape sets.

- ◆ NDMP supports multiple topology configurations, allowing efficient sharing of secondary storage (tape library) resources through the use of three-way network data connections. NDMP backup applications typically provide user-friendly interfaces that simplify the management of data protection services.

Some common NDMP tape backup topologies

NDMP supports a number of topologies and configurations between backup applications and storage systems or other NDMP servers providing data (file systems) and tape services.

Storage system-to-local-tape: In the simplest configuration, a backup application backs up data from a storage system to a tape subsystem attached to the storage system. The NDMP control connection exists across the network boundary. The NDMP data connection that exists within the storage system between the data and tape services is called an NDMP local configuration.

Storage system-to-storage system-to-tape: A backup application can also back up data from a storage system to a tape library (a medium changer with one or more tape drives) attached to another storage system. In this case, the NDMP data connection between the data and tape services is provided by a TCP/IP network connection. This is referred to as an NDMP three-way storage system-to-storage system configuration.

Storage system-to-network attached tape library: NDMP-enabled tape libraries provide a variation of the three-way configuration. In this case, the tape library attaches directly to the TCP/IP network and communicates with the backup application and the storage system through an internal NDMP server.

Storage system-to-data-server-to-tape or data-server-to-storage system-to-tape: NDMP also supports storage system-to-data-server and data-server-to-storage system three-way configurations, although these variants are less widely deployed. Storage system-to-server allows storage system data to be backed up to a tape library attached to the backup application host or to another data server system. The server-to-storage system configuration allows server data to be backed up to a storage system-attached tape library.

NDMP security

Data ONTAP provides features for preventing or monitoring for unauthorized use of NDMP connections to your storage system. You can restrict the set of backup application hosts permitted to start NDMP sessions on a storage system. You can specify the authentication method to use (text or challenge) in order to allow NDMP requests. You can enable or disable monitoring of NDMP connection requests.

Data ONTAP also generates an NDMP-specific passwords for administrators who do not have root privilege on the target storage system.

Considerations when using NDMP

Take the following considerations into account when starting NDMP service on your storage system:

- ◆ Data ONTAP supports a maximum of 16 concurrent backup and/or restore operations. The maximum includes backups initiated by NDMP as well as by the storage system's `dump` and/or `restore` command.
- ◆ NDMP supports a maximum of 128 concurrent sessions on NearStore systems, and 40 on other systems.
- ◆ NDMP backup applications require specification of a target system password. To enable successful authentication by NDMP services on the storage system, you must use either the storage system's root password or, to authenticate a non-root user administrator, a system-generated NDMP-specific password.
- ◆ NDMP services optionally generate file history data at the request of NDMP backup applications. File history is used by backup applications to enable optimized recovery of selected subsets of data from a backup image. File history generation and processing is time-consuming and CPU-intensive for both the storage system and the backup application. If your data protection needs are limited to disaster recovery, where the entire backup image will be recovered, you may want to disable file history generation to reduce backup times. See your backup application documentation to determine if it is possible to disable NDMP file history generation.
- ◆ If enhanced Direct Access Restore (DAR) functionality is enabled, an offset map is generated at backup time. The processing required to generate the offset map is directly related to the number of files that need to be processed and can result in a significant impact to performance if a large number of small files are backed up. See [“Enabling or disabling enhanced DAR functionality”](#) on page 161.
- ◆ NDMP does not support Unicode format. If you use a third-party tool to back up a storage system using the Unicode directory format, the NDMP primary storage displays the file names in the NDMP file history database in two ways:
 - ❖ If possible, the storage system replaces the Unicode format with the NFS character set.
 - ❖ If the storage system cannot replace the Unicode format with the NFS character set, the NDMP client displays the file names as: `XXXXXXXX`, where `XXXXXXXX` is a character string.

- ◆ Because NDMP uses the `restore` command to recover files, you can use NDMP to restore files with correct file names regardless of how the client displays the file names in the NDMP file history database.

Using tape devices with NDMP

Tape devices and configurations you can use with the storage system

The storage system can read from or write to these devices when using NDMP:

- ◆ Standalone tape drives or tapes within a tape library attached to the storage system
- ◆ Tape drives or tape libraries attached to the workstation that runs the backup application
- ◆ Tape drives or tape libraries attached to a workstation or storage system on your network
- ◆ NDMP-enabled tape libraries attached to your network

Restriction for tape libraries

If you use NDMP to back up the storage system to attached tape libraries, be aware that if you want to use NDMP to manage your tape library, you must set the tape stacker autoloading setting to Off. If the autoloading setting is set to On, the storage system uses the tape stacker the same way it uses a stand-alone tape drive and does not allow medium changer operations to be controlled by the NDMP backup application.

Naming conventions for tape libraries

Historically, the following names were always used to refer to tape libraries:

- ◆ `mcn` or `/dev/mcn`
- ◆ `sptn` or `/dev/sptn`

In a specific tape library name, *n* is a number. For example, `mc0`, `spt0`, `/dev/mc0`, and `/dev/spt0` all refer to the same library.

Now, tape libraries can also be aliased to worldwide names (WWNs). For more information about tape aliasing, see “[Assigning tape aliases](#)” on page 16.

To see what tape libraries are recognized by your system, use the `sysconfig -m` command on the storage system console. To see what names are currently assigned to any libraries, use the `storage show mc` command on the storage system. Tape aliasing is also used to refer to tape drives, and you can see the aliases of tape drives using the `storage show tape` command.

Example

Below is an example of the output you might see from the `storage show mc` command:

```
filerA> storage show mc
Media Changer:      2.3
Description:        SPECTRA 10000
Serial Number:      7030290500
World Wide Name:    WWN[2:000:0090a5:00011c]
Alias Name(s):      mc0
Device State:       available (does not support reservations)
```

Preparing a storage system for basic NDMP backup application management

About NDMP preparation

To enable a storage system for basic management by a commercial NDMP backup application, you must enable the storage system’s NDMP support, specify the backup application’s configured NDMP version, host IP address, and authentication method. If an operator without root privileges to the storage system is using a backup application, a storage system-generated NDMP-specific password must be used by that user to carry out backup operations on that storage system.

To prepare a storage system for NDMP management, complete the following steps.

Step	Action
1	Enable NDMP. At the console command line of target storage system, enter: <code>ndmpd on</code> For details see “ Enabling or disabling NDMP service ” on page 150.
2	Specify the NPMD version to support on your storage system. The version must match the version configured for your NDMP backup application. Enter: <code>ndmpd version {2 3 4}</code> For details see “ Specifying a preferred network interface ” on page 159.

Step	Action
3	If you want to specify a restricted set of NDMP backup application hosts that can connect to the filer, enter: <pre>options ndmpd.access hosts</pre> <i>hosts</i> is a comma-separated list of host names or IP addresses of nodes permitted to start NDMP sessions with the storage system. If you do not wish to restrict NDMP access to this storage system by host, do not set this option. By default NDMP does not restrict access by host. For details see “Specifying NDMP access by host or interface” on page 151.
4	Specify the authentication type (plaintext, challenge, or plaintext <i>and</i> challenge) required for an NDMP connection to this storage system. For example: <pre>options ndmpd authtype plaintext,challenge</pre> This setting must include the authentication type supported by the NDMP backup application. For details see “Specifying the NDMP authentication type” on page 152.
5	If operators without root privilege on the storage system will be carrying out tape backup operations through the NDMP backup application ◆ Make sure they have a useradmin account on the storage system. If they do not, enter: <pre>useradmin useradd username</pre> ◆ Display that user’s system-generated NDMP-specific password. Enter: <pre>ndmpd password username</pre> When running the NDMP backup application, the user must input this NDMP-specific password to connect to this storage system and carry out NDMP backup and restore operations. For details, see “Generating an NDMP-specific password for non-root administrators” on page 155.

Enabling or disabling NDMP service

What enabling NDMP service does

Enabling NDMP service on your storage system allows NDMP-compliant data protection applications to communicate with the storage system.

Enabling or disabling NDMP service

To enable or disable NDMP service, complete the following step.

Step	Action
1	Enter the following command: <code>ndmpd {on off}</code> Use <code>on</code> to enable NDMP. Use <code>off</code> to disable NDMP. After you enter <code>ndmpd off</code>, the storage system continues processing all requests on already established sessions, but rejects new sessions.

Note

The setting of the `ndmpd on|off` command is persistent across reboots.

Managing NDMP security features

About NDMP security features

- Data ONTAP provides the following commands to enable you to manage and monitor the security of NDMP connections to the storage system.
- ◆ The options `ndmpd.access` command enables you to restrict which hosts can run NDMP sessions with the storage system. For details see [“Specifying NDMP access by host or interface”](#) on page 151.
 - ◆ The `ndmpd.authtype` command enables you to specify the authentication method (plaintext, challenge, or both) through which users are allowed to start NDMP sessions with the storage system. For details see [“Specifying the NDMP authentication type”](#) on page 152.
 - ◆ The options `ndmpd.connectlog` command allows you to enable or disable logging of NDMP connections attempts with the storage system. For details see [“Enabling or disabling NDMP connection logging”](#) on page 153.
 - ◆ The options `ndmpd.password_length` command allows you specify an 8-character or 16-character NDMP password. For details see [“Specifying the NDMP password length”](#) on page 154.
 - ◆ The `ndmpd.password` command enables even administrators with non-root level privilege on the storage system to obtain a system-generated secure password. This is the password that will allow them to carry out NDMP operations, either through an NDMP-compliant backup application or through the storage system command-line interface (CLI) that require a password. For details see [“Generating an NDMP-specific password for non-root administrators”](#) on page 155.

Specifying NDMP access by host or interface

You can use the options `ndmpd.access` configuration to specify the hosts or interfaces through which NDMP sessions are permitted. If you use this option to list specific authorized hosts or interfaces, NDMP sessions will be refused through any host or interface that is not listed by this option. Conversely, you can also specify hosts or interfaces to block from NDMP sessions.

To specify permitted hosts or interfaces, complete the following steps.

Step	Action
1	Start a console session on the storage system to which you want to restrict NDMP access.

Step	Action
2	Enter the following command: <pre>options ndmpd.access {all legacy host[!]=hosts if[!]=interfaces}</pre> <code>all</code> is the default value, which permits NDMP sessions with any host. <code>legacy</code> restores previous values in effect before a Data ONTAP version upgrade. In the case of Data ONTAP 6.2, the legacy value is equal to <code>all</code>. <code>host=hosts</code> is a parameter string that allows a specified host or a comma-separated list of hosts to run NDMP sessions on this storage system. The hosts can be specified by either hostname or IP address. <code>host!=hosts</code> is a parameter string that blocks a specified host or a comma-separated list of hosts from running NDMP sessions on this storage system. The hosts can be specified by either hostname or IP address. <code>if=interfaces</code> is a parameter string that allows NDMP sessions through a specified interface or a comma-separated list of interfaces on this storage system. <code>if!=interfaces</code> is a parameter string that blocks NDMP sessions through a specified interface or a comma-separated list of interfaces on this storage system.

Specifying the NDMP authentication type

Data ONTAP supports two methods, plaintext and challenge, for authenticating NDMP access to a storage system.

You can use the `options ndmpd.authtype` command to specify whether a storage system will accept plaintext, challenge, or both plaintext and challenge methods to authenticate NDMP session requests. To specify the authentication type, complete the following steps.

Step	Action
1	Start a console session on the storage system whose NDMP authentication method you want to specify.

Step	Action
2	Enter <i>one</i> of the following command lines: <pre>options ndmpd.auththtype challenge</pre> <pre>options ndmpd.auththtype plaintext</pre> <pre>options ndmpd.auththtype plaintext,challenge</pre> The first command line accepts the challenge authentication method, generally the preferred and more secure authentication method. The second command line accepts the plaintext authentication method, in which the login password is transmitted as clear text. The third command line accepts both challenge and plaintext authentication methods. If you are carrying out NDMP operations through a backup application, the authentication type or types you specify on this command line must include the types supported by that backup application.

Enabling or disabling NDMP connection logging

Data ONTAP can log NDMP connection attempts in the /etc/messages file. These entries enable an administrator to determine whether and when authorized or unauthorized individuals are attempting to start NDMP sessions. The default is off.

To enable or disable logging, complete the following steps.

Step	Action
1	Start a console session on the storage system on which you want to enable or disable NDMP connection monitoring.
2	Enter the following command: <pre>options ndmpd.connectlog.enabled {on off}</pre> The value you set for this option will persist across storage system reboots.

Step	Action
3	If you want to check attempted NDMP connection activity, use your UNIX or Windows Admin host to view your storage system's /etc/messages file. Entries recording attempted NDMP connections or operations will display the following fields: ◆ time ◆ thread ◆ NDMP request and action (allow or refuse) ◆ NDMP version ◆ session ID ◆ source IP (address of where the NDMP request originated) ◆ destination IP (address of the storage system receiving the NDMP request) ◆ source port (port through which the NDMP request was transmitted) ◆ storage system port through which the NDMP request was received Example entry: <pre>Friday Sept 13:16:45:17GMT ndmpd.access allowed for version =4, sessid=34, from src ip = 172.29.19.40, dst ip =172.29.19.95, src port = 63793, dst port = 10000.</pre>

Specifying the NDMP password length

Administrators who have an account on a storage system but do not have root status on that storage system need to input a special NDMP-specific password when carrying out NDMP-related operations on the storage system. This password is a system-generated string derived from that administrator's regular storage system account password.

The NDMP password can be either 8 or 16 characters long. The default value is 16 characters. To specify the NDMP password length, complete the following step.

Step	Action
1	On the storage system console, enter the following command: <pre>options ndmpd.password_length length</pre> <i>length</i> is either 8 or 16. If you enter a value other than 8 or 16, the storage system prompts you with the following message: <pre>options ndmpd.password_length: Length must be either 8 or 16</pre> Note _____ If this option is set to 8, all NDMP applications managing backups for the storage system must use a 8 character password for authentication.

Generating an NDMP-specific password for non-root administrators

The administrator without root privilege uses this NDMP-specific password for any NDMP backup and restore operation that requires password input in either a backup application or command-line interface (CLI) environment.

Note _____

For information about creating accounts for administrators without root privilege, see the section about managing local user accounts in the introduction to storage system administration chapter of the *System Administration Guide*.

The administrator without root privilege obtains an NDMP-specific password by completing the following steps:

Step	Action
1	Start a console session on the storage system you want to access.

Step	Action
2	Enter the following command: <pre>ndmpd password username</pre> <i>username</i> is the username of the administrator. The system returns an 8 or 16 character string, depending on the <code>ndmpd.password_length</code> option value. For example: <pre>>ndmp password barbaraD >ndmp password: QM12N%\$cnaFWPBVe</pre> You use this password in any current or future NDMP operation that requires password input. Note_____ This NDMP-specific password is valid until you change the password to your regular account.
3	If you change the password to your regular storage system account, repeat this procedure to obtain your new system-generated NDMP-specific password.

Specifying the NDMP version

Why you need to specify the NDMP version

Data ONTAP provides full support for NDMP versions 2 to 4. Data ONTAP is shipped with the NDMP version set to 4, as both the default version and the maximum version.

The storage system and the backup application must agree on a version of NDMP to be used for each NDMP session. When the backup application connects to the storage system, the storage system sends the default version back. The application can choose to use that default version and continue with the session. However, if the backup application uses an earlier version, it begins version negotiation, asking if each version is supported, to which the storage system responds yes or no.

The `ndmp version` command sets an upper limit (maximum version) as well as a default version for the versions that the storage system says, during the negotiation process, that it supports.

Because some backup applications do not support version negotiation, the `ndmp version` command controls the maximum and default NDMP version allowed. If you know your backup application does not support NDMP version 4 and does not negotiate versions, you can use this command to define the maximum version Data ONTAP supports so that the application can operate correctly.

Displaying the NDMP version

To see the maximum version of NDMP the storage system is currently set to use, complete the following step.

Step	Action
1	Enter the following command: <code>ndmpd version</code> Result: The maximum version that NDMP currently allows to be used is displayed.

Specifying the NDMP version

To specify the NDMP version you want, complete the following step.

Step	Action
1	Enter the following command: ndmpd version <i>n</i> <i>n</i> is the number of the version you want to specify. Options available are 2, 3, and 4. The default maximum version is 4.

Note_____

The `ndmpd version` command is persistent across storage system reboots.

Specifying a preferred network interface

Why to specify a preferred network interface

You can specify the preferred storage system network interface to be used when establishing an NDMP data connection to another storage system.

By default, an NDMP data connection uses the same network interface as the NDMP control connection established by the NDMP backup application. However, when a data connection between NDMP-enabled storage systems needs to be established over an alternate network, it is necessary to specify the storage system’s interface through which the alternate network will be accessed.

For example, a UNIX or NT resident NDMP backup application and multiple storage systems can be interconnected via a corporate network. The same storage systems can also be interconnected via an isolated private network. To minimize load on the corporate network, the `ndmpd.preferred_interface` option can be used to direct all NDMP data connections over the isolated private network.

How to specify a preferred network interface

To specify the preferred network interface to be used for NDMP data connections, complete the following step.

Step	Action
1	Enter the following command: <pre>options ndmpd.preferred_interface interface</pre> <i>interface</i> identifies the network interface to be used for all NDMP data connections. Any network interface providing TCP/IP access can be specified. If no argument is specified, the command returns the name of the interface currently configured for data connections. If no interface is currently set, it reports <i>disable</i>. You can find the available network interfaces by using the <code>ifconfig -a</code> command.

Note_____

The `ndmpd.preferred_interface` option is persistent across storage system reboots.

How to turn off data connection specification

To disable a preferred network interface specification and force the NDMP default interface to be used for data connections, complete the following step.

Step	Action
1	Enter the following command: <pre>options ndmpd.preferred_interface disable</pre> Result: NDMP does not use the preferred interface for data connection. The default value for the <code>ndmpd.preferred_interface</code> option is <i>disable</i>.

Enabling or disabling enhanced DAR functionality

What enhanced DAR functionality is

Enhanced Direct Access Restore (DAR) functionality provides support for directory DAR, DAR of files with NT Streams, and capacity and performance improvements.

By default, enhanced DAR functionality is disabled in Data ONTAP; however, you can enable it by using the `ndmpd.offset_map.enable` option.

Enabling enhanced DAR functionality

To enable enhanced DAR functionality on your storage system, complete the following step.

Note

Before enabling enhanced DAR functionality, make sure you are aware of the performance impact this functionality can cause. See “[Considerations when using NDMP](#)” on page 144 for more information.

Step	Action
1	Enter the following command: <code>options ndmpd.offset_map.enable on</code>

Disabling enhanced DAR functionality

To disable enhanced DAR functionality on your storage system, complete the following step.

Step	Action
1	Enter the following command: <code>options ndmpd.offset_map.enable off</code>

Specifying environment variables for backup applications

What environment variables are

Environment variables are used to communicate information about a backup or restore operation between an NDMP-enabled backup application and a storage system. For example, if a user specifies that a backup application should back up /vol/vol0/etc, the backup application sets the FILESYSTEM environment variable to /vol/vol0/etc. Similarly, if a user specifies that a backup should be a level-1 backup, the backup application sets the LEVEL environment variable to 1 (one).

Note _____
Setting and examining of environment variables are typically transparent to backup administrators; that is, the backup application sets them automatically.

When backup administrators might specify environment variables

A backup administrator rarely specifies environment variables; however, you might want to change the value of an environment variable from that set by the backup application to characterize or work around a functional or performance problem. For example, an administrator may want to temporarily disable file history generation to determine if the backup application's processing of file history information is contributing to performance issues or functional problems.

Many backup applications provide a means to override or modify environment variables or to specify additional environment variables. See your backup application documentation for information.

Supported environment variables

The following table contains descriptions of what the environment variables supported by Data ONTAP do if they are used. The default values listed in the table are set by Data ONTAP and might be modified by your backup application. If you modify manually the values set by the backup application, as needed for it to work properly, the application might behave unpredictably because the dump or restore operations might not be doing what the backup application expected them to do. But in some cases, judicious modifications might help in identifying or working around problems.

Note _____
In most cases, variables that have Y or N values also accept T or F values, respectively.

Environment Variable	Valid Values	Default	Description
ACL_START	<i>return only</i>	none	Created by the backup operation, the ACL_START variable is an offset value used by a direct access restore or restartable NDMP backup operation. The offset value is the byte offset in the dump file where the ACL data (Pass V) begins and is returned at the end of a backup. For a direct access restore operation to correctly restore backed up data, the ACL_START value must be passed to the restore operation when it begins. An NDMP restartable backup operation uses the ACL_START value to tell the backup application where the nonrestartable portion of the backup stream begins.
BASE_DATE	0, -1, or <i>DUMP_DATE</i> <i>value</i>	-1	Specifies the start date for incremental backups. There is no equivalent option for the <code>dump</code> command. When set to -1, the BASE_DATE incremental specifier is disabled. When set to 0 on a level 0 backup, incremental backups are enabled. Subsequent to the initial backup, the value of the DUMP_DATE variable from the previous incremental backup is assigned to the BASE_DATE variable. These variables are an alternative to the <code>/etc/dumpdates</code> file for controlling incremental backups.
DATA_BLOCK_SIZE	4-64	63	Specifies the tape blocking factor in kilobytes. Note _____NDMP backups do not allow tape block sizes bigger than 64KB. An attempt to use block sizes of greater than 64KB will result in failed backups._____

Environment Variable	Valid Values	Default	Description
DEBUG	Y or N	N	Specifies that debugging information is printed. Debug information is not printed if only the VERBOSE variable is set. The messages printed when the DEBUG or VERBOSE variables are set are mutually exclusive, debug messages are not a superset of verbose messages and vice versa. There is no equivalent option for the <code>dump</code> or <code>restore</code> command.
DIRECT	Y or N	N	Specifies that a restore should fast-forward directly to the location on the tape where the file data resides instead of scanning the entire tape. For direct access restoration to work, the backup application must provide positioning information; therefore, you would rarely modify the DIRECT variable. If you specified the value Y, the backup application would specify the file or directory names and the positioning information in its <code>nlist</code> . The <code>nlist</code> information cannot be modified by changing environmental variable values.
DMP_NAME	<i>string</i>	none	Specifies the name for a multiple subtree backup. The DMP_NAME variable is equivalent to the <code>n</code> option of the <code>dump</code> command. This variable is mandatory for multiple subtree backups.
DUMP_DATE	<i>return value</i>	none	You do not change this variable directly. It is created by the backup if the BASE_DATE variable is set to a value other than -1. The BASE_DATE variable is derived by prepending the 32-bit level value to a 32-bit time value computed by the dump software. The level is incremented from the last level value passed into the BASE_DATE variable. The resulting value is used as the BASE_DATE value on a subsequent incremental backup.

Environment Variable	Valid Values	Default	Description
ENHANCED_DAR_ENABLED	Y or N	N	Specifies if enhanced Direct Access Restore (DAR) functionality is instantiated. Enhanced functionality consists of support for directory DAR, DAR of files with NT Streams, and capacity and performance improvements. The ENHANCED_DAR_ENABLED variable is not checked at the beginning of a backup and, for correct operation, should not be set. Enhanced DAR is instantiated with a T value, not Y, at the conclusion of a backup only if the following is true: ◆ Data ONTAP supports enhanced DAR (Data ONTAP 6.4 or later) ◆ File history is enabled (HIST=Y) ◆ Creation of enhanced DAR information on the dump file completed successfully When restoring, the value that was returned by the Data Get Environment NDMP protocol message after the backup is passed to the restore operation. If you want to disable the enhanced DAR features when restoring data, set the ENHANCED_DAR_ENABLED variable to F or N. Note_____ You must not set the ENHANCED_DAR_ENABLED variable to T when restoring data if enhanced DAR was not instantiated because the data to support enhanced DAR is not on the dump file. _____

Environment Variable	Valid Values	Default	Description
EXCLUDE	<i>pattern string</i>	none	Specifies files or directories that are excluded when backing up data. The EXCLUDE variable is equivalent to the <code>x</code> option of the <code>dump</code> command. The exclude list is a comma-separated list of file or directory names. If the name of a file or directory matches one of the names in the list, it is excluded from the backup. The following are rules for specifying names in the exclude list: ◆ The exact name of the file or directory must be used. ◆ An asterisk (*) is a wildcard character. The asterisk must be either the first or the last character of the string. Each string can have up to two asterisks. ◆ A comma in a file or directory name must be preceded with a backslash. ◆ The exclude list can contain up to 32 names.
EXTRACT	Y or N	N	Specifies that subtrees of a backed-up dataset are to be restored. The EXTRACT variable is equivalent to the <code>x</code> option of the <code>restore</code> command. The backup application specifies in the <code>nlist</code> the names of the subtrees to be extracted. If a file named in the <code>nlist</code> matches a directory whose contents were backed up, the directory is recursively extracted.
EXTRACT_ACL	Y or N	Y	Specifies that ACLs from the backed up file are restored on a restore operation. The EXTRACT_ACL variable is equivalent to the <code>A</code> option of the <code>restore</code> command. The default is to restore ACLs when restoring data, except for direct access restores (DIRECT=Y).
FILESYSTEM	<i>string</i>	None	Specifies the pathname of the root of the data that is being backed up. For example, <code>/vol/vol0/etc</code>.

Environment Variable	Valid Values	Default	Description
FORCE	Y or N	N	Specifies that a restore operation continues, regardless of inode limitations. The FORCE variable is equivalent to the F option of the <code>restore</code> command. When this variable is set to N, if the restore operation determines that there are fewer free inodes than the number of files it needs to create, it aborts. Setting the variable to Y causes the restore operation to proceed on the assumption that new files overwrite older files and that the file system will not run out of inodes. If the restore operation runs out of inodes, the restore operation aborts during its run.
HIST	Y or N	N	Specifies that file history information is sent to the backup application. Most commercial backup applications set the HIST variable to Y. If you want to increase the speed of a backup operation, or you want to troubleshoot a problem with file history collection, you can set this variable to N. Note _____ You should not set the HIST variable to Y if the backup application does not support file history. _____
IGNORE_CTIME	Y or N	N	Specifies that a file is not incrementally backed up if only its ctime value has changed since the previous incremental backup. Some applications, such as virus scanning software, change the ctime value of a file within the inode, even though the file or its attributes have not changed. As a result, an incremental backup might back up files which have not changed. The IGNORE_CTIME variable should be specified only if incremental backups are taking an unacceptable amount of time or space because the ctime value was modified.

Environment Variable	Valid Values	Default	Description
IGNORE_QTREES	Y or N	N	Specifies that the restore operation does not restore qtree information from backed up qtrees. The IGNORE_QTREES variable is equivalent to the <code>Q</code> option of the <code>restore</code> command.
LEVEL	0-9	0	Specifies the backup level. Level 0 copies the entire dataset. Incremental backup levels, specified by values above 0, copy all files new or modified since the last incremental backup. For example, a level 1 backs up new or modified files since the level 0 backup, a level 2 backs up new or modified files since the level 1 backup, and so on.
LIST	Y or N	N	Specifies that backed-up file names and inode numbers are listed as they are restored. The LIST variable is equivalent to the <code>t</code> option of the <code>restore</code> command.
LIST_QTREES	Y or N	N	Specifies that backed-up qtrees are listed as are restored. The LIST_QTREES variable is equivalent to the <code>T</code> option of the <code>restore</code> command.
MULTI_SUBTREE_NAMES	<i>string</i>	none	Specifies that the backup is a multiple subtree backup. The MULTI_SUBTREE_NAMES variable is equivalent to the <code>l</code> option of the <code>dump</code> command. Multiple subtrees are specified in the string which is a newline-separated, null-terminated list of subtree names. Subtrees are specified by path names relative to their common root directory, which must be specified as the last element of the list. If you use this variable, you must also use the <code>DMP_NAME</code> variable.

Environment Variable	Valid Values	Default	Description
NDMP_UNICODE_FH	Y or N	N	Specifies that a Unicode name is included in addition to the NFS name of the file in the file history information. This option is not used by most backup applications and should not be set unless the backup application is designed to receive these additional file names. The HIST variable must also be set.
NDMP_VERSION	<i>return only</i>	none	You should not modify the NDMP_VERSION variable. Created by the backup operation, the NDMP_VERSION variable returns the NDMP version. Data ONTAP sets the NDMP_VERSION variable during a backup for internal use and to pass to a backup application for informational purposes. The NDMP version of an NDMP session is not set with this variable.
NO_ACLS	Y or N	N	Specifies that Access Control Lists (ACLs) are not copied when backing up data. The NO_ACLS variable is equivalent to the A option of the <code>dump</code> command. Ordinarily a backup using the <code>dump</code> command writes out metadata related to Windows ACLs. The NO_ACLS variable stops this information from being backed up.
NON_QUOTA_TREE	Y or N	N	Specifies that files and directories in qtrees are ignored when backing up data. The NON_QUOTA_TREE variable is equivalent to the Q option of the <code>dump</code> command. When set to Y, items in qtrees in the dataset specified by the FILESYSTEM variable are not backed up. This variable has an effect only if the FILESYSTEM variable specifies an entire volume. The NON_QUOTA_TREE variable only works on a level-0 backup and does not work if the MULTI_SUBTREE_NAMES variable is specified.

Environment Variable	Valid Values	Default	Description
NOWRITE	Y or N	N	Specifies that the restore operation does not write data to the disk. The NOWRITE variable is equivalent to the <code>N</code> option of the <code>restore</code> command. This variable is used for debugging.
REPLICATE	Y or N	N	Specifies an unlimited number of incremental backups. There is no equivalent option for the <code>dump</code> command. The REPLICATE variable is ignored when the <code>BASE_DATE</code> and <code>DUMP_DATE</code> variables are used to manage incrementals.
RESTART	Y or N	N	Not supported by NDMP software; the RESTART variable is ignored.
UPDATE	Y or N	N	Specifies that the <code>dumpdates</code> file is updated after backing up data. The UPDATE variable is equivalent to the <code>u</code> option of the <code>dump</code> command. When set to Y, the <code>/etc/dumpdates</code> file (used for managing incremental backups) is updated after a successful backup.
VERBOSE	Y or N	N	Specifies that detailed information about dump and restore operations is printed. The VERBOSE variable is equivalent to the <code>v</code> option of the <code>dump</code> and <code>restore</code> commands.

Displaying NDMP session status

Reasons to display NDMP session status

You can view either general status information or detailed session information about one or more NDMP sessions:

- ◆ View general status information to determine whether the session is operating as expected.
- ◆ View detailed session information to help you debug an NDMP session if there are problems.

How to display general status information about an NDMP session

To display general NDMP status information, complete the following step.

Step	Action
1	Enter the following command: ndmpd status [<i>session</i>] <i>session</i> is the specific session number you want the status of, from 0 to 99. Leave <i>session</i> blank to display the status of all current sessions.

Example

In the following example, the command displays information about session 4:

```
filerA> ndmpd status 4

ndmpd ON.
Session: 4
  Active
version:          3
Operating on behalf of primary host.
tape device:      not open
mover state:      Idle
data state:       Idle
data operation:   None
```

How to display
detailed NDMP
session information

To display detailed NDMP session information, complete the following step.

Step	Action
1	Enter the following command: ndmpd probe [session] <i>session</i> is the number of the session you want to probe, from 0 to 99. Leave <i>session</i> blank to display detailed information about all sessions.

Example

In the following example, the command shows detailed status of session 4:

```
filerA> ndmpd probe 4

ndmpd ON.
Session: 4
  isActive:                TRUE
  protocol version:        3
  effHost:                  Local
  authorized:               FALSE
  client addr:              10.10.10.12.47154
  spt.device_id:            none
  spt.ha:                   -1
  spt.scsi_id:              -1
  spt.scsi_lun:             -1
  tape.device:              rst0a
  tape.mode:                Read/Write
  mover.state:              Active
  mover.mode:               Read
  mover.pauseReason         N/A
  mover.haltReason          N/A
  mover.recordSize:         10240
  mover.recordNum:          315620
  mover.dataWritten:        3231948800
  mover.seekPosition:       0
  mover.bytesLeftToRead:    0
  mover.windowOffset:       0
  mover.windowLength:       -1
  mover.position:           0
  mover.connect.addr_type:  LOCAL
  data.operation:           Backup
  data.state:               Active
```

```
data.haltReason:      N/A
data.connect.addr_type: LOCAL
data.bytesProcessed:  3231989760
```

Displaying NDMP debug messages

What debug messages are

NDMP debug messages provide a detailed description of all active NDMP sessions. The amount of information displayed by a debug message is determined by the `ndmpd` debug level specified by the user. By default, debug messages are disabled.

Where the storage system saves debug messages

Debug messages can be output to the storage system console, the NDMP session log file, or both. The `ndmpd debug` command is used to specify where debug messages will be output. By default, debug messages are output to both the storage system console and the NDMP session log.

About the NDMP session log files

NDMP session log files are stored in the `/etc/log` directory. The file name is `ndmplog.yyyymmdd`, where `yyyy` is the year, `mm` is the month, and `dd` is the date. For example, the session log file generated on December 5, 2000, is named `ndmplog.20001205`. The session log file can contain information about one or more NDMP sessions. If multiple NDMP sessions take place on the same day, Data ONTAP saves the information about all sessions to the same session log file. Before generating a fresh NDMP session log file, Data ONTAP deletes all files more than eight days old. Data ONTAP keeps a maximum of eight session log files: one each for the past seven days plus the current day.

Enabling the NDMP debug log messages

To enable NDMP debug logging, complete the following step.

Step	Action
1	Enter the following command: <pre>ndmpd debug <i>n</i></pre> <i>n</i> specifies the debug level from 0 to 70. The default level is 0. Use 0 for the level to turn the debug messages off. Use a nonzero value for the level to turn the debug messages on. The levels supported are: 10—Displays connections being made and connections being closed. 30—Displays information regarding the actual NDMP messages such as the message type, sequence numbers, and timestamps. This level also prints out the NDMP errors and some of the relevant fields of the NDMP message. 50—Same as level 30, but includes display of environment variables as well as any exceptions issued by the NDMP server implementation. 70—Same as level 50, but includes display of tape and SCSI command descriptor blocks (CDBs) sent. Note _____ CDBs are used for low-level tape and medium changer control. _____

Displaying the NDMP debug log level

To see the NDMP debug levels currently set, complete the following step.

Step	Action
1	Enter the following command: <pre>ndmpd debug</pre> Result: The current NDMP debug level and toggles are displayed.

Example:

```
filerA> ndmpd debug
ndmpd debug verbose: 0
ndmpd debug stack trace: false
ndmpd debug screen trace: true
ndmpd debug file trace: true
```

Changing NDMP debug log messages

You use the `ndmp debug` command to choose whether to display the debug messages on the storage system console, to write the debug messages into the NDMP debug log file, and to print out stack traces for any exceptions issued by the NDMP server implementation. To toggle these options, complete the following step:

Step	Action
1	Enter the following command: ndmpd debug stack screen file stack toggles stack trace printouts on or off. screen toggles printouts to the storage system console on or off. file toggles printouts to the NDMP log file on or off.

Displaying an NDMP session log file

To display an NDMP session log file, complete the following step.

Step	Action	
1	From a UNIX environment...	From an NT environment...
	Mount the root volume of the storage system to a UNIX client and view the contents of the NDMP session log file using the <code>cat</code> or <code>more</code> UNIX commands or a text editor.	Map the root volume of the storage system to the NT system and view the contents of the NDMP session log file using WordPad, Notepad, or an equivalent text-viewing application.

Terminating an NDMP session

How to terminate an NDMP session

You terminate an NDMP session if the session is not responding.
To terminate an NDMP session, complete the following step.

Step	Action
1	Enter the following command: <code>ndmpd kill session #</code> <i>session #</i> is the specific NDMP session you want to terminate, from 0 to 99. If you want to terminate all NDMP sessions, enter the following command. <code>ndmpd killall</code> These commands allow nonresponding sessions to be cleared without the need for a reboot, since the <code>ndmpd off</code> command waits until all sessions are inactive before turning off the NDMP service.

Copying data with ndmpcopy

What ndmpcopy is

The `ndmpcopy` command enables a storage system administrator to transfer file system data between storage systems that support NDMP v3 or v4 and the UNIX file system (UFS) dump format.

Using the `ndmpcopy` command, you can carry out both full and incremental data transfers; however, incremental transfers are limited to a maximum of two levels (one full and up to two incremental). You can transfer full or partial volumes, qtrees, or directories, but not individual files.

`ndmpcopy` functions as a simple NDMP data management application (backup application) that performs data transfers by initiating a backup operation on the source storage system and recovery operation on the destination storage system. It establishes separate NDMP control connections with the source and destination storage systems, creates a single NDMP data connection between the storage systems, and initiates backup and recovery operations that result in the desired data transfer between the source and destination.

You can invoke it at the command line of the source storage system, the destination storage system, or a storage system that is neither the source nor the destination of the data transfer.

You can also invoke `ndmpcopy` on a single storage system that is both the source and the destination of the data transfer. See the `ndmpcopy` examples, next, for more information on supported data transfer topologies.

How to copy data using ndmpcopy

To copy data within a storage system or between storage systems using `ndmpcopy`, complete this step.

Note

The command can also be entered from a storage system that is not the source or the destination.

Step	Action
1	Enter the following command: <pre>ndmpcopy [options] [source_filer:]source_path [destination_filer:]destination_path</pre> Descriptions of the options and arguments follow.

options can be one or more of the following:

-sa *username:[password]* source authorization specifies the user name and password for connecting to the source storage system.

Note

If you specify a user without root privilege who requires a password, you must specify that user's system-generated NDMP-specific password rather than that user's regular storage system account password. See [“Generating an NDMP-specific password for non-root administrators”](#) on page 155.

-da *username:[password]* destination authorization specifies the user name and password for connecting to the destination storage system.

-st {challenge|text} sets the source authentication type to be used when connecting to the source storage system.

-dt {challenge|text} sets the destination authentication type to be used when connecting to the destination storage system.

By default, challenge is the authentication type used. The text authentication type exchanges the user name and password in clear text. The challenge authentication type exchanges the user name and password in encrypted form.

-l *level* sets the dump level used for the transfer to the specified value of *level*. Valid values for *level* are 0, 1, and 2, where 0 indicates a full transfer and 1 or 2 an incremental transfer. The default is 0.

-d enables generation of `ndmpcopy` debug log messages. `ndmpcopy` debug log files appear in the root volume `/etc/log` directory. The `ndmpcopy` debug log file names are in the form `ndmpcopy.yyyymmdd`.

-f enables forced mode. This mode enables overwriting system files in the `/etc` directory on the root volume. See Example 5 for details.

-h prints the help message.

source_filer and *destination_filer* can be host names or IP addresses

If *destination_path* does not specify a volume, or specifies a non-existent volume, the root volume is used.

Examples of the ndmcopy command

In these examples, network host names are used for the storage systems. (“myhost” is used for a local storage system and “remotehost1” and “remotehost2” are used for remote storage systems.) If you specify host names when you use the `ndmcopy` command, the storage system running the `ndmcopy` command should be able to resolve these names to their IP addresses. In UNIX, you can use the `ping` command to make sure that host name resolution works.

Example 1: This command migrates data from a source path (*source_path*) to a different destination path (*destination_path*) on the same storage system (myhost).

```
myhost> ndmcopy -sa username:password -da username:password  
myhost:/vol/vol0/source_path myhost:/vol/vol0/destination_path
```

This shorter form of the command achieves the same purpose:

```
myhost> ndmcopy /vol/vol0/source_path /vol/vol0/destination_path
```

Because you are running the `ndmcopy` command on myhost and the source and destination storage system are the same as myhost, you can omit the source and destination storage system names on the `ndmcopy` command line. When your `ndmcopy` command is running on the same storage system as the source storage system or destination storage system, you can also omit the `-sa` or `-da` options.

Example 2: This command migrates data from a source path (*source_path*) to a different destination path (*destination_path*) on remotehost1.

```
myhost> ndmcopy -da username:password /vol/vol0/source_path  
remotehost1:/vol/vol0/destination_path
```

The destination storage system must be specified in this case, because it is a remote storage system. The destination authorization is needed, but not the source authorization.

Example 3: This command migrates data from a source path (*source_path*) on remotehost2 to a destination path (*destination_path*) on myhost.

```
myhost> ndmcopy -sa username:password -st text  
remotehost2:/vol/vol0/source_path /vol/vol0/destination_path
```

The source authentication type specified by `-st` is text. The `ndmcopy` command tool running on `myhost` will authenticate with the source storage system using text authentication.

Example 4: This command migrates data from a source path (*source_path*) on `remotehost1` to a destination path (*destination_path*) on `remotehost2`.

```
myhost> ndmcopy -sa username:password -da username:password -l 1  
remotehost1:/vol/vol0/source_path  
remotehost2:/vol/vol0/destination_path
```

The `-l 1` option is used to do a level 1 transfer.

Example 5: Without the `-f` option, the `/etc` directory and its contents on the root volume of `remotehost1` are protected from being overwritten with the `/etc` directory from `myhost`, preventing unintentional changing of the system characteristics after the root volume migration is completed.

```
myhost> ndmcopy -da username:password /vol/rootvol  
remotehost1:/vol/rootvol
```

To intentionally overwrite the `/etc` directory during the root volume migration, use the `-f` flag as in the following example.

```
myhost> ndmcopy -da username:password -f /vol/rootvol  
remotehost1:/vol/rootvol
```


About this chapter

This chapter describes how Data ONTAP automatically records the events that occur during dump and restore operations in log files. You can use these event logs to monitor, manage, and diagnose problems with your data backup and restore operations.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“About event logging for backup and restore”](#) on page 184
- ◆ [“Event log messages”](#) on page 185
- ◆ [“Enabling or disabling event logging”](#) on page 190

About event logging for backup and restore

What event logging is

Data ONTAP automatically logs significant events and the times at which they occur during dump and restore operations. All dump and restore events are recorded in a log file named backup in the /etc/log/ directory. Event logging is turned on by default.

Reasons for viewing event log files

You might want to view event log files for the following reasons:

- ◆ To find out whether a nightly backup was successful
- ◆ To gather statistics on backup operations
- ◆ To use information contained in past event log files to help diagnose problems with dump and restore operations

Log file rotation

Once every week, the log files are rotated. The /etc/log/backup file is copied to /etc/log/backup.0, the /etc/log/backup.0 file is copied to /etc/log/backup.1, and so on. The system saves the log files for up to six weeks; therefore, you can have up to seven message files (/etc/log/backup.0 through /etc/log/backup.5 and the current /etc/log/backup file).

Event log files in takeover mode

If a takeover occurs in a cluster, the set of backup log files for the takeover storage system remains separate from the backup log files for the failed storage system.

Event log messages

Event log message format For each event, a message is written to the backup log file in the following format:

type timestamp identifier event (event_info)

Example: dmp Thu Apr 5 18:54:56 PDT /vol/vol0/home(5) Start (level 0, NDMP)

Type field Each log message begins with one of the type indicators described in the following table.

Type	Meaning
log	Logging event
dmp	Dump event
rst	Restore event

Timestamp field The timestamp field shows the date and time of the event.

Identifier field The identifier field for a dump event includes the dump path and the unique ID for the dump. The identifier field for a restore event uses only the restore destination path name as a unique identifier. Logging-related event messages do not include an identifier field.

Start and stop logging events

The event field of a message that begins with `log` contains one of the events shown in the following table.

Event	Meaning
Start_Logging	Indicates the beginning of logging or that logging has been turned back on after being disabled.
Stop_Logging	Indicates that logging has been turned off.

Dump events

The event field for a dump event contains an event type followed by event-specific information in parentheses. The following table describes the events, their meanings, and the related event information that might be recorded for a dump operation.

Event	Meaning	Event Information
Start	A dump or NDMP dump begins.	Dump level and the type of dump
Restart	A dump restarts.	Dump level
End	Dumps completed successfully.	Amount of data processed
Abort	The operation aborts.	Amount of data processed
Options	Specified options are listed.	All options and their associated values, including NDMP options
Tape_open	The tape is open for read/write.	The new tape device name
Tape_close	The tape is closed for read/write.	The tape device name
Phase-change	Dump is entering a new processing phase.	The new phase name
Error	Dump encounters an unexpected event.	Error message

Event	Meaning	Event Information
Snapshot copy	A Snapshot copy is created or located.	The name and time of the Snapshot copy
Base_dump	A base dump entry in the etc/dumpdates files has been located.	The level and time of the base dump (for incremental dumps only)

Example for dump

The following is an example of the output for a dump operation:

```
dmp Thu Sep 20 01:11:22 GMT /vol/vol0/(1) Start (Level 0)

dmp Thu Sep 20 01:11:22 GMT /vol/vol0/(1) Options (b=63, B=1000000,
u)

dmp Thu Sep 20 01:11:22 GMT /vol/vol0/(1) Snapshot
(snapshot_for_backup.6, Sep 20 01:11:21 GMT)

dmp Sep 20 01:11:22 GMT /vol/vol0/(1) Tape_open (nrst0a)

dmp Sep 20 01:11:22 GMT /vol/vol0/(1) Phase_change (I)

dmp Sep 20 01:11:24 GMT /vol/vol0/(1) Phase_change (II)

dmp Sep 20 01:11:24 GMT /vol/vol0/(1) Phase_change (III)

dmp Sep 20 01:11:26 GMT /vol/vol0/(1) Phase_change (IV)

dmp Sep 20 01:14:19 GMT /vol/vol0/(1) Tape_close (nrst0a)

dmp Sep 20 01:14:20 GMT /vol/vol0/(1) Tape_open (nrst0a)

dmp Sep 20 01:14:54 GMT /vol/vol0/(1) Phase_change (V)

dmp Sep 20 01:14:54 GMT /vol/vol0/(1) Tape_close (nrst0a)

dmp Sep 20 01:14:54 GMT /vol/vol0/(1) End (1224 MB)
```

There are five phases in a dump operation (map files, map directories, dump directories, dump files, and dump ACLs).

The log file for a dump operation begins with either a Start or Restart event and ends with either an End or Abort event.

Restore events

The event field for a restore event contains an event type followed by event-specific information in parentheses. The following table describes the events, their meanings, and the related event information that might be recorded for a restore operation.

Event	Meaning	Event Information
Start	A restore or NDMP restore begins.	Restore level and the type of restore
Restart	A restore restarts.	Restore level
End	Restores completed successfully.	Number of files and amount of data processed
Abort	The operation aborts.	Number of files and amount of data processed
Options	Specified options are listed.	All options and their associated values, including NDMP options
Tape_open	The tape is open for read/write.	The new tape device name
Tape_close	The tape is closed for read/write.	The tape device name
Phase-change	Restore is entering a new processing phase.	The new phase name
Error	Restore encounters an unexpected event.	Error message

Example for restore

The following is an example of the output for a restore operation:

```
rst Thu Sep 20 02:24:22 GMT /vol/rst_vol/ Start (level 0)

rst Thu Sep 20 02:24:22 GMT /vol/rst_vol/ Options (r)

rst Thu Sep 20 02:24:22 GMT /vol/rst_vol/ Tape_open (nrst0a)

rst Thu Sep 20 02:24:23 GMT /vol/rst_vol/ Phase_change (Dirs)
```

```
rst Thu Sep 20 02:24:24 GMT /vol/rst_vol/ Phase_change (Files)
rst Thu Sep 20 02:39:33 GMT /vol/rst_vol/ Tape_close (nrst0a)
rst Thu Sep 20 02:39:33 GMT /vol/rst_vol/ Tape_open (nrst0a)
rst Thu Sep 20 02:44:22 GMT /vol/rst_vol/ Tape_close (nrst0a)
rst Thu Sep 20 02:44:22 GMT /vol/rst_vol/ End (3516 files, 1224 MB)
```

There are two phases in a restore operation (restore directories and restore files).

The log file for a restore operation begins with either a Start or Restart event and ends with either an End or Abort event.

Example of an aborted restore

The following is an example of the output of an aborted restore operation:

```
rst Thu Sep 20 02:13:54 GMT /rst_vol/ Start (Level 0)
rst Thu Sep 20 02:13:54 GMT /rst_vol/ Options (r)
rst Thu Sep 20 02:13:54 GMT /rst_vol/ Tape_open (nrst0a)
rst Thu Sep 20 02:13:55 GMT /rst_vol/ Phase_change (Dirs)
rst Thu Sep 20 02:13:56 GMT /rst_vol/ Phase_change (Files)
rst Thu Sep 20 02:23:40 GMT /vol/rst_vol/ Error (Interrupted)
rst Thu Sep 20 02:23:40 GMT /vol/rst_vol/ Tape_close (nrst0a)
rst Thu Sep 20 02:23:40 GMT /vol/rst_vol/ Abort (3516 files, 598 MB)
```

Enabling or disabling event logging

Turning event logging on or off

To turn event logging on or off, complete the following step.

Step	Action
1	Enter the following command to enable or disable event logging: <code>options backup.log.enable [on off]</code> <i>on</i>—turns event logging on. <i>off</i>—turns event logging off. Note _____ Event logging is turned on by default. _____

Glossary

ACL	Access Control List. A list that contains the users' or groups' access rights to each share.
active file system	A file system excluding its Snapshot copies.
blocking factor	The number of tape blocks that are transferred in each write operation.
CIFS	Common Internet File System. A protocol for networking PCs.
console	A terminal that is attached to a storage system's serial port and is used to monitor and manage storage system operation.
DNS	Domain Naming Service. An Internet service for finding IP addresses.
dump path	A path that specifies one volume, qtree, or subtree to back up.
file mark	Data on a tape that signals a boundary of a tape file.
HTTP	HyperText Transfer Protocol. An Internet Web transfer protocol.
increment chain	A series of incremental backups of the same path.
inode	A data structure containing information about files on a storage system and in a UNIX file system.

local tape device	A program-based functionality associated with a tape drive that is directly attached to a storage system that is performing a tape operation.
MD5	Message Digest 5. A checksum algorithm described in RFC (Request For Comments) 1321, a proposal for an Internet standard.
mirror	A volume that contains a read-only copy of data in the active file system of another volume, usually on another storage system.
NDMP	Network Data Management Protocol. A protocol that allows Network Appliance storage systems to communicate with backup applications, and provides capabilities for controlling the robotics of multiple tape backup devices.
NFS	Network File System. A protocol for networking UNIX-based computers.
NIS	Network Information Service, formerly called Yellow Pages. An administrative database for networks.
NVFAIL	Software that warns you of compromised database validity and automatically renames the database so that it does not restart automatically.
NVRAM	Nonvolatile RAM in the storage system, used for logging incoming write data and NFS requests. Improves system performance and prevents loss of data in case of a storage system or power failure.
Open Systems platform	A non-storage system system, such as a server running Solaris, HP-UX, Windows NT, or Windows 2000, whose data can be backed up to a SnapVault secondary storage system.
Open Systems SnapVault agent	A software module that can be installed on a non-storage system system and enables that system to back up its data to a SnapVault secondary storage system.

plex	A physical copy of a file system. An unmirrored volume has one plex; a mirrored volume has two identical plexes.
primary storage system	A system whose data is to be backed up by SnapVault.
quota	A limit placed on a file system that restricts disk space usage by files with a given UID or GID.
qtree	A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.
RAID	Redundant Array of Independent Disks. A technique that protects against disk failure by computing parity information based on the contents of all the disks in the array. NetApp storage systems use RAID Level 4, which stores all parity information on a single disk.
Remote Shell	A program that enables a user on one system to execute a program on another system. Remote Shell connections are usually not interactive.
remote tape device	A program-based functionality associated with a tape drive that is not directly attached to a storage system that is performing a tape operation.
root volume	The volume that contains information that controls the entire storage system, usually in the /etc/rc file.
secondary storage system	A storage system or NearStore® system to which data is backed up by SnapVault® software.

share	A directory or directory structure on the storage system that has been made available to network users and can be mapped to a drive letter on a CIFS client.
SnapMirror®	Software that performs automated file system replication of a volume onto the same or a separate disk or storage system.
SnapRestore®	Software that restores an entire volume to the state recorded in a previously taken Snapshot copy.
Snapshot copy	An online, read-only copy of the entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshot copies enable users to restore files and enable administrators to back up the storage system to tape while the storage system is in use.
Snapshot reserve	The portion of a volume's disk space that is reserved for Snapshot copies.
source storage system	The storage system from which you are replicating data.
subtree	A directory in a volume or qtree.
tape block	1,024 bytes of data.
tape device	A specific functionality of a physical tape drive that you create by specifying information in the tape device name when you install a tape drive or tape stacker.
tape file	Data on a tape delimited by file marks.
tape library	Hardware that can access tape cartridges randomly.

tape medium changer	Hardware that can load and unload tape cartridges.
tape stacker	Hardware that can access tape cartridges from a stack.
vFiler™	A virtual filer created using the MultiStore® software that allows you to partition the storage and networking resources of a single storage system so that it appears as multiple storage systems on the network.
volume	A file system.
volume copy	A way of copying both data in the active file system and data in Snapshot copies from one volume to another.

Index

Symbols

/dev/mcn file 146
/etc/dumpdates file 77
/etc/log/backup file 184
/etc/tape_config directory, contents 34
/etc/tape_config files 34

A

ACLs (Access Control Lists)
 including in tape backups 50
 omitting from tape backup 73
 omitting from tape restores 128
aliases, tape
 on multiple filers 21
 purposes for 16
aliasing commands. *See* storage (aliasing)
commands
audience, intended for this book xi
automatic confirmations, defined 127

B

backup terminate command (terminates a dump) 98
backups
 comparison of methods 47
 copying from tape with restore command 112
 creating snapshot_for_backup file for 50
 Data ONTAP process for 50
 defined 46
 estimating size of 53
 event logging for 183
 excluding specified files 92
 file information backed up in 49
 full 47
 incremental 47
 disadvantages of 60
 restoring incremental 131
 specifying with dumpdates file 77
 levels defined 63
 metadata 49
 methods Data ONTAP supports 47

nonconsecutive, contents of 64
parallel 56
restarting 97
restoring data from tapes 103
restoring files from multiple tapes
 on single-tape drive, example 138
 on tape stacker, example 139
 on two single-tape drives, example 138
restoring files from remote filer, example 137
restoring files from tape drive on a Solaris
 system, example 137
restoring from tape
 displaying table of contents (files) 116
 displaying table of contents (qtrees) 116
 ignoring inode limitations 125
 omitting ACLs 128
 omitting qtree information 129
 preparing for other than original path 107
 restoring entire filer to different location,
 example 134
 restoring entire filer to original location,
 example 134
 restoring entire filer, procedure 132
 restoring individual files, example 136
 running a test restore 130
 specifying a full restore 114
 specifying a restore destination 122
 specifying a resumed restore, procedure
 117
 specifying a resumed restore, rules 117
 specifying a tape file 121
 specifying automatic confirmations 127
 specifying individual files to restore 113
 specifying restore status 124
 specifying tape devices 119
 when to restore from tape 100
types of 50
using df to estimate size of 53
using dump command. *See* backups to tape
 (dump command)
why necessary 2
Windows NT ACL information, including 50

- backups to tape (dump command)
 - appending backups on one tape 79
 - benefits of entering at console 55
 - benefits of using Remote Shell 55
 - blocking factor in 75
 - data format 49
 - disadvantages 4
 - estimating tapes required for 54
 - example 58
 - excluding data in qtrees 70
 - excluding files 71
 - file information backed up 49
 - from existing snapshots 51
 - full backup excluding specified files, example 92
 - full backup from a snapshot, example 88
 - full backup of a list of files or directories, example 89
 - full backup specifying blocking factor, example 91
 - full backup to local tape
 - example 81
 - minus qtrees, example 82
 - with record, example 81
 - full backup to multiple tapes on single-tape drives, example 86
 - full backup to multiple tapes on tape stacker, example 86
 - full backup to remote Solaris system, example 84
 - full backup to remote tape, example 84
 - full backup to Solaris system with tape file size limit, example 85
 - full backup without ACLs, example 90
 - incremental backup to local tape, example 83
 - insufficient tapes for 54
 - labeling backup tapes 61
 - leaving volumes online for backup 50
 - maximum tape blocks per tape file 76
 - minimizing tapes used for 61
 - number of simultaneous dumps possible 56
 - objects backed up 49
 - options for 58
 - order of tape devices specified by 66
 - organizing efficiently 60

- preparing for 52
 - reasons for 46
 - recommendations for 60
 - rules for 57
 - rules for excluding files from 71
 - simultaneous dumps 51
 - specifying backup level 63
 - specifying backup names 72
 - specifying dump path 68
 - specifying files and directories 69
 - specifying local tape device names of 66
 - specifying remote tape device names of 67
 - specifying tape blocks per tape file 76
 - syntax 57
 - unattended 54
 - updating dumpdates file 78
 - use of snapshots 50
 - verifying 93
 - where to enter command 55
 - Windows NT ACL information, omitting 73
- blocking factor
 - specifying when backing up data 75
 - specifying when restoring data 123

C

- CIFS attributes, not backed up by dump command 49
- commands
 - See also* dump command
 - See also* NDMP commands
 - See also* restore command
 - See also* storage (aliasing) commands
 - See also* tape drive configuration commands
- cpio 47
- df (estimates amount of data) 53
- maxfiles (sets maximum number of files per volume) 125
- mt -bsf (moves tape backward, skips files) 31
- mt -eom (positions tape to end of data) 30
- mt -fsf (moves tape forward, skipping files) 31
- mt -offline (rewinds tape, unloads tape media) 32
- mt -rewind (rewinds tape) 32
- mt -status (shows information about tape)

- devices and their drives) 33
 - options `backup.log.enable` (turns event logging on or off) 190
 - storage alias (assigns tape alias) 20
 - storage alias (displays tape aliases) 19
 - storage show (displays tape drive information) 19
 - storage unalias (removes tape alias) 21
- compression type
 - displaying 25
 - specifying for tape devices 10
 - specifying in restores from tape 119
- confirmations, automatic
 - defined 127
 - specifying 127
- console, using `dump` command through 55
- conventions
 - command xi
 - formatting xi
 - keyboard xii
- `cpio` command, using for backup 47

D

- DAR functionality 161
- data
 - backups, defined 46
 - protection methods
 - defining 2
 - need for 2
- Data Management Application (DMA) 142
- devices, tape. *See* tape devices
- `df` command, finding size of data for backups 53
- Direct Access Restore (DAR) functionality 161
- DMA. *See* Data Management Application 142
- drive number, specifying for tape devices 9
- `dump` command
 - appending backups on one tape 79
 - backup levels, defined 63
 - CIFS attributes, not backed up by 49
 - deleting a restartable dump 98
 - example 58
 - excluding data in qtrees from backups 70

- excluding files from backups 71
- full backup excluding specified files, example 92
- full backup from a snapshot, example 88
- full backup of a list of files or directories, example 89
- full backup specifying blocking factor, example 91
- full backup to local tape minus qtrees, example 82
- full backup to local tape with record, example 81
- full backup to local tape, example 81
- full backup to multiple tapes on single-tape drives, example 86
- full backup to multiple tapes on tape stacker, example 86
- full backup to remote Solaris system with tape file size limit, example 85
- full backup to remote Solaris system, example 84
- full backup to remote tape, example 84
- full backup without ACLs, example 90
- how it works 50
- incremental backup to local tape, example 83
- leaving volumes online for 50
- maximum tape blocks per tape file 76
- objects backed up 49
- options for 58
- order of tape devices specified by 66
- permissions limitations 49
- preparing for use 52
- restarting 97
- rules for 57
- rules for excluding files from backups 71
- simultaneous dumps 51
- specifying backup level 63
- specifying backup names 72
- specifying dump path 68
- specifying files and directories 69
- specifying local tape device names 66
- specifying remote tape device names of 67
- specifying tape blocks per tape file 76
- syntax 57
- updating `dumpdates` file 78

- use of snapshots 50
- using for backups
 - labeling backup tapes 61
 - minimizing tapes used for 61
 - organizing efficiently 60
- verifying 93
- where to enter 55
- Windows NT ACL information, including 50
- Windows NT ACL information, omitting 73
- with blocking factor 75
- dump commands
 - backup terminate (terminates a dump) 98
 - dump restart (restarts interrupted backups) 97
- dump event log, example 187
- dump events 186
- dump restart command (restarts interrupted backup) 97
- dumpdates file
 - example 77
 - purpose 77
 - rules for 77
 - updating 78

E

- emulation of qualified tape drives 38
- environment variables
 - ACL_START 163
 - BASE_DATE 163
 - DATA_BLOCK_SIZE 163
 - DEBUG 164
 - DIRECT 164
 - DMP_NAME 164
 - DUMP_DATE 164
 - ENHANCED_DAR_ENABLED 165
 - EXCLUDE 166
 - EXTRACT 166
 - EXTRACT_ACL 166
 - FILESYSTEM 166
 - FORCE 167
 - HIST 167
 - IGNORE_CTIME 167
 - IGNORE_QTREES 168
 - LEVEL 168
 - LIST 168

- LIST_QTREES 168
- MULTI_SUBTREE_NAMES 168
- NDMP_UNICODE_FH 169
- NDMP_VERSION 169
- NO_ACLS 169
- NON_QUOTA_TREE 169
- NOWRITE 170
- REPLICATE 170
- RESTART 170
 - specifying 162
 - supported 162
- UPDATE 170
- VERBOSE 170
 - what they are 162
- event log files
 - effect of takeover mode on 184
 - rotation 184
 - viewing, reasons for 184
- event log messages
 - aborted restore, example 189
 - dump events 186
 - dump, example 187
 - format 185
 - identifier field 185
 - restore events 188
 - restore, example 188
 - start and stop logging events 186
 - timestamp field 185
 - type field 185
- event logging
 - defined 184
 - enabling or disabling 190
 - for backup and restore 183
- examples
 - dump command 58
 - dumpdates file 77
 - event log
 - aborted restore 189
 - restore 188
 - full backup
 - excluding specified files 92
 - from a snapshot 88
 - of a list of files or directories 89
 - specifying blocking factor 91
 - to local tape 81

- to local tape minus qtrees 82
- to local tape with record 81
- to multiple tapes on single-tape drives 86
- to multiple tapes on tape stacker 86
- to remote Solaris system 84
- to remote tape 84
- to Solaris system with tape file size limit 85
- without ACLs 90

- incremental backup to local tape 83
- restore command. *See* restore command
- restoring files from multiple tapes on single-tape drive 138
- tape configuration file 35
- tape device names 10

F

filers

- adding Fibre-Channel-attached drives
 - dynamically 23
- displaying compression capacity of tape drives 25
- displaying information about tape drive connections to 26
- displaying SCSI IDs of tape drives 25
- displaying tape drive information 25
- restoring all volumes on from backup 132
- showing which tape drives supported 14

files

- backing up using dump command 49
- displaying table of contents of for restoration 116
- dumpdates 77
- event log 184
- excluding from dump command 71
- restoring all from tape 114
- restoring specified files from tape 113
- tape qualification 34
- using Remote Shell for backups 69

- full backups 47

- full restores 114

H

- hosts, remote 8

I

- increment chains, of backups 64

- incremental backups, nonconsecutive, contents of 64

- inodes

- defined 125

- displaying for restored files 112

- ignoring limits when restoring files 125

L

- levels of backups 63

- log files

- for backup and restore events 184

- NDMP 175

- LUN (logical unit number) 17

M

- maxfiles command (sets maximum number of files per volume) 125

- metadata backups 49

- methods of backup 47

- mt command (controls tape movement on tape drives) 29

- mt command syntax 29

N

NDMP

- advantages of 144

- copying with local tool (ndmpcopy) 178

- debug log file, displaying 175

- debug messages 174

- defined 142

- disabling preferred network interface 160

- disabling service (ndmpd off) 150

- enabling or disabling service 150

- enabling service (ndmpd on) 150

- killing sessions (ndmpd kill command) 177

- log files 175

- preparing a filer for basic management 148

session information

- displaying (ndmpd status command) 171
- displaying detailed (ndmpd probe) 172
- types 171

setting preferred network interface 159

showing max version supported (ndmpd version command) 157

tape devices used for 146

terminating session (ndmpd kill) 177

Unicode format, file names in 144

using with tape libraries 146

versions of 158

NDMP commands

ndmp on 150

ndmcopy (uses local copy tool) 178

ndmpd debug (outputs debug log file) 175

ndmpd kill (terminates NDMP session) 177

ndmpd off (disables service) 150

ndmpd on (enables service) 150

ndmpd probe (displays detailed status) 172

ndmpd status (displays status) 171

ndmpd version (shows max version supported) 157

ndmp.preferred_interface disable option (disables preferred network interface) 160

ndmp.preferred_interface option (sets preferred network option) 159

ndmcopy command 178

ndmpd debug command (outputs debug log file) 175

ndmpd kill command (terminates NDMP session) 177

ndmpd on/off command 150

ndmpd probe command (displays detailed status) 172

ndmpd status command (displays status) 171

ndmpd version command (shows max version supported) 157

ndmpd.offset_map.enable option 161

NFS clients, using cpio command with 47

nonqualified tape drives

- defined 38

- emulating qualified tape drives 39

O

offset map 144

options

- ndmpd.offset_map.enable 161

- ndmpd.preferred_interface (sets preferred network interface) 159

- ndmpd.preferred_interface disable (disables preferred network interface) 160

- tape.persistent_reservations (associates tape device with one filer) 42

options backup.log.enable command (turns event logging on or off) 190

P

parallel backups 56

permissions limitations, in dump and restore 49

persistent tape reservations 42

physical path names (PPNs)

- definition 16

- example 17

- format 17

Q

qtrees

- backing up with dump command 50

- displaying a table of contents for restoring from tape 116

- excluding data from backup 70

- omitting data from dump command 70

- omitting information from tape restores 129

qualified tape drives, defined 34

R

remote hosts 8

Remote Shell

- using automatic confirmation when restoring from tape 127

- using dump command through, benefits of 55
- using to display table of contents for restores from tape 115

- using with specified files to dump 69

reservations, tape 42

- restartable backups
 - deleting 98
 - finding when needed 94, 96
 - qualifications 96
- restore command
 - Data ONTAP version compatibility 101
 - data that cannot be restored 101
 - disk space required for 105
 - displaying a table of contents (list of files) 116
 - displaying a table of contents (list of qtrees) 116
 - for other than original path 107
 - how it works 103
 - how to run 112
 - ignoring inode limitations 125
 - information required for using 105
 - omitting ACLs 128
 - omitting qtree information 129
 - options 110
 - permissions limitations 49
 - preparing for 104
 - restoring entire filer
 - procedure 132
 - to different location, example 134
 - to original location, example 134
 - restoring files
 - from multiple tapes on single-tape drive, example 138
 - from multiple tapes on tape stacker, example 139
 - from multiple tapes on two single-tape drives, example 138
 - from remote filer, example 137
 - from tape drive on a Solaris system, example 137
 - restoring incremental backups 131
 - restoring individual files, example 136
 - rules 109
 - running a test restore 130
 - specifying a full restore 114
 - specifying a restore destination 122
 - specifying a resumed restore 117
 - specifying a tape file 121
 - specifying automatic confirmations 127
 - specifying individual files to restore 113
 - specifying restore status 124
 - specifying tape devices 119
 - syntax 109
 - tapes required for 105
 - types of restores 109, 110
 - using Remote Shell for displaying table of contents 115
 - using with Remote Shell 106
 - when to use 100
 - Windows NT ACLs 110
- restore event log
 - aborted restore, example 189
 - successful restore, example 188
- restore events 188
- restores
 - displaying a table of contents (files) 115
 - displaying a table of contents (qtrees) 116
 - event logging for 183
 - ignoring inode limitations 125
 - incremental 131
 - omitting ACLs 128
 - omitting qtree information 129
 - preparing for 104
 - restoring entire filer
 - procedure 132
 - to different location, example 134
 - to original location, example 134
 - restoring files
 - from multiple tapes on tape stacker, example 139
 - from multiple tapes on two single-tape drives, example 138
 - from remote filer, example 137
 - from tape drive on a Solaris system, example 137
 - restoring individual files, example 136
 - running a test restore 130
 - specifying a full restore 114
 - specifying a restore destination 122
 - specifying a resumed restore, procedure 117
 - specifying a resumed restore, rules 117
 - specifying a tape file 121
 - specifying automatic confirmations 127
 - specifying individual files to restore 113
 - specifying restore status 124

- specifying tape devices 119
- types of 110
- rewind type, specifying for tape devices 9
- rmt protocol 48
- rules
 - dump command 57
 - for dumpdates file 77
 - for excluding files from dump command 71
 - for restore command 109
 - for specifying a resumed restore 117

S

- SCSI ID, of tape drives, displaying 25
- single-tape tape drives 86
- snapshots
 - snapshot_for_backup file 50
- Solaris ufsrestore command 49
- special messages xiii
- stackers, tape, backups to 86
- storage (aliasing) commands
 - storage alias (assigns tape alias) 20
 - storage alias (displays tape aliases) 19
 - storage show (displays tape drive information) 19
 - storage unalias (removes tape alias) 21
- storage alias command (assigns tape alias) 20
- storage alias command (displays tape aliases) 19
- storage show command (displays tape drive information) 19
- storage show tape supported command (shows acceptable tape drives) 14
- storage stats tape command (displays tape statistics) 12
- storage stats tape zero command (restarts tape statistics display) 13
- storage unalias command (removes tape alias) 21
- subtrees, defined 68
- sysconfig -m command (shows information on tape media changers) 26
- sysconfig -t command (shows tape drive information) 25
- sysconfig -v command (shows tape drive connections to filer) 26

T

- tape aliases
 - definition 16
 - how to display 18
 - on multiple filers 21
 - physical path names (PPNs) 16
- tape backups. *See* backups to tape (dump command)
- tape blocks, specifying the size for dump 75
- tape configuration file, example 35
- tape devices
 - defined 7
 - example of name listing 10
 - examples of names 10
 - format of names 8
 - local, defined 8
 - on remote Solaris systems 8
 - remote, defined 8
 - reserving 42
 - specifying compression type of 10
 - specifying for a resumed restore of 119
 - specifying rewind number of 9
 - specifying rewind type of 9
 - specifying when restoring data 111
 - tape libraries, qualified tape drives in 24
 - tape stackers, backing up to multiple tapes on 86
 - use for 8
- tape drive configuration commands
 - sysconfig -m (shows information about tape media changers) 26
 - sysconfig -t (shows tape drive information) 25
 - sysconfig -t (shows what tape drives are on a filer) 24
 - sysconfig -v (shows tape drive connections to filer) 26
- tape drive slot on the filer, displaying 25
- tape drives
 - adding Fibre-Channel-attached drives
 - dynamically 23
 - connections to filer, displaying information about 26
 - controlling tape movement 29
 - displaying compression capacity 25
 - displaying SCSI IDs 25

- displaying which ones are on a filer 24
- displaying which supported 14
- how filer qualifies 36
- in tape stackers, listing qualified 24
- local, specifying backups on 81
- moving tape backward (mt -bsf) 31
- moving tape forward, skipping files (mt -fsf) 31
- multiple tapes, using for backup on a tape stacker 86
- nonqualified
 - defined 38
 - displaying information 38
 - emulating qualified tape drives 39
 - getting information for emulation 39
 - information needed for emulation 38
 - using 24
- positioning tape to end of data (mt -eom) 30
- qualified
 - defined 34
 - list of 24
- qualifying nonqualified drives 34
- remote, specifying 8
- reserving 42
- rewinding tape (mt -rewind) 32
- showing status (mt -status) 33
- single-tape, using multiple tapes on 86
- skipping files (mt -bsf) 31
- tape media changers, displaying information about 26
- unloading tape after rewind (mt -offline) 32
- usable by NDMP 146
- tape files
 - compression types 28
 - definition of 28
 - maximum size for dump 76
 - specifying local tape device names of 66
 - specifying remote tape device names of 67
 - specifying size for dump of 76
- tape libraries
 - defined 143
 - qualified tape drives in 24
 - recognized by filers 146
 - showing names assigned to 146
 - using with NDMP 146
- tape media changers, displaying information about 26
- tape qualification files
 - creating 34
 - error messages related to, example 36
 - format 34
 - how used by filer 36
 - list of SCSI command timeout values 34
 - requirements 34
- tape restores
 - displaying a table of contents (files) 115
 - displaying a table of contents (qtrees) 116
 - ignoring inode limitations 125
 - incremental 131
 - omitting ACLs 128
 - omitting qtree information 129
 - preparing for 104
 - restoring entire filer
 - procedure 132
 - to different location, example 134
 - to original location, example 134
 - restoring files
 - from multiple tapes on tape stacker, example 139
 - from multiple tapes on two single-tape drives, example 138
 - from remote filer, example 137
 - from tape drive on a Solaris system, example 137
 - restoring files from multiple tapes on single-tape drive, example 138
 - restoring individual files, example 136
 - running a test restore 130
 - specifying a full restore 114
 - specifying a restore destination 122
 - specifying a resumed restore, rules 117
 - specifying a tape file 121
 - specifying automatic confirmations 127
 - specifying individual files to restore 113
 - specifying restore status 124
 - specifying tape devices 119
 - types of 110
- tape stackers
 - backing up to multiple tapes on 86
 - displaying information about 26

- tape statistics
 - displaying 12
 - restarting display 13
- tape_alias file 21
- tapes
 - backing data up to 45
 - displaying statistics 12
 - estimating number needed for backups 54
 - insufficient number of for backups 54
 - reasons to use for backups 46
 - required for restore 105
 - restarting tape statistics display 13
 - restoring data from 103
- terminology xi

V

- variables, environment
 - about 162
 - ACL_START 163
 - BASE_DATE 163
 - DATA_BLOCK_SIZE 163
 - DEBUG 164
 - DIRECT 164
 - DMP_NAME 164
 - DUMP_DATE 164
 - ENHANCED_DAR_ENABLED 165
 - EXCLUDE 166
 - EXTRACT 166

- EXTRACT_ACL 166
- FILESYSTEM 166
- FORCE 167
- HIST 167
- IGNORE_CTIME 167
- IGNORE_QTREES 168
- LEVEL 168
- LIST 168
- LIST_QTREES 168
- MULTI_SUBTREE_NAMES 168
- NDMP_UNICODE_FH 169
- NDMP_VERSION 169
- NO_ACLS 169
- NON_QUOTA_TREE 169
- NOWRITE 170
- REPLICATE 170
- RESTART 170
- specifying 162
- supported 162
- UPDATE 170
- VERBOSE 170
 - what they are 162
- volumes, backing up with dump command 50

W

- Windows NT ACLs, omitting from restores from tape 128
- worldwide names (WWNs) 18