

Data ONTAP® 7.3

Data Protection

Tape Backup and Recovery Guide

NetApp, 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 210-04762_A0
Updated for Data ONTAP 7.3.3 on 15 January 2010

Contents

Copyright information	9
Trademark information	11
About this guide	13
Audience	13
Accessing Data ONTAP man pages	13
Terminology	14
Where to enter commands	15
Keyboard and formatting conventions	15
Special messages	16
How to send your comments	17
Data protection using tape	19
Advantages and disadvantages of tape backup	19
Tape drive management	21
What tape devices are	21
Tape device name format	22
Supported number of simultaneous tape devices	24
Displaying tape device statistics	24
Displaying supported tape devices	25
What assigning tape aliases is	26
What physical path names are	27
What worldwide names are	28
Displaying existing aliases of tape drives	29
Displaying information about tape drives or libraries	29
Assigning tape aliases	30
Removing tape aliases	31
Propagating tape aliases to multiple storage systems	31
UNIX shell scripts for propagating tape aliases	32
How to add Fibre Channel-attached tape drives and libraries	32
How to display tape drive and tape library information	33
Displaying information about tape drives	33
Displaying information about tape medium changers	34

Displaying information about tape drive connections to the storage system	34
Controlling tape drives	35
Moving a tape to the end of data	36
Moving forward to a file	37
Moving backward to the beginning of a file	37
Rewinding a tape	38
Taking a tape drive offline	38
Displaying status information	39
Qualified tape drives	40
Format of the tape configuration file	40
How the storage system qualifies a new tape drive dynamically	42
How to use a nonqualified tape drive	42
Displaying information about nonqualified tape drives	43
Tape drive information required for emulation	43
Emulating a qualified tape drive	44
What tape reservations are	45
Enabling tape reservations	45
Disabling tape reservations	46
NDMP management	47
What the advantages of NDMP are	47
What NDMP security is	48
Specifying NDMP access by host or interface	49
Specifying the NDMP authentication type	49
Enabling or disabling NDMP connection logging	50
Specifying the NDMP password length	51
Generating an NDMP-specific password for non-root administrators	51
How to manage NDMP	52
Enabling and disabling NDMP services	52
Specifying a preferred network interface	53
Turning off a data connection specification	53
Displaying the general status information about NDMP sessions	54
Displaying detailed NDMP session information	54
Optimizing NDMP communication performance	55
Terminating an NDMP session	56
What NDMP debug messages are	56

Enabling the NDMP debug log messages	57
Displaying the NDMP debug log level	57
Changing NDMP debug log messages	58
Displaying an NDMP session log file	58
Why you need to specify the NDMP version	58
Displaying the NDMP version	59
Specifying the NDMP version	59
NDMP extensions supported by Data ONTAP	59
Tape backup using NDMP services	60
Common NDMP tape backup topologies	60
Considerations when using NDMP	61
Tape devices and configurations you can use with the storage system	61
Preparing for basic NDMP backup application management	62
What environment variables do	64
Data backup using the dump engine	65
How a dump backup works	66
What the dump engine backs up	66
What increment chains are	67
How to specify tape devices for the backup	69
What the /etc/dumpdates file is	69
What the blocking factor is	70
How to use the dump backup	71
How to minimize backup time and data loss	71
How to decrease tape backup time	72
How to minimize the number of tape drives	72
What to label on the backup tapes	72
Considerations before using the dump backup	73
Determining the amount of backup data	73
Estimating the number of tapes for the backup	73
When to restart a dump backup	74
How a dump restore works	75
What the dump engine restores	75
Considerations before restoring data	76
How to prepare the destination for a dump restore	77
How online migration affects tape backup	77
How to perform a dump backup and restore using NDMP services	78

Environment variables supported for dump	78
Enabling or disabling enhanced DAR functionality	87
What the ndmpcopy command does	88
Displaying file history statistics	92
How to perform a dump backup using the CLI	93
What the dump command syntax is	94
Where to enter the dump command	96
Specifying the backup level	97
Improving incremental dump performance	98
Updating the /etc/dumpdates file	98
Specifying a local tape device	99
Specifying a tape device on a remote storage system	99
Specifying the dump path	101
Specifying a list of files for backup	102
Backing up all data that is not in a qtree	103
Excluding specified files and directories	104
Omitting ACLs from a backup	105
Specifying a name for a backup	106
Specifying a blocking factor	106
Specifying the tape file size	107
Appending backups to tapes	108
Verifying the files backed up by a dump command backup	108
Checking the status of a dump backup	109
Finding out whether a backup has to be restarted	111
How to get details about a specific backup	112
Restarting a dump command backup	113
Deleting restartable dump command backups	114
How to perform a dump restore using the CLI	114
Restore command syntax	115
What restore types are	115
What modifiers are	116
Where to enter the restore command	117
Executing a restore command	117
Restoring incremental backups	118
Restoring each volume backed up as separate subtrees or qtrees	118
Restoring individual files and directories	119

Specifying a full restore	119
What a table-of-contents restore is	120
Specifying a resume restore	121
Specifying tape devices in the restore command	122
Specifying a single tape file on a multifile tape	123
Specifying the restore destination	124
Specifying the blocking factor during restore	124
Displaying detailed status output	125
Ignoring inode limitations	126
Specifying automatic confirmations	127
Specifying no ACLs to be restored	127
Specifying not to restore qtree information	128
Specifying a test restore	129
Restore examples: Restoring using a remote tape drive	129
Restore examples: Multiple tape restores	131
What event logging is	133
What the dump and restore event log message format is	133
What logging events are	134
What dump events are	134
What restore events are	136
Enabling or disabling event logging	137
Index	139

Copyright information

Copyright © 1994–2010 NetApp, 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.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP 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.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. 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 NetApp.

The product described in this manual may be protected by one or more U.S.A. 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

All applicable trademark attribution is listed here.

NetApp, the Network Appliance logo, the bolt design, NetApp-the Network Appliance Company, Cryptainer, Cryptoshred, DataFabric, DataFort, Data ONTAP, Decru, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, NOW NetApp on the Web, SANscreen, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, StoreVault, SyncMirror, Topio, VFM, and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The NetApp arch logo; the StoreVault logo; ApplianceWatch; BareMetal; Camera-to-Viewer; ComplianceClock; ComplianceJournal; ContentDirector; ContentFabric; Data Motion; EdgeFiler; FlexShare; FPolicy; Go Further, Faster; HyperSAN; InfoFabric; Lifetime Key Management, LockVault; NOW; ONTAPI; OpenKey, RAID-DP; ReplicatorX; RoboCache; RoboFiler; SecureAdmin; SecureView; Serving Data by Design; Shadow Tape; SharedStorage; Simplicore; Simulate ONTAP; Smart SAN; SnapCache; SnapDirector; SnapFilter; SnapMigrator; SnapSuite; SohoFiler; SpinMirror; SpinRestore; SpinShot; SpinStor; vFiler; VFM Virtual File Manager; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, Inc. in the U.S.A.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. 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 U.S.A. and/or other countries.

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetCache is certified RealSystem compatible.

About this guide

You can use your product more effectively when you understand this document's intended audience and the conventions that this document uses to present information.

This guide describes how to protect, back up, restore, and copy data between storage systems that run Data ONTAP software.

Next topics

Audience on page 13

Accessing Data ONTAP man pages on page 13

Terminology on page 14

Where to enter commands on page 15

Keyboard and formatting conventions on page 15

Special messages on page 16

How to send your comments on page 17

Audience

This document is written with certain assumptions about your technical knowledge and experience.

This guide is for system administrators who are familiar with operating systems that run on the storage system clients, such as UNIX, Linux, Solaris, Windows NT, Windows 2000, Windows XP, and Windows Vista.

It also assumes that you are familiar with how to configure the storage system 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 storage system.

Accessing Data ONTAP man pages

You can use the Data ONTAP manual (man) pages to access technical information.

About this task

Data ONTAP manual pages are available for the following types of information. They are grouped into sections according to standard UNIX naming conventions.

Types of information	Man page section
Commands	1
Special files	4
File formats and conventions	5
System management and services	8

Step

1. View man pages in the following ways:

- Enter the following command at the storage system command line:
`man command_or_file_name`
- Click the manual pages button on the main Data ONTAP navigational page in the FilerView user interface.
- Use the *Commands: Manual Page Reference*, Volumes 1 and 2 (which can be downloaded or ordered through the NOW site).

Note: All Data ONTAP man pages are stored on the storage system in files whose names are prefixed with the string "na_" to distinguish them from client man pages. The prefixed names are used to distinguish storage system man pages from other man pages and sometimes appear in the NAME field of the man page, but the prefixes are not part of the command, file, or services.

Terminology

To understand the concepts in this document, you might need to know how certain terms are used.

Storage terms

array LUN	Refers to storage that third-party storage arrays provide to storage systems running Data ONTAP software. One array LUN is the equivalent of one disk on a native disk shelf.
LUN (logical unit number)	Refers to a logical unit of storage identified by a number.
native disk	Refers to a disk that is sold as local storage for storage systems that run Data ONTAP software.
native disk shelf	Refers to a disk shelf that is sold as local storage for storage systems that run Data ONTAP software.

storage controller	Refers to the component of a storage system that runs the Data ONTAP operating system and controls its disk subsystem. Storage controllers are also sometimes called <i>controllers</i> , <i>storage appliances</i> , <i>appliances</i> , <i>storage engines</i> , <i>heads</i> , <i>CPU modules</i> , or <i>controller modules</i> .
storage system	Refers to the hardware device running Data ONTAP that receives data from and sends data to native disk shelves, third-party storage, or both. Storage systems that run Data ONTAP are sometimes referred to as <i>filers</i> , <i>appliances</i> , <i>storage appliances</i> , <i>V-Series systems</i> , or <i>systems</i> .
third-party storage	Refers to the back-end storage arrays, such as IBM, Hitachi Data Systems, and HP, that provide storage for storage systems running Data ONTAP.

Where to enter commands

You can use your product more effectively when you understand how this document uses command conventions to present information.

You can perform common administrator tasks in one or more of the following ways:

- You can enter commands either at the system console or from any client computer that can obtain access to the storage system using a Telnet or Secure Shell (SSH) session.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.
- You can enter commands either at the switch console or from any client that can obtain access to the switch using a Telnet session.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.

Keyboard and formatting conventions

You can use your product more effectively when you understand how this document uses keyboard and formatting conventions to present information.

Keyboard conventions

Convention	What it means
The NOW site	Refers to <i>NetApp On the Web</i> at http://now.netapp.com/ .

Convention	What it means
<i>Enter, enter</i>	<ul style="list-style-type: none"> Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. Used to mean pressing one or more keys on the keyboard and then pressing the Enter key, or clicking in a field in a graphical interface and then typing information into the field.
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
type	Used to mean pressing one or more keys on the keyboard.

Formatting conventions

Convention	What it means
<i>Italic font</i>	<ul style="list-style-type: none"> Words or characters that require special attention. Placeholders for information that you must supply. For example, if the guide says to enter the <code>arp -d hostname</code> command, you enter the characters "arp -d" followed by the actual name of the host. Book titles in cross-references.
Monospaced font	<ul style="list-style-type: none"> Command names, option names, keywords, and daemon names. Information displayed on the system console or other computer monitors. Contents of files. File, path, and directory names.
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.

Special messages

This document might contain the following types of messages to alert you to conditions that you need to be aware of.

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

Attention: An attention notice contains instructions that you must follow to avoid a system crash, loss of data, or damage to the equipment.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by e-mail to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the name of your product and the applicable operating system. For example, *FAS6070—Data ONTAP 7.3*, or *Host Utilities—Solaris*, or *Operations Manager 3.8—Windows*.

Data protection using tape

You use tape backup and recovery to create tape archives and to retrieve data from tape archives.

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 copy 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).

Advantages and disadvantages of tape backup

Data backed up to tape requires fewer resources to maintain. However, restoring data from tape might take a long time.

Following are the advantages of tape backup over online storage:

- Tape backups require fewer resources to maintain.
- You can place the archives in a more secure place than you can place a storage system.
- You can recover data from any release of Data ONTAP.

Following are the disadvantages of tape archives over online storage:

- Restoring data from tape takes a long time.
- Finding a particular file or directory on tape is time consuming.

Tape drive management

You need to manage tape drives when you back up data from the storage system to tape or when you restore data from tape to the storage system.

When you back up data to tape, the data is stored in tape files. File marks separate the tape files and the files have no names. You specify a tape file by its position on the tape. You write a tape file using a tape device. When you read the tape file, you must specify a device that has the same compression type that you used to write it.

Next topics

[*What tape devices are*](#) on page 21

[*Tape device name format*](#) on page 22

[*Supported number of simultaneous tape devices*](#) on page 24

[*Displaying tape device statistics*](#) on page 24

[*Displaying supported tape devices*](#) on page 25

[*What assigning tape aliases is*](#) on page 26

[*Displaying existing aliases of tape drives*](#) on page 29

[*Displaying information about tape drives or libraries*](#) on page 29

[*Assigning tape aliases*](#) on page 30

[*Removing tape aliases*](#) on page 31

[*Propagating tape aliases to multiple storage systems*](#) on page 31

[*How to add Fibre Channel-attached tape drives and libraries*](#) on page 32

[*How to display tape drive and tape library information*](#) on page 33

[*Controlling tape drives*](#) on page 35

[*Qualified tape drives*](#) on page 40

[*How to use a nonqualified tape drive*](#) on page 42

[*What tape reservations are*](#) on page 45

What tape devices are

A tape device is a representation of a tape drive. It is a specific combination of rewind type and compression capability of a tape drive.

A tape device is created for each combination of rewind type and compression capability. Therefore, a tape drive or tape library can have several tape devices associated with it. You must specify a tape device to move, write, or read tapes.

When you install a tape drive or tape library on a storage system, Data ONTAP creates tape devices associated with the tape drive or tape library.

Data ONTAP detects tape drives and tape libraries and assigns logical numbers and tape devices to them. Data ONTAP detects the Fibre Channel tape drives and libraries when they are connected to

the Fibre Channel interface ports. Data ONTAP detects these drives when their interfaces are enabled subsequently.

There are two types of tape devices:

- A local tape device on the storage system, which performs a tape operation
- A remote tape device on a storage system or Solaris machine that fulfills the following criteria:
 - 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 (remote magnetic tape) protocol (which is a bundled component of Data ONTAP)
 - Has a trust relationship with the storage system that is performing the tape operation

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

Tape device name format

Each tape device has an associated name that appears in a defined format. The format includes information about the type of device, its alias, and compression type.

The format of a tape device name is as follows:

[remote_host:]rewind_type st alias_number compression_type

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

rewind_type is the rewind type.

The following list describes the various rewind type values:

- r** Data ONTAP rewinds the tape after it finishes writing the tape file.
- nr** Data ONTAP does not rewind the tape after it finishes writing the tape file. Use this rewind type when you want to write multiple tape files on the same tape.
- ur** This is the unload/reload rewind type. When you use this rewind type, the tape library 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 in a tape library or is in a medium changer that is in the library 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 library 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.

`alias_number` is the alias that Data ONTAP assigns to the tape drive. When Data ONTAP detects a new tape drive, it assigns an alias to it. You can modify an alias using the `storage alias` command. An alias assigned by Data ONTAP or modified by the user persists through reboots.

`compression_type` is a drive-specific code for the density of data on the tape and the type of compression.

The following list describes the various values for `compression_type`:

- 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.

Attention: When using the `ur` device with the `dump` or `restore` command, ensure that you use tape libraries and that there are sufficient tapes in the library sequence. Otherwise, the tape drives involved terminate the command sequence or overwrite the same tape multiple times.

Example of a listing of tape devices

The following example shows the tape devices associated with HP Ultrium 2-SCSI:

```
Tape drive (fc202_6:2.126L1) HP      Ultrium 2-SCSI
rst0l - rewind device,          format is: HP (200GB)
nrst0l - no rewind device,      format is: HP (200GB)
urst0l - unload/reload device,  format is: HP (200GB)
rst0m - rewind device,          format is: HP (200GB)
nrst0m - no rewind device,      format is: HP (200GB)
urst0m - unload/reload device,  format is: HP (200GB)
rst0h - rewind device,          format is: HP (200GB)
nrst0h - no rewind device,      format is: HP (200GB)
urst0h - unload/reload device,  format is: HP (200GB)
rst0a - rewind device,          format is: HP (400GB w/comp)
nrst0a - no rewind device,      format is: HP (400GB w/comp)
urst0a - unload/reload device,  format is: HP (400GB w/comp)
```

The following list describes the abbreviations in the preceding example:

- GB**—Gigabytes; this is the capacity of the tape.

- w/comp—With compression; this shows the tape capacity with compression.

Related tasks

[Assigning tape aliases](#) on page 30

Supported number of simultaneous tape devices

Data ONTAP software supports a maximum of 64 simultaneous tape drive connections for each storage system in any mix of Fibre Channel or SCSI attachment.

Tape drives can be devices in tape libraries or stand-alone devices. Virtual Tape Libraries, such as NetApp VTL and their tape drives, are treated as actual tape drives; therefore, Data ONTAP supports a maximum of 64 simultaneous connections.

Note: Though a storage system can detect 64 tape drive connections, only 16 concurrent backup or restore sessions with local tapes are allowed.

Displaying tape device statistics

The tape device statistics help understand tape performance and check usage pattern. You reset the statistics reading and restart the process of displaying the statistics whenever you want.

Step

1. To display the statistics for a specified tape device, enter the following command:

```
storage stats tape tape_name
```

tape_name is the name of a tape device.

Example

```
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

Displaying supported tape devices

You can view a list of tape devices supported by a storage system using the `storage show tape supported` command. You can use a tape device only if it is listed in the output of this command.

Step

1. To display a list of the tape drives supported by the storage system, enter the following command:

`storage show tape supported [-v]`

The `-v` option gives you more detailed information about each tape drive.

Examples

```
filer1>storage show tape supported
Supported Tapes
```

```
-----
Exabyte 8500C 8mm
Exabyte 8505 8mm
Exabyte 8900 8mm
Exabyte 8500 8mm
Exabyte Mammoth-2 8mm
Digital DLT2000
Quantum DLT2000
Sun DLT2000
```

```
storage show tape supported -v
```

```
IBM ULTRIUM-TD1
Density Compression
Setting      Setting
-----
0x40         0x00      LTO Format 100 GB
0x40         0x00      LTO Format 100 GB
0x40         0x00      LTO Format 100 GB
0x40         0x01      LTO Format 200 GB comp
```

```
IBM 03590B
Density Compression
Setting      Setting
-----
0x29         0x00      B Format 10 GB
0x29         0x00      B Format 10 GB
0x29         0x00      B Format 10 GB
0x29         0xFF      B Format 20 GB comp
```

```
IBM 03590E
```

Density Setting	Compression Setting	
0x2A	0x00	E Format 20 GB
0x2A	0x00	E Format 20 GB
0x2A	0x00	E Format 20 GB
0x2A	0xFF	E Format 40 GB comp

IBM 03590H

Density Setting	Compression Setting	
0x2C	0x00	H Format 30 GB
0x2C	0x00	H Format 30 GB
0x2C	0x00	H Format 30 GB
0x2C	0xFF	H Format 60 GB comp

Certance Ultrium 2 - Dynamically Qualified

Density Setting	Compression Setting	
0x00	0x00	LTO-1 100GB
0x00	0x01	LTO-1 200GB comp
0x00	0x00	LTO-2 200GB
0x00	0x01	LTO-2 400GB comp

Certance Ultrium 3 - Dynamically Qualified

Density Setting	Compression Setting	
0x00	0x00	LTO-1(ro)/2 1/200GB
0x00	0x01	LTO-1(ro)/2 2/400GB comp
0x00	0x00	LTO-3 400GB
0x00	0x01	LTO-3 800GB comp

What assigning tape aliases is

Aliasing binds a tape or a medium changer device address, or a WWN, to a persistent, but modifiable alias name.

Aliasing simplifies the process of device identification. The following table describes how tape aliasing enables you to ensure that a tape drive (or tape library or medium changer) is always associated with a single alias name.

Scenario	Reassigning of the alias
When the system reboots	The tape drive is automatically reassigned its previous alias.

Scenario	Reassigning of the alias
When a tape device moves to another port	The alias can be adjusted to point to the new address.
When more than one system uses a particular tape device	The user can set the alias to be the same for all the systems.

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

Note: `st0` and `st00` are different logical names.

You can use tape aliases as parameters to the `dump` and `restore` commands.

Note: Logical names and WWNs are used only to access a device. After the device is accessed, it returns all error messages using the physical path name.

There are two types of names available for aliasing: PPNs and WWNs.

Next topics

[What physical path names are](#) on page 27

[What worldwide names are](#) on page 28

Related tasks

[Assigning tape aliases](#) on page 30

[Removing tape aliases](#) on page 31

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 the SCSI-2/3 adapter or switch (specific location) they are connected to, on the storage system. PPNs are also known as electrical names.

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: The LUN value is displayed only for tape and medium changer devices whose LUN values are not zero; that is, if the LUN value is zero the `lun` part of the PPN is not displayed.

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. Fibre Channel, SCSI tape drives and libraries, and disks have 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 `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 the switch MY_SWITCH, the new tape drive would be accessible using 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.

Accessing a tape drive or library using 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 can have different PPNs for the same device. Using the WWN in these cases eliminates any confusion.

Also, if you rename a switch or move a tape drive in the storage system, the WWN of the tape drive or library does not change. The scripts or backup programs do not need to change the name of the tape drive or library to which they are backing up.

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, the L## part of the string is not displayed.

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

Example of a dump command that uses the logical name of a tape drive

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

Example of a dump command that uses the worldwide name of a tape drive

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

Displaying existing aliases of tape drives

You can determine the existing aliases of tape drives using the `storage alias` command.

Step

1. To determine the existing aliases of tape drives, enter the following command:

```
storage alias
```

Example

```
filer1>storage alias
Alias                Mapping
-----
st0                  MY_SWITCH:5.3L3
st2                  MY_SWITCH:5.4L6
mc1                  2:4e3:38fe3f:758eab
mc348                MY_SWITCH:5.3L0
```

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 WWN 2:4e3:38fe3f:758eab.

Displaying information about tape drives or libraries

Information about tape drives and tape libraries helps you to assign tape aliases.

Step

1. To display information about tape drives and tape libraries (medium changers), enter the following command:

```
storage show {tape | mc} [{alias | PPN | WWN}]
```

alias is the logical name of the tape drive or medium changer.

PPN is the physical path name.

WWN is the worldwide name.

Examples

```
filer1>storage show tape
Tape Drive:      MY_SWITCH:5.3L4
Description:     Quantum DLT7000
Serial Number:   12345679
World Wide Name: WWN[2:333:4444444:555555]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
```

```
filer1>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
```

```
filer1>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

You can assign aliases to tape drives or medium changers using the `storage alias` command.

Step

1. To assign an alias to a tape drive or medium changer, enter the following command:

```
storage alias [alias {PPN | WWN}]
```

alias is the logical name of the tape drive or medium changer to which you want to add the alias.

PPN is the physical path name to which you want to assign the tape drive or medium changer.

WWN is the WWN to which you want to assign the tape drive or medium changer.

Examples

```
storage alias st0 MY_SWITCH:5.3L3
```

The tape device `st0` is assigned to the physical path name `MY_SWITCH:5.3L3`.

```
storage alias mc80 WWN[2:4e3:38fe3f:758eab]
```

The medium changer `mc80` is assigned to the worldwide name `WWN[2:4e3:38fe3f:758eab]`.

Removing tape aliases

You can remove aliases from tape drives or medium changers, or both, using the `storage unalias` command.

Step

1. To remove an alias from a tape drive or medium changer, enter the following command:

```
storage unalias {alias | -a | -m | -t}
```

alias 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.

Examples

```
storage unalias st0
```

```
storage unalias 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. You can then propagate the aliases to multiple storage systems.

Steps

1. To propagate tape aliases to multiple storage systems, create a file named `tape_alias` containing the tape alias information.

Example

```
storage unalias -a
storage alias st0 8.6
storage alias st1 8.7
storage alias mc0 8.1
```

2. Copy the file to the root volume of each storage system.
3. Execute the following command on each storage system:

```
source /vol/root_volume_name/tape_alias
```

root_volume_name specifies the root volume.

All the storage systems contain the same configuration information.

Note: 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.

UNIX shell scripts for propagating tape aliases

UNIX users can use a shell script to propagate the source script information to the storage systems.

You can use a shell script similar to the following one 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
```

How to add 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 you add a new medium changer, 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.

How to display tape drive and tape library information

You can view information about tape drives, tape medium changers, and tape drive connections to the storage system.

You can use this information to verify that the storage system detects the tape drive associated with the tape device. You can also verify the available tape device names associated with the tape drive. You can view information about qualified and nonqualified tape drives, tape libraries, and tape drive connections to the storage system.

Next topics

[Displaying information about tape drives](#) on page 33

[Displaying information about tape medium changers](#) on page 34

[Displaying information about tape drive connections to the storage system](#) on page 34

Displaying information about tape drives

You can view information about the tape drives on a storage system, such as the slot on the storage system and the tape drive's SCSI ID.

Step

1. Enter the following command:

```
sysconfig -t
```

Example

```
filer1>sysconfig -t
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)
```

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 preceding example, the Exabyte 8900 has SCSI ID 1 and is attached to a controller in slot 0b.

Note: Compression capacity in the display is an estimate; actual capacity depends on how much the data being written to the tape can be compressed.

Displaying information about tape medium changers

You can view the details about a tape medium changer, such as the slot to which it is attached in the storage system.

Step

1. To view details about tape medium changers, enter the following command:

```
sysconfig -m
```

Example

```
filer1>sysconfig -m
Medium changer (UC060000834:49.126) EXABYTE EXB-440
mc0 - medium changer device
```

Note: If the autoloader option of the medium changer is set to On, the medium changer information might not appear.

Displaying information about tape drive connections to the storage system

You can view the information about a tape drive connection to the storage system. You can view information such as the SCSI ID, Vendor ID, Product ID, and firmware version.

Step

1. Enter the following command:

```
sysconfig -v
```

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.

```
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

Controlling tape drives

You can control tape drives using the `mt` command. You can use the command to move and position the tape.

You can use the `mt` command to perform any of the following tasks:

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

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

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

Variables and options	Description
<code>-f</code> and <code>-t</code>	Indicate that the next parameter is a device. These options are interchangeable.
<i>device</i>	Is a tape device.
<i>command</i>	Is a command that controls the tape drives.
<i>count</i>	Specifies the number of times to execute a command that supports multiple operations.

The *command* option can be any one of the following.

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.
<code>bsf</code>	Move the tape backward, skipping a specified number of files.
<code>fsr</code>	Move the tape forward and position the tape on the end-of-tape side of the records.
<code>bsr</code>	Move the tape backwards and position the tape on the beginning-of-tape side of the records.

Command	Task
rewind	Rewind the tape.
offline	Rewind the tape and unload the tape medium, if possible.
status	Display information about a device and the drive associated with it.

Note: Use a no-rewind (nrst) devices for all tape status and movement operations. Using other rewind types can produce unwanted results.

Attention: When you use a unload/reload (urst) device with the `mt` command, you must use tape libraries for the backup and there must be enough tapes in the tape library. Otherwise, the tape drives involved terminate the command sequence or overwrite the same tape multiple times.

Next topics

[Moving a tape to the end of data](#) on page 36

[Moving forward to a file](#) on page 37

[Moving backward to the beginning of a file](#) on page 37

[Rewinding a tape](#) on page 38

[Taking a tape drive offline](#) on page 38

[Displaying status information](#) on page 39

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.

Step

1. Enter the following command:

```
mt -f device eom
```

device 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.

Step

1. To move forward to the beginning of a tape file, enter the following command:

```
mt -f device fsf n
```

device is the name of a tape device used on the tape.

n is the number of tape file marks you want to skip over going forward. The tape moves forward to the beginning of the *n*th file from its current file location.

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 positioned towards the beginning of tape from the current position.

Steps

1. Enter the following command:

```
mt -f device bsf n
```

device is the name of a tape device used on the tape.

n is the number of tape file marks you want to skip over going backward.

The tape moves backward to the end of the *n*th file from its current file location.

2. Enter the following command:

```
mt -f device fsf 1
```

The tape moves forward one file mark to the beginning of the desired file.

Example

If you enter the following commands in the middle of file 5 on the tape, the tape moves to the beginning of file 2 on the tape:

```
mt -f nrst0a bsf 4
mt -f nrst0a fsf 1
```

Rewinding a tape

If you use a no-rewind tape device to back up the data, the tape device does not automatically rewind the tape after the backup. To restore data backed up using such a tape device, you should rewind the tape when you load the tape drive.

Step

1. To rewind a tape, enter the following command:

```
mt -f device rewind
```

device is the name of a tape device used on the tape.

Example

```
mt -f nrst0a rewind
```

Related concepts

[Tape device name format](#) on page 22

Taking a tape drive offline

You take a drive offline to remove or change the tape cartridge. This operation rewinds the tape cartridge and ejects it from the tape drive. The device is still available to the system, but is not ready for I/O or tape movement.

About this task

You use a urst tape device to unload and reload a tape cartridge during a backup or restore operation. When you use a urst device, Data ONTAP waits for you to insert the new cartridge before continuing the operation. However, when you want to remove the current cartridge when no other operation is ongoing, you must use the `mt offline` command with an `nrst` tape device.

Step

1. To rewind the tape and take the tape drive offline by unloading the tape, enter the following command:

```
mt -f device offline
```

device is the name of a tape device.

Example

```
mt -f nrst0a offline
```

Related concepts

[Tape device name format](#) on page 22

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.

Step

1. To display status information about a tape device and the drive associated with it, enter the following command:

```
mt -f device status
```

device is the name of the tape device.

Example

```
filer1>mt -f nrst0a status
Tape drive: CERTANCEULTRIUM 3
Status: ready, write enabled
Format: LTO-3 800GB cmp
fileno = 0  blockno = 0  resid = 0
```

The following list describes the output of the command:

Tape drive	The model of the tape drive.
Status	Whether the tape drive is ready and write-enabled.
Format	The tape drive type, 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 it reached the end of the tape.

Qualified tape drives

A qualified tape drive is a tape drive that has been tested and found to work properly on storage systems. A qualified tape drive appears in the Data ONTAP kernel's internal tape qualification list or is represented by a valid tape configuration file in the controller's `/etc/tape_config` directory.

You can add support for tape drives to existing Data ONTAP releases using the tape configuration file. You can also view the current list of supported tape drives at the NOW Web site.

To add support to Data ONTAP for a tape drive that was qualified after the release of the Data ONTAP version you are using, copy the corresponding tape configuration file into the controller's `/etc/tape_config` directory.

Only qualified tape drives are listed in the tape qualification list and the tape libraries are not listed. For example, the tape library IBM TS3500 is not listed. However, the IBM LTO 4 tape drives that the IBM TS3500 contains are listed.

You can display information about qualified and nonqualified tape drives, tape libraries, and tape drive connections to the storage system.

Next topics

[Format of the tape configuration file](#) on page 40

[How the storage system qualifies a new tape drive dynamically](#) on page 42

Related information

<http://www.netapp.com/us/solutions/a-z/data-protection-devices.html>

http://now.netapp.com/NOW/download/tools/tape_config/

Format of the tape configuration file

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

The following table displays the format of the tape configuration file.

Item	Size	Description
<code>vendor_id</code> (string)	up to 8 bytes	The vendor ID as reported by the SCSI Inquiry command.
<code>product_id</code> (string)	up to 16 bytes	The product ID as reported by the SCSI Inquiry command.

Item	Size	Description
<i>id_match_size</i> (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.
<i>vendor_pretty</i> (string)	up to 16 bytes	If this parameter is present, it is specified by the string displayed by <code>sysconfig -v</code> or <code>sysconfig -t</code> ; otherwise, INQ_VENDOR_ID is displayed.
<i>product_pretty</i> (string)	up to 16 bytes	If this parameter is present, it is specified by the string displayed by <code>sysconfig -v</code> or <code>sysconfig -t</code> ; otherwise, INQ_PRODUCT_ID is displayed.

Note: The `vendor_pretty` and `product_pretty` fields are optional, but if one of these fields has a value, the other must also have a value.

The following table explains the description, density code, and compression algorithm for the various compression types such as l, m, h, and a.

Item	Size	Description
{l m h a}_description= (string)	up to 16 bytes	The string to print for <code>sysconfig -t</code> 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 table describes the optional fields available in the tape configuration file.

Field	Description
<code>autoload=</code> (Boolean yes/no)	Set this field to <code>yes</code> 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 <code>load</code> (start/stop unit) command. The default for this field is <code>no</code> .

Field	Description
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).

How the storage system qualifies a new tape drive dynamically

The storage system qualifies a tape drive dynamically by matching its *vendor_id* and *product_id* with the information contained in the tape qualification table.

The storage system's `/etc/tape_config` directory is created automatically when the storage system boots. When a tape configuration file is added 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.

Information about the tape persists as long as the file is in the directory or until the file is altered.

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 *vendor_id* and *product_id* match between information obtained during the tape discovery process and information contained in the internal tape qualification table. If the storage system discovers a match, it marks the tape drive as qualified and can access the tape drive. If the storage system cannot find a match, the tape drive remains in the unqualified state and is not accessed.

How to 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.

For a nonqualified tape drive to emulate a qualified tape drive, you must enter the nonqualified tape drive information in the `/etc/cloned_tapes` file. This file enables the storage system to register the drive as a clone of a qualified drive.

Next topics

[Displaying information about nonqualified tape drives](#) on page 43

Tape drive information required for emulation on page 43

Emulating a qualified tape drive on page 44

Displaying information about nonqualified tape drives

To make use of a nonqualified tape drive, you must determine whether it emulates any of the qualified tape drives.

Steps

1. If the storage system has accessed the tape drive through the `dump` or `mt` command, go directly to Step 3. If the storage system has not accessed the tape drive through the `dump` or `mt` command, go to Step 2.

2. To access the tape drive, enter the following command:

```
mt -f device status
```

device is any device that contains the tape drive number that you think is assigned to the tape drive.

Example

```
mt -f nrstla status
```

3. Enter the following command:

```
sysconfig -t
```

If the storage system has registered a tape drive as emulating a qualified tape drive, it displays a message similar to 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 similar to the following:

```
Tape drive (6.5) DLTXXXX (Non-qualified tape drive)
```

Tape drive information required for emulation

To emulate a qualified tape drive, you must know certain specific information about your nonqualified tape drive.

The required information is as follows:

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

Emulating a qualified tape drive

You can use a nonqualified tape drive by making it emulate a qualified tape drive.

Steps

1. Ensure that you have a tape adapter available on the storage system.
2. Disable the adapter port to which the tape drive will be attached.
3. Connect the tape drive to the storage system according the tape drive manufacturer's instructions.
4. Turn on the tape drive and wait for the tape drive to complete its power-on activities.
5. Enable the adapter interface. When the adapter is enabled, it will discover the device.

An error message is displayed, which tells you that the tape drive is unsupported.

6. Enter the following command:

```
sysconfig -t
```

This command creates the `/etc/cloned_tapes` file, if it does not exist already. Observe the vendor ID and product ID of the nonqualified devices.

Note: The cloned tapes emulation method cannot be used if the product ID contains spaces.

For example, the product ID *Ultrium 4-SCSI* cannot be used for cloning because it has a space between *Ultrium* and *4*. In such a case, you must use a configuration file.

7. Open the storage system's `/etc/cloned_tapes` file in a text editor on a client that can access it.
8. For each nonqualified tape drive, create a line with the following format in the `/etc/cloned_tapes` file:

```
[clone_vendor_ID] clone_product_ID EMULATES [vendor_ID] product_ID
```

clone_vendor_ID is the vendor of the nonqualified tape drive.

clone_product_ID is the model number of the nonqualified tape drive.

vendor_ID is the vendor of a qualified tape drive that you want the nonqualified tape drive to emulate.

product_ID 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:

```
QUANTUM DLT9000 EMULATES QUANTUM DLT7000
```

9. Enter the following command:

sysconfig -t

The system reads the `cloned_tapes` file and puts emulation into effect. Verify that the new device appears as an emulated device.

Related concepts

[Qualified tape drives](#) on page 40

What tape reservations are

Multiple storage systems can share access to tape drives, medium changers, bridges, or tape libraries. Tape reservations ensure that only one storage system accesses a device at any particular time by enabling either the SCSI Reserve/Release mechanism or SCSI Persistent Reservations for all Fibre Channel-attached tape drives, medium changers, bridges, and tape libraries.

Note: All of the systems that share devices in a Fibre Channel library, whether switches are involved or not, must use the same reservation method.

The SCSI Reserve/Release mechanism for reserving devices works well under normal conditions. However, during the interface error recovery procedures, the reservations can be lost. If this happens, initiators other than the reserved owner can access the device.

Reservations made with SCSI Persistent Reservations are not affected by error recovery mechanisms, such as loop reset; however, not all devices implement SCSI Persistent Reservations correctly.

Next topics

[Enabling tape reservations](#) on page 45

[Disabling tape reservations](#) on page 46

Enabling tape reservations

You can enable tape reservation using the `options tape.reservations` command. By default, tape reservation is turned off.

Step

1. To use either the SCSI Reserve/Release mechanism or SCSI Persistent Reservations, enter the following command:

```
options tape.reservations {scsi | persistent}
```

`scsi` selects the SCSI Reserve/Release mechanism.

`persistent` selects SCSI Persistent Reservations.

Disabling tape reservations

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

Step

1. To turn off tape reservations, enter the following command:

```
options tape.reservations off
```

NDMP management

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.

Next topics

[*What the advantages of NDMP are*](#) on page 47

[*What NDMP security is*](#) on page 48

[*How to manage NDMP*](#) on page 52

[*What NDMP debug messages are*](#) on page 56

[*Why you need to specify the NDMP version*](#) on page 58

[*NDMP extensions supported by Data ONTAP*](#) on page 59

[*Tape backup using NDMP services*](#) on page 60

What the advantages of NDMP are

Accessing 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.

What NDMP security is

Data ONTAP provides features for preventing or monitoring 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.

All non-root NDMP users on the root vFiler unit and all NDMP users on vFiler units are required to use NDMP passwords that are distinct from the password of the user. This password can be generated using the `ndmpd password userid` command.

Starting with Data ONTAP 7.3.2, NDMP users must have the `login-ndmp` capability to be able to successfully authenticate NDMP sessions. A predefined role named `backup`, by default, has the `login-ndmp` capability. To provide a user with the `login-ndmp` capability, the `backup` role can be assigned to the group to which the user belongs. However, when a group is assigned the `backup` role, all users within the group get the `login-ndmp` capability. Therefore, it is best to group all NDMP users in a single group that has the `backup` role.

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

Data ONTAP provides a set of commands that enable you to manage and monitor the security of NDMP connections to the storage system.

The following are the commands that monitor the security of NDMP connections to storage systems.

- The options `ndmpd.access` command enables you to restrict which hosts can run NDMP sessions with the storage system.
- The options `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.
- The options `ndmpd.connectlog` command allows you to enable or disable logging of NDMP connections attempts with the storage system.
- The options `ndmpd.password_length` command allows you specify an 8- or 16-character NDMP password.
- The `ndmpd password` command generates a secure NDMP password for administrators who do not have root privileges on the storage system. This password allows them to carry out NDMP operations through an NDMP-compliant backup application. For the NDMP password to be generated, the NDMP user must have the `login-ndmp` capability.

Next topics

[*Specifying NDMP access by host or interface*](#) on page 49

[*Specifying the NDMP authentication type*](#) on page 49

[Enabling or disabling NDMP connection logging](#) on page 50

[Specifying the NDMP password length](#) on page 51

[Generating an NDMP-specific password for non-root administrators](#) on page 51

Specifying NDMP access by host or interface

You can use the `options ndmpd.access` command to specify the hosts or interfaces through which NDMP sessions are permitted. Conversely, you can also specify hosts or interfaces to block from NDMP sessions.

Steps

1. Start a console session on the storage system to which you want to restrict NDMP access.
2. Enter the following command:

```
options ndmpd.access {all|legacy|host[!]=hosts|if [!]=interfaces}
```

`all` is the default value, which permits NDMP sessions with any host.

`legacy` restores previous values in effect before a Data ONTAP version upgrade.

Note: In the case of Data ONTAP 6.2, the `legacy` value is equal to `all`.

`host=hosts` 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 the host name or by an IPv4 or IPv6 address.

`host!=hosts` 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 the host name or by an IPv4 or IPv6 address.

`if=interfaces` is a parameter string that allows NDMP sessions through a specified interface or a comma-separated list of interfaces on this storage system.

`if!=interfaces` 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 for authenticating NDMP access to a storage system: plaintext and challenge. You can use the `options ndmpd.authtype` command to specify whether a storage system will accept plaintext, challenge, or both to authenticate NDMP session requests.

Steps

1. Start a console session on the storage system whose NDMP authentication method you want to specify.
2. Enter the following command:

```
options ndmpd.authtype {challenge|plaintext|plaintext,challenge}
```

`challenge` sets the challenge authentication method, generally the preferred and more secure authentication method.

`plaintext` sets the plaintext authentication method, in which the login password is transmitted as clear text.

`plaintext,challenge` sets both challenge and plaintext authentication methods.

Note: 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`.

Steps

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:

```
options ndmpd.connectlog.enabled {on|off}
```

Note: The value you set for this option will persist across storage system reboots.

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 IPv4 or IPv6 address (address from where the NDMP request originated)
- Destination IPv4 or IPv6 address (address of the storage system receiving the NDMP request)
- Source port (through which the NDMP request was transmitted)
- Storage system port (through which the NDMP request was received)

Example

```
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.
```

Specifying the NDMP password length

Administrators who have an account on a storage system but do not have root status on that storage system must 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.

About this task

The NDMP password can be either 8 or 16 characters long. The default value is 16 characters.

Step

1. To specify the NDMP password length, enter the following command on the storage system console:

```
options ndmpd.password_length length
```

length is either 8 or 16. If you enter a value other than 8 or 16, the storage system prompts you with the following message:

```
options ndmpd.password_length: Length must be either 8 or 16
```

Note: If this option is set to 8, all NDMP applications managing backups for the storage system must use an 8-character password for authentication.

Generating an NDMP-specific password for non-root administrators

An administrator without root privileges uses the NDMP-specific password for any NDMP backup and restore operation that requires password input in either a backup application or CLI environment.

Steps

1. Start a console session on the storage system you want to access.
2. Enter the following command:

```
ndmpd password username
```

username is the user name of the administrator.

The system returns an 8- or 16- character string, depending on the password length set using the `ndmpd.password_length` command. For example:

```
filer>ndmp password barbaraD
filer>ndmp password: QM12N%$cnaFWPBVe
```

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.

How to manage NDMP

You can enable or disable NDMP services, specify a preferred network interface, turn off a data connection specification, optimize performance, or terminate a session by using a set of `ndmpd` commands. You can also view the status of NDMP sessions using the `ndmpd` command.

Next topics

[Enabling and disabling NDMP services](#) on page 52

[Specifying a preferred network interface](#) on page 53

[Turning off a data connection specification](#) on page 53

[Displaying the general status information about NDMP sessions](#) on page 54

[Displaying detailed NDMP session information](#) on page 54

[Optimizing NDMP communication performance](#) on page 55

[Terminating an NDMP session](#) on page 56

Enabling and disabling NDMP services

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

Step

1. To enable or disable NDMP service, enter the following command:

```
ndmpd {on|off}
```

Use `on` to enable NDMP.

Use `off` to disable NDMP.

After you disable the NDMP service, the storage system continues processing all requests on already established sessions, but rejects new sessions.

Note: This setting is persistent across reboots.

Specifying 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.

About this task

By default, an NDMP data connection uses the same network interface as the NDMP control connection established by the NDMP backup application. However, to establish a data connection between NDMP-enabled storage systems over an alternate network, you need 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 through a corporate network. The same storage systems can also be interconnected through an isolated private network. To minimize load on the corporate network, the `options ndmpd.preferred_interface` command can be used to direct all NDMP data connections over the isolated private network.

Step

1. To specify the preferred network interface to be used for NDMP data connections, enter the following command:

```
options ndmpd.preferred_interface interface
```

interface identifies the network interface to be used for all NDMP data connections. Any network interface providing TCP/IP access can be specified. If no parameter is specified, the command returns the name of the interface currently configured for data connections. If no interface is currently set, it reports `disable`.

You can find the available network interfaces by using the `ifconfig -a` command.

Note: The preferred network interfaces that are set using the `options ndmpd.preferred_interface` command are persistent across storage system reboots.

Turning off a data connection specification

You can disable a preferred network interface specification and force the NDMP default interface to be used for data connections.

Step

1. To disable a preferred network interface specification and force the NDMP default interface to be used for data connections, enter the following command:

```
options ndmpd.preferred_interface disable
```

Note: The default value is `disable`.

Displaying the general status information about NDMP sessions

You can view the general status information to determine whether the NDMP session is operating as expected.

Step

1. To display general NDMP status information, enter the following command:

```
ndmpd status [session]
```

session is the specific session number for which you want the status. To display the status of all current sessions, do not enter any value for *session*.

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
```

Displaying detailed NDMP session information

You can view detailed NDMP session information to help you debug errors encountered during an NDMP session.

Step

1. To display detailed NDMP session information, enter the following command:

```
ndmpd probe [session]
```

session is the number of the session you want to probe. To display the detailed information about all sessions, do not enter any value for *session*.

Example of detailed status information

In the following example, the command shows the detailed status of session 4 with an IPv4 control connection and IPv6 data connection.

```
filer1>ndmpd probe 7
ndmpd ON.
Session: 7
```

```

isActive:                TRUE
protocol version:        4
effHost:                 Local
authorized:              TRUE
client_addr:             192.0.2.1
spt.device_id:           none
spt.ha:                  -1
spt.scsi_id:             -1
spt.scsi_lun:            -1
tape.device:             not open
tape.mode:               Read only
mover.state:             Idle
mover.mode:              Read
mover.pauseReason:       N/A
mover.haltReason:        N/A
mover.recordSize:        0
mover.recordNum:         0
mover.bytesMoved:        0
mover.seekPosition:      0
mover.bytesLeftToRead:   0
mover.windowOffset:      0
mover.windowLength:      0
mover.position:          0
mover.setRecordSizeFlag: false
mover.setWindowFlag:     false
mover.connect.addr_type: LOCAL
data.operation:          None
data.state:              Connected
data.haltReason:         N/A
data.connect.addr_type:  TCP_IPV6
data.connect.addr:        [2001:0db8::10]
data.connect.port:        63920
data.bytesProcessed:     0

```

Optimizing NDMP communication performance

You can optimize the performance of the NDMP socket through which the storage system communicates with the DMA.

About this task

You can optimize performance for either minimal transmission delay or throughput. By default, the performance is optimized for overall throughput, in which the packets are grouped in time blocks of 200 ms. If the communication performance is optimized for minimal transmission delay, the queued packets are sent immediately.

Step

1. To optimize NDMP communication performance, enter the following command:

```
options ndmp.tcpnodelay.enable {on|off}
```

on optimizes for minimal transmission delay.

off optimizes for overall throughput.

Terminating an NDMP session

If an NDMP session is not responding, you can terminate it using the `ndmpd kill` command. The `ndmp kill` command allows nonresponding sessions to be cleared without the need for a reboot.

Step

1. To terminate an NDMP session, enter the following command:

```
ndmpd kill session
```

session is the specific NDMP session you want to terminate.

Note: If you want to terminate all NDMP sessions, use the `ndmpd killall` command.

What NDMP 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.

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.

The NDMP session log files are stored in the `/etc/log` directory. The file name is `ndmpdlog.yyyyymmdd`, where *yyyy* is the year, *mm* is the month, and *dd* is the date. For example, the session log file generated on December 5, 2008, is named `ndmpdlog.20081205`. 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 previous seven days and the current day.

Next topics

[Enabling the NDMP debug log messages](#) on page 57

[Displaying the NDMP debug log level](#) on page 57

[Changing NDMP debug log messages](#) on page 58

[Displaying an NDMP session log file](#) on page 58

Enabling the NDMP debug log messages

To update the NDMP session log files stored in the `/etc/log` directory, you have to enable the NDMP debug logging.

Step

1. To enable NDMP debug logging, enter the following command:

```
ndmpd debug n
```

`n` specifies the debug level from 0 to 70. To turn off the debug messages, use 0. To turn on the debug messages, use a nonzero value. The default level is 0.

The following list describes the debug levels that are supported:

- 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 the display of environment variables as well as any exceptions issued by the NDMP server implementation.
- 70** Same as level 50, but includes the 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

You can display the currently set NDMP debug levels using the `ndmpd debug` command.

Step

1. To see the NDMP debug levels currently set, enter the following command:

```
ndmpd debug
```

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 can use the `ndmp debug` command to display the debug messages on the storage system console, to write the debug messages into the NDMP debug log file, or to print stack traces for any exceptions issued by the NDMP server implementation.

Step

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

You can display an NDMP session log file from a UNIX environment or an NT environment.

Step

1. Depending on your operating system, choose one of the following methods of displaying an NDMP session log file:
 - In a UNIX 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 `cat` or `more` UNIX commands or a text editor.
 - In an NT environment, 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.

Why you need to specify the NDMP version

Data ONTAP provides full support for NDMP versions 2, 3, and 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 with a `yes` or a `no`.

Next topics

Displaying the NDMP version on page 59

Specifying the NDMP version on page 59

Displaying the NDMP version

The `ndmp version` command displays the highest version of NDMP that the storage system is currently set to use.

Step

1. Enter the following command:

```
ndmpd version
```

The highest version that NDMP currently allows to be used is displayed.

Specifying the NDMP version

You can use the `ndmpd version` command to control the highest and default NDMP version allowed.

About this task

If you know that your backup application does not support NDMP version 4 and does not negotiate versions, you can use this command to specify the highest version that Data ONTAP supports, so that the application can operate correctly.

The NDMP version that is set using the `ndmpd version` command is persistent across storage system reboots.

Step

1. To specify the NDMP version you want, enter the following command:

```
ndmpd version n
```

`n` is the version you want to specify. The options available are 2, 3, and 4. The default highest version is 4.

NDMP extensions supported by Data ONTAP

NDMP version 4 provides a mechanism for creating NDMP v4 protocol extensions without requiring modifications to the core NDMP v4 protocol.

Following are some of the NDMP v4 extensions supported by Data ONTAP:

- Restartable backup.

- SnapVault management
- SnapMirror management
- Snapshot extension

To benefit from these NDMP v4 extensions, the NDMP backup applications must support these extensions.

Related information

<http://www.ndmp.org/>

Tape backup using NDMP services

You can use NDMP-enabled commercial backup applications to perform network-based tape backup and recovery.

Next topics

Common NDMP tape backup topologies on page 60

Considerations when using NDMP on page 61

Tape devices and configurations you can use with the storage system on page 61

Preparing for basic NDMP backup application management on page 62

What environment variables do on page 64

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-tape attached to another storage system

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 called 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.

Considerations when using NDMP

You have to take into account a list of considerations when starting the NDMP service on your storage system.

- Data ONTAP supports a maximum of 16 concurrent backups, restores, or both on a local tape drive.
This includes backups initiated by NDMP as well as by the storage system's `dump` and `restore` commands.
However, a storage system supports a maximum of 32 dump or restore sessions.
- NDMP supports a maximum of 128 concurrent sessions on NearStore 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 a system-generated NDMP-specific password (to authenticate a non-root user or administrator).
- NDMP services can 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 might be 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 can disable file history generation to reduce backup time. See your backup application documentation to determine if it is possible to disable NDMP file history generation.

Tape devices and configurations you can use with the storage system

You can use different types of tape devices and configurations on your storage system.

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

- Stand-alone 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

When you use NDMP to back up the storage system to attached tape libraries, you need to set the tape library autoload setting to `Off`. If the autoload setting is `On`, the storage system uses the tape library 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:

- `mc n` or `/dev/mc n`
- `spt n` or `/dev/spt n`

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 WWNs.

To view the tape libraries 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.

Examples

The following is an example of an output 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 for basic NDMP backup application management

To enable a storage system for basic management by a commercial NDMP backup application, you must enable the storage system's NDMP support and specify the backup application's configured NDMP version, host IP address, and authentication method.

About this task

If an operator without root privileges to the storage system is using a backup application, that user must use a storage system-generated NDMP-specific password to carry out backup operations on that storage system.

Steps

1. To enable NDMP, enter the following command at the console command line of the target storage system:

```
ndmpd on
```

2. To specify the NDMP version to support on your storage system, enter the following command:

```
ndmpd version {2|3|4}
```

Note: The version must match the version configured for your NDMP backup application.

3. To specify a restricted set of NDMP backup application hosts that can connect to the storage system, enter the following command:

```
options ndmpd.access hosts
```

hosts is a comma-separated list of host names or IP addresses of nodes permitted to start NDMP sessions with the storage system.

Note: By default, all hosts have NDMP access.

4. Specify the authentication type (plaintext, challenge, or plaintext and challenge) required for an NDMP connection to this storage system. For example:

```
options ndmpd authtype plaintext,challenge
```

This setting must include the authentication type supported by the NDMP backup application.

Note: The challenge authentication type is the default for this option.

5. If operators without root privilege on the storage system are carrying out tape backup operations through the NDMP backup application, make sure they have a user administration account on the storage system.

- a. If the operator does not have a user administration account on the storage system, enter the following command:

```
useradmin useradd username
```

- b. If you want to know the system-generated NDMP-specific password, enter the following command:

```
ndmpd password username
```

Use this user name and password to connect to the storage system to carry out NDMP backup and restore operations.

Related tasks

[Enabling and disabling NDMP services](#) on page 52

[Specifying a preferred network interface](#) on page 53

[Specifying NDMP access by host or interface](#) on page 49

[Specifying the NDMP authentication type](#) on page 49

[*Generating an NDMP-specific password for non-root administrators*](#) on page 51

What environment variables do

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: The setting and examining of environment variables are typically transparent to backup administrators; that is, the backup application sets them automatically.

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 might 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. For information, see your backup application documentation.

Related concepts

[*Environment variables supported for dump*](#) on page 78

Data backup using the dump engine

Dump is a Snapshot copy-based backup and recovery solution from Data ONTAP that helps you to back up files and directories from a Snapshot copy to a tape device and restore the backed up data to a storage system.

You can back up your file system data, such as directories, files, and their associated security settings to a tape device by using the dump backup. You can backup an entire volume, an entire qtree, or a subtree that is neither an entire volume nor an entire qtree.

You can perform a dump backup or restore by using NDMP-compliant backup applications or by using the dump and restore CLI commands.

When you perform a dump backup, you can specify the Snapshot copy to be used for a backup. If you do not specify a Snapshot copy for the backup, a base Snapshot copy is created for the backup.

You can perform level-0, incremental, or differential backups to tape by using the dump engine.

Next topics

[*How a dump backup works*](#) on page 66

[*What the dump engine backs up*](#) on page 66

[*What increment chains are*](#) on page 67

[*How to specify tape devices for the backup*](#) on page 69

[*What the /etc/dumpdates file is*](#) on page 69

[*What the blocking factor is*](#) on page 70

[*How to use the dump backup*](#) on page 71

[*Considerations before using the dump backup*](#) on page 73

[*When to restart a dump backup*](#) on page 74

[*How a dump restore works*](#) on page 75

[*What the dump engine restores*](#) on page 75

[*Considerations before restoring data*](#) on page 76

[*How to prepare the destination for a dump restore*](#) on page 77

[*How online migration affects tape backup*](#) on page 77

[*How to perform a dump backup and restore using NDMP services*](#) on page 78

[*How to perform a dump backup using the CLI*](#) on page 93

[*How to perform a dump restore using the CLI*](#) on page 114

How a dump backup works

A dump backup writes file system data from disk to tape using a predefined process. It is optimized for data restoration to a storage system using the dump restore.

You can back up an entire volume, an entire qtree, or a subtree that is neither an entire volume nor an entire qtree.

The following table describes the process that Data ONTAP uses 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 the files to be backed up. If you are backing up an entire volume or qtree, Data ONTAP combines this stage with Stage 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 ACL information (if applicable) to tape.

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

The dump backup names each Snapshot copy it creates as `snapshot_for_backup.n`, where *n* is an integer starting at 0. Each time the dump backup 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 performs multiple dump backups simultaneously, the dump engine creates multiple Snapshot copies. For example, if Data ONTAP is running two dump backups simultaneously, you find the following 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, the dump engine does not create an additional Snapshot copy.

What the dump engine backs up

The dump engine can back up a file, directory, qtree, or an entire volume to a tape.

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

- UNIX GID, owner UID, and file permissions
- UNIX access, creation, and modification time
- File type
- File size
- DOS name, DOS attributes, and creation time
- Access Control Lists (ACLs)
- Qtree information
- LUN and LUN clones

You can back up only an entire LUN object; you cannot back up a single file within the LUN object. Similarly, you can restore an entire LUN object but not a single file within the LUN.

Note: The dump engine backs up LUN clones as independent LUNs.

When you back up data to tape, the `dump` command does not back up the LUN clones that are inconsistent. For all other LUN clones, the `dump` command locks their backing Snapshot copies to ensure that they do not become inconsistent during the backup.

When you back up a qtree SnapMirror destination to tape, only the data on the qtree is backed up. The associated metadata is not backed up. Therefore, when you try to restore the qtree, only the data on that qtree is restored. Information about the qtree SnapMirror relationships is not available in the backup and therefore is not restored.

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

Other dumps and restores preserve permissions.

What increment chains are

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 and restores effectively.

There are two types of increment chains:

- A consecutive increment chain is a sequence of incremental backups that starts with level 0 and is raised by 1.
- A nonconsecutive increment chain is one in which incremental backups skip levels or have levels that are out of sequence, such as 0, 2, 3, 1, 4, or more commonly, 0,1,1,1 or 0,1,2,1,2.

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 following table explains the bases of the incremental backups.

Back-up order	Increment level	Increment chain	Base	Files backed up
1	0	Both	Files on the storage system	All files in the back up 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
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

An incremental backup has certain limitations:

- During an incremental backup, the `dump` command backs up only files that have a timestamp later than the backup timestamp stored in the `/etc/dumpdates` file or the `BASE_DATE` environment variable.

Dump is a timestamp-based backup. During an incremental backup, the `dump` command determines the changed or modified files since the previous backup, using the timestamp stored in the `/etc/dumpdates` file or in the `BASE_DATE` environment variable. In Data ONTAP, there can be instances where files are replaced with their earlier version, for example, when using `snapmirror resync` and `snap restore`.

- If you attempt an incremental backup of a volume SnapMirror destination after breaking the SnapMirror relationships, you might lose data.

In these cases, you must perform a level-0 backup instead.

How to specify tape devices for the backup

You must specify at least one tape device to do a 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.

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

If the backup requires more tape devices than the number specified, the last tape device is used for all remaining tape files.

Attention: 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 that the 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.

What the `/etc/dumpdates` file is

The `/etc/dumpdates` file enables you to keep track of backups.

It records the following information:

- The name of the backup, which can be 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

Reasons 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 the 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

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.

```
/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
```

What the blocking factor is

A tape block is 1,024 bytes of data. During a tape backup or restore, you can specify the number of tape blocks that are transferred in each read/write operation. This number is called the blocking factor.

Data ONTAP 6.5.3 and later releases support a blocking factor between the range of 4 KB and 256 KB. The default blocking factor is 63 KB.

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 select.

If you plan to restore a backup 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.

During an NDMP backup, the `MOVER_RECORD_SIZE` determines the blocking factor. Data ONTAP allows a maximum value of 256 KB for `MOVER_RECORD_SIZE`.

Related information

<http://www.ndmp.org/>

How to use the dump backup

To use the dump engine for a successful backup, you have to follow certain recommendations.

To reduce the risk of encountering an error that requires restarting the dump backup, avoid backing up too much data in a single dump backup.

However, if the dump backup encounters an error, you might be able to correct the error and proceed from the point where the backup operation terminated.

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.

If you suspect that a backup might have errors, you should verify the backup by performing a test restore.

Next topics

[*How to minimize backup time and data loss*](#) on page 71

[*How to decrease tape backup time*](#) on page 72

[*How to minimize the number of tape drives*](#) on page 72

[*What to label on the backup tapes*](#) on page 72

Related tasks

[*Specifying a test restore*](#) on page 129

How to minimize backup time and data loss

You can minimize both the time required to perform a backup and the possibility of data loss by following certain guidelines.

The shorter the time for the dump backup to finish, the more incremental backups you can perform. Follow these guidelines to minimize the backup time and data loss:

- 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 backup runs faster if the dump path specifies one of the following:

- A full volume

- Full qtrees
- A full volume excluding qtrees

How to decrease tape backup time

You can decrease the tape backup time in various ways.

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 moderate.
- Do not divide a backup into more than 15 qtrees.

How to minimize the number of tape drives

Data ONTAP supports the RMT protocol and therefore several storage systems can share the same tape drive. You can minimize the number of such shared tape drives.

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 backup on each storage system so that it starts only when no other storage systems are using the tape drive.

What to label on the backup tapes

For ease of use during a restore, you must label the backup tapes with certain information.

You have to 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 multfile tape
A brief description of the contents of each tape file on a multfile tape helps you locate a desired tape file for restoring.
- The sequence of tape files on a multfile tape

This enables you to specify which file to restore. To specify a tape file, you must know the location of the tape file in the sequence of tape files.

- The Data ONTAP version of each backup

Considerations before using the dump backup

Before backing up data using the `dump` command, you must have a clear idea of how much data you will be backing up and how many tapes you will need to store the data.

Next topics

[*Determining the amount of backup data*](#) on page 73

[*Estimating the number of tapes for the backup*](#) on page 73

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.

Step

1. For each item that you want to back up, enter the following command:

```
df path_name
```

path_name is the name of the path.

Note: For multiple items, such as multiple volumes, add the data for each item to determine the total amount of data to be backed up.

Estimating the number of tapes for the backup

You must estimate the number of tapes required for the backup before executing the `dump` command. This estimate helps you to ensure that the `dump` command does not fail because it runs out of tapes. It also helps you to load the required number of tapes in the tape drives or libraries in advance for an unattended backup.

About this task

If you initiate the `dump` command from the console and have not loaded enough tapes, Data ONTAP prompts you to load additional tapes. However, if you initiate the `dump` command from a Remote Shell connection and have not loaded enough tapes, you do not see the prompts from Data ONTAP and the `dump` command terminates.

Steps

1. Determine the capacity of the tape device you are using for the backup by entering the following command:

```
sysconfig -t
```

2. Determine the amount of data to be backed up.
3. Divide the amount of data by the capacity of the tape.
4. If your estimate indicates that your data will nearly fill the last tape, add a tape to the estimate. This avoids a backup failure if the backup exceeds your estimate. This is especially important when using compression, because compression rates vary based on the data.

Related tasks

[Determining the amount of backup data](#) on page 73

When to restart a dump backup

A dump backup sometimes does not finish because of internal or external errors, such as tape write errors, power outages, accidental user interruptions, or internal inconsistency on the storage system. If your backup fails for one of these reasons, you can restart it.

You can 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. You can interrupt a long backup and restart it later if a more urgent restore (or backup) requires the same tape drive. Restartable backups persist across reboots.

Starting with Data ONTAP 7.2.3, you can restart dumps of volumes containing qtree SnapMirror destinations.

Dumps of volumes containing qtree SnapMirror destinations read data from multiple Snapshot copies and write them onto a tape. When such a dump operation is aborted and left in a restartable state, the associated Snapshot copies are locked. These Snapshot copies are released after the backup context is deleted. To view the list of locked Snapshot copies, run the `backup status` command.

Example

```
filer> backup status 2

State: RESTARTABLE                Type:      ndmp
Path: /vol/vol1                   Level:     0
Snapshot: filer(0101184236)_vol1_filer_svp-dst.0
Snapshot: snapshot_for_backup.9 [Dec 27 00:41]
Options:      b=63, X
Devices:      [none]
```

```
Completed: 1 tapefile(s)
Last Update: Thu Dec 27 00:41:23 2007
```

The backup status output provides the following information:

State	The state of the dump: ACTIVE or RESTARTABLE.
Type	The type of invocation of dump: CLI or NDMP.
Path	The dump path.
Level	The level of the dump (0 through 9).
Snapshot	The Snapshot copies of the path that is being backed up.
Options	All the options specified for the backup and their respective parameters.
Devices	The current device to which the dump is writing.
Completed	The number tape files that have already been copied.
Last Update	The time and date of the last completed update.

Related tasks

[Restarting a dump command backup](#) on page 113

How a dump restore works

A dump restore writes file system data from tape to disk using a predefined process.

The process in the following table shows how the dump restore 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 stages 2 and 3 until all the specified files are copied from the tape.

What the dump engine restores

The dump engine enables you to recover all the information that you backed up.

The dump engine can recover the following data:

- Contents of files and directories
- UNIX file permissions
- ACLs

If you restore a file that has only UNIX file permissions into an 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.

Attention: Data ONTAP 7.3 and later releases support more than 192 Access Control Entries (ACEs) per ACL, whereas earlier versions support only a maximum of 192. Therefore, any data migration from Data ONTAP 7.3 or later releases to an earlier release will result in loss of ACLs.

- Qtree information
Qtree information is used only if a qtree is restored to the root of a volume. Qtree information is not used if a qtree is restored to a lower directory, such as `/vol/vol0/subdir/lowerdir`, and it ceases to be a qtree.
- All other file and directory attributes
- Windows NT streams
- LUNs

A LUN must be restored to a volume level or a qtree level for it to remain as a LUN. If it is restored to a directory, it is restored as a file because it does not contain any valid metadata.

You can restore data from previous versions of Data ONTAP using the dump engine. If you want to 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, you can do 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, and then performing incremental restores will not restore the incremental backups because the data is in a different format from the level 0 restore. In such a case, you have to repeat the level 0 restore before you can restore incremental backups.

Considerations before restoring data

Before performing a dump restore, you need to ensure that you have the required information and prepare the destination for the restore.

Before restoring data, 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 during the backup

Required tape drives and tapes

You must meet the following requirements for the restore operation to be successful:

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

Required space on the destination storage system

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

Attention: The restore operation will not start if there are insufficient inodes and space available. If you use the `F` option to force a restore to occur, the restore operation will fill up the available space and then abort.

How to prepare the destination for a dump restore

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.

If you are restoring a volume, you must create a new volume. If you are restoring a qtree or a directory, you must rename or move files that are likely to have the same names as files you are restoring.

Attention: If a restored file has the same name as an existing file, the existing file is overwritten by the restored file. However, the directories are not overwritten.

How online migration affects tape backup

You cannot perform a tape backup or restore of a vFiler volume that is currently undergoing online migration.

Online migration affects tape backup in the following ways:

- Backup or restore of a vFiler volume currently under online migration results in the following error:
The volume 'vol_name' is currently under migration
- Transfer of file system data using the `ndmcopy` command results in a failure.
- Incremental backup of a volume is not possible, if the volume was migrated after the level-0 backup.

How to perform a dump backup and restore using NDMP services

You can perform a dump backup or restore by using NDMP-compliant backup applications.

Data ONTAP provides a set of environment variables that enable you to perform a tape backup and restore using NDMP services. The dump engine-based restore using NDMP also supports enhanced direct access recovery (DAR), which enables directory DAR and DAR of files with NT streams.

You can also transfer file system data between storage systems by using the `ndmpcopy` command.

Next topics

[Environment variables supported for dump](#) on page 78

[Enabling or disabling enhanced DAR functionality](#) on page 87

[What the `ndmpcopy` command does](#) on page 88

[Displaying file history statistics](#) on page 92

Environment variables supported for dump

Data ONTAP supports environment variables for dump, which have an associated default value. However, you can manually modify these default values.

If you manually modify the values set by the backup application, the application might behave unpredictably. This is 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.

The following table contains descriptions of what the environment variables supported by Data ONTAP do if they are used.

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> value	-1	Specifies the start date for incremental backups. There is no equivalent option for the dump 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 <i>/etc/dumpdates</i> file for controlling incremental backups.
DEBUG	Y or N	N	Specifies that debugging information is printed. Note: There is no command line equivalent for the DEBUG variable.

Environment variable	Valid values	Default	Description
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 recovery to work, the backup application must provide positioning information. If this variable is set to Y, the backup application will specify the file or directory names and the positioning information.
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 DUMP_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	<p>Specifies if enhanced DAR functionality is instantiated. Enhanced DAR functionality supports directory DAR, and DAR of files with NT Streams. It provides performance improvements. Enhanced DAR during restore is possible only if the following conditions are met:</p> <ul style="list-style-type: none"> • Data ONTAP supports enhanced DAR (Data ONTAP 6.4 or later) • File history is enabled (HIST=Y) during the backup • The <code>ndmpd.offset_map.enable</code> option is set to "on" • ENHANCED_DAR_ENABLED variable is set to "Y" during restore
EXCLUDE	<i>pattern_string</i>	none	<p>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:</p> <ul style="list-style-type: none"> • 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.

Environment variable	Valid values	Default	Description
EXTRACT	Y or N	N	Specifies that subtrees of a backed-up data set 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 the names of the subtrees to be extracted. If a file name specified 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 DARs (DIRECT=Y).
FILESYSTEM	<i>string</i>	none	Specifies the path name of the root of the data that is being backed up. For example, <code>/vol/vol0/etc</code> .
FORCE	Y or N	N	Specifies that a restore operation continues, regardless of inode limitations. The FORCE variable is equivalent to the <code>F</code> 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.

Environment variable	Valid values	Default	Description
HIST	Y or N	N	<p>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 the file history collection, you can set this variable to N.</p> <p>Note: You should not set the HIST variable to Y if the backup application does not support file history.</p>
IGNORE_CTIME	Y or N	N	<p>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.</p>
IGNORE_QTREES	Y or N	N	<p>Specifies that the restore operation does not restore qtree information from backed up qtrees. The IGNORE_QTREES variable is equivalent to the Q option of the restore command.</p>
LEVEL	0-9	0	<p>Specifies the backup level. Level 0 copies the entire data set. 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.</p>

Environment variable	Valid values	Default	Description
LIST	Y or N	N	Specifies that backed-up file names and inode numbers be 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 be 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 DMP_NAME variable.
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.

Environment variable	Valid values	Default	Description
NO_ACLS	Y or N	N	Specifies that ACLs not be 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 be 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 data set 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.
NOWRITE	Y or N	N	Specifies that the restore operation not write data to the disk. The NOWRITE variable is equivalent to the N option of the <code>restore</code> command. This variable is used for debugging.

Environment variable	Valid values	Default	Description
RECURSIVE	Y or N	Y	<p>Specifies that directory entries during a DAR restore be expanded. The <code>DIRECT</code> and <code>ENHANCED_DAR_ENABLED</code> environment variables must be enabled (set to <code>Y</code>) as well. If the <code>RECURSIVE</code> variable is disabled (set to <code>N</code>), only the permissions and ACLs for all the directories in the original source path are restored from tape, not the contents of the directories. If the <code>RECURSIVE</code> variable is <code>N</code> or the <code>RECOVER_FULL_PATHS</code> variable is <code>Y</code>, the recovery path must end with the original path.</p> <p>Note: If the <code>RECURSIVE</code> variable is disabled and if there are more than one recovery path, all the recovery paths must be contained within the longest of the recovery paths. Otherwise, an error message is displayed.</p> <p>For example, the following are valid recovery paths as all the recovery paths are within <code>foo/dir1/deepdir/myfile</code> :</p> <ul style="list-style-type: none"> • <code>/foo</code> • <code>/foo/dir</code> • <code>/foo/dir1/deepdir</code> • <code>/foo/dir1/deepdir/myfile</code> <p>The following are invalid recovery paths:</p> <ul style="list-style-type: none"> • <code>/foo</code> • <code>/foo/dir</code> • <code>/foo/dir1/myfile</code> • <code>/foo/dir2</code> • <code>/foo/dir2/myfile</code>

Environment variable	Valid values	Default	Description
RECOVER_FULL_PATHS	Y or N	N	Specifies that full recovery path will have their permissions and ACLs restored after the DAR. <code>DIRECT</code> and <code>ENHANCED_DAR_ENABLED</code> must be enabled (set to Y) as well. If <code>RECOVER_FULL_PATHS</code> is Y, recovery path must end with the original path. If directories already exist on the destination volume, their permissions and ACLs will not be restored from tape.
UPDATE	Y or N	Y	Updates the <code>/etc/dumpdates</code> file.

Enabling or disabling enhanced DAR functionality

Enhanced direct access recovery (DAR) functionality provides support for directory DAR and DAR of files with NT Streams. This feature is supported only for the NDMP-initiated dump backup and restore and provides improved restore performance.

About this task

By default, enhanced DAR functionality is enabled in Data ONTAP; however, you can enable or disable it using the `options ndmpd.offset_map.enable` command.

Note: You should enable or disable this functionality before you initiate the NDMP dump operation.

Because an offset map has to be created and written onto tape, enabling enhanced DAR functionality might impact the backup performance.

Step

1. To enable enhanced DAR functionality on your storage system, enter the following command:

```
options ndmpd.offset_map.enable [on|off]
```

`on` enables enhanced DAR functionality.

`off` disables enhanced DAR functionality

Related concepts

[Considerations when using NDMP](#) on page 61

What the ndmcopy command does

The `ndmcopy` command enables a storage system administrator to transfer file system data between storage systems that support NDMP v3 or v4 and the UFS dump format.

The `ndmcopy` command 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 a recovery operation on the destination storage system. The command establishes control connections to the storage systems and facilitates data connection establishment. After connections are established, it facilitates data transfer. You can use host names or IP addresses of source and destination storage systems in the `ndmcopy` command.

Starting with Data ONTAP 7.3.3, the `ndmcopy` command supports IPv6 addresses of storage systems also. You can use IPv6 addresses to establish control connections to source and destination storage systems and can request the `ndmcopy` command to use an IPv6 address mode to establish the data connection.

Using the `ndmcopy` command, you can perform both full and incremental data transfers; however, incremental transfers are limited to a maximum of two levels (one full and up to two incremental backups). You can transfer full or partial volumes, qtrees, directories, or individual files.

You cannot perform a block-level transfer using the `ndmcopy` command.

Next topics

[Copying data using ndmcopy](#) on page 88

[Examples of the ndmcopy command](#) on page 90

Copying data using ndmcopy

You can invoke the `ndmcopy` command 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 `ndmcopy` on a single storage system that is both the source and the destination of the data transfer. The command can also be entered from a storage system that is not the source or the destination.

Step

1. To copy data within a storage system or between storage systems using `ndmcopy`, enter the following command:

```
ndmcopy [options][source_filer:]source_path
[destination_filer:]destination_path [-mcs {inet|inet6}][-mcd {inet|inet6}][-md {inet|inet6}]
```

options can be one or more of the following:

- `-sa username:[password]` is the source authorization that specifies the user name and password for connecting to the source storage system.

Note: For a user without root privilege, you must specify the user's system-generated NDMP-specific password and not the regular storage system account password.

- `-da username:[password]` is the destination authorization that specifies the user name and password for connecting to the destination storage system.
- `-st {md5|text}` sets the source authentication type to be used when connecting to the source storage system.
- `-dt {md5|text}` sets the destination authentication type to be used when connecting to the destination storage system.

Note: md5 is the default authentication type used. The md5 authentication type exchanges the user name and password in encrypted form. The text authentication type exchanges the user name and password in clear text.

- `-l` 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 ndmcopy debug log messages. ndmcopy debug log files appear in the root volume `/etc/log` directory. The ndmcopy debug log file names are in the form `ndmcopy.yyyyymmdd`.
- `-f` enables forced mode. This mode enables overwriting system files in the `/etc` directory on the root volume.
- `-h` prints the help message.

source_filer and *destination_filer* can be host names or IP addresses.

The `ndmcopy` command determines the address mode for control connections as follows:

- When an IP address (IPv4 or IPv6) is specified instead of the host name, the addressing mode for the control connection is the corresponding IP address type.
- When a host name is specified and it resolves to both IPv6 and IPv4 addresses, IPv6 addressing mode is used.

You can override these rules by using the `-mcs` and `-mcd` options.

source_path and *destination_path* are the absolute path names of the directories to be used during the data transfer.

`-mcs` specifies the preferred addressing mode for the control connection to the source storage system. *inet* indicates an IPv4 address mode and *inet6* indicates an IPv6 address mode.

`-mcd` specifies the preferred addressing mode for the control connection to the destination storage system. *inet* indicates an IPv4 address mode and *inet6* indicates an IPv6 address mode.

`-md` specifies the preferred addressing mode for communication between the source and the destination storage system. *inet* indicates an IPv4 address mode and *inet6* indicates an IPv6 address mode.

If you do not use the `-md` option in the `ndmcopy` command, the addressing mode for the data connection is determined as follows:

- If either of the addresses specified for the control connections is an IPv6 address, the data connection address mode is IPv6.
- If both the addresses specified for the control connections are IPv4 addresses, the `ndmcopy` command first attempts an IPv6 address mode for the data connection. If that fails, the command uses an IPv4 address mode.
- When a DNS name is specified for the control connections, the `ndmcopy` command attempts an IPv6 DNS lookup followed by an IPv4 DNS lookup. The data connection address mode is determined by the outcome of the DNS lookup.

Note: An IPv6 address, if specified, must be enclosed within square brackets.

Related tasks

[Generating an NDMP-specific password for non-root administrators](#) on page 51

Related references

[Examples of the `ndmcopy` command](#) on page 90

Examples of the `ndmcopy` command

You can migrate data from the source path to a destination path on the same storage system or to a different destination path on a remote host. You can also migrate data from a source path on a remote host to a destination path on the same host or to a destination path on a remote host.

In these examples, *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.

Example of migrating data from a source path to a different destination path on the same storage system

This sample 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
```

The following 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 of migrating data from a source path to a different destination path on a remote host

This sample 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 of migrating data from a source path on remote host to a destination path on the local storage system

This sample 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 of migrating data from a source path on a remote host to a destination path on another remote host

This sample 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 of overwriting the /etc directory during the root volume migration

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*. This helps prevent 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>ndmpcopy -da username:password -f /vol/rootvol
remotehost1:/vol/rootvol
```

Example of the ndmpcopy command where the address modes are explicitly set to IPv6

This sample command explicitly sets the control connections and the data connection to use IPv6 address mode. In this command *remotehost1* is the host name that resolves to an IPv6 address.

```
myhost>ndmpcopy -sa username:password -da username:password
-l 0 -mcs inet6 -mcd inet6 -md inet6 remotehost1:/vol/vol0/
source_path [2001:0db8::10]:/vol/vol0/destination_path
```

Displaying file history statistics

You can view detailed statistics about file history performance of currently active dump sessions using the `stats show ndmp` command.

Step

1. Enter the following command:

```
stats show ndmp
```

The output of the `stats show ndmp` command includes the following statistics:

- Total number of directory file history entries generated
- Total number of normal file history entries generated
- Total number of messages sent to the file history thread
- Minimum, maximum, and average delay times for adding file history entries
- Minimum, maximum, and average delay times for the file history thread to send messages to the NDMP thread
- Total number of file history flush calls
- Minimum, maximum, and average flush times
- Total number of times the dump thread had to block because of slow processing by the file history thread
- Maximum number of outstanding buffers to the file history thread

Sample output of the stat show ndmp command

```
filer*> stats show ndmp
ndmp:Session 01:dir_buffers_sent:19
ndmp:Session 01:node_buffers_sent:0
ndmp:Session 01:dir_send_was_blocked:2
ndmp:Session 01:node_send_was_blocked:0
ndmp:Session 01:dir_flush_calls:0
```

```

ndmp:Session 01:node_flush_calls:0
ndmp:Session 01:num_node_entries:2731
ndmp:Session 01:num_dir_entries:104362
ndmp:Session 01:num_dir_entries_2fh:104362
ndmp:Session 01:dir_entry_2fh_min_latency:0ms
ndmp:Session 01:dir_entry_2fh_max_latency:200ms
ndmp:Session 01:dir_entry_2fh_ave_latency:0ms
ndmp:Session 01:dir_entry_2fh_tot_latency:419ms
ndmp:Session 01:num_node_entries_2fh:2731
ndmp:Session 01:node_entry_2fh_min_latency:0ms
ndmp:Session 01:node_entry_2fh_max_latency:1ms
ndmp:Session 01:node_entry_2fh_ave_latency:0ms
ndmp:Session 01:node_entry_2fh_tot_latency:1ms
ndmp:Session 01:num_dir_entries_2ndmp:36
ndmp:Session 01:dir_entry_2ndmp_min_latency:19ms
ndmp:Session 01:dir_entry_2ndmp_max_latency:212ms
ndmp:Session 01:dir_entry_2ndmp_ave_latency:61ms
ndmp:Session 01:dir_entry_2ndmp_tot_latency:2598ms
ndmp:Session 01:num_node_entries_2ndmp:0
ndmp:Session 01:node_entry_2ndmp_min_latency:0ms
ndmp:Session 01:node_entry_2ndmp_max_latency:0ms
ndmp:Session 01:node_entry_2ndmp_ave_latency:0ms
ndmp:Session 01:node_entry_2ndmp_tot_latency:0ms
ndmp:Session 01:max_queue_depth:16
ndmp:Session 01:fh_queue_full_cnt:2

```

At the end of the backup session, the file history statistics is updated in the `etc/log/backup` file.

How to perform a dump backup using the CLI

You can perform a file system backup of your data to tape by using the dump command.

Next topics

[What the dump command syntax is](#) on page 94

[Where to enter the dump command](#) on page 96

[Specifying the backup level](#) on page 97

[Improving incremental dump performance](#) on page 98

[Updating the `/etc/dumpdates` file](#) on page 98

[Specifying a local tape device](#) on page 99

[Specifying a tape device on a remote storage system](#) on page 99

[Specifying the dump path](#) on page 101

[Specifying a list of files for backup](#) on page 102

[Backing up all data that is not in a `qtree`](#) on page 103

[Excluding specified files and directories](#) on page 104

[Omitting ACLs from a backup](#) on page 105

[Specifying a name for a backup](#) on page 106

[*Specifying a blocking factor*](#) on page 106

[*Specifying the tape file size*](#) on page 107

[*Appending backups to tapes*](#) on page 108

[*Verifying the files backed up by a dump command backup*](#) on page 108

[*Checking the status of a dump backup*](#) on page 109

[*Finding out whether a backup has to be restarted*](#) on page 111

[*How to get details about a specific backup*](#) on page 112

[*Restarting a dump command backup*](#) on page 113

[*Deleting restartable dump command backups*](#) on page 114

What the dump command syntax is

The Data ONTAP `dump` command has a defined syntax that consists of a set of options.

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 the 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.

The `dump` command syntax is as follows:

```
dump options parameters dump_path
```

The following list describes the various `dump` command options.

backup level	Level 0 is a full backup; levels 1 through 9 are for incremental backups.
A	Does not back up ACLs.
b	The blocking factor. Parameter: The number of 1-KB blocks in each write operation. For a storage system, the range is 4 through 64, and the default is 63.
B	Specifies the number of tape blocks to be written to a tape file before starting a new tape file. Parameter: The number of tape blocks in a tape file.
f	Specifies the tape device for the backup. (mandatory) Parameter: At least one tape device name as a parameter. Separate additional tape device names with commas.
l	Backs up only specific files and directories in the dump path. You must use the <code>n</code> option when using the <code>l</code> option.
n	Specifies to provide a name for the backup to be recorded in the <code>/etc/dumpdates</code> file. It takes a string as a parameter. It is required if you use the <code>l</code> option.

- Q** Backs up all data in the specified volume that does not reside in a qtree.
- u** Updates the `/etc/dumpdates` file. You must use this option if you plan to perform incremental backups in the future.
- X** Excludes specified files from the backup.
- Parameter: A string that specifies the exclusion prefixes or suffixes.

Note: Not all options are mandatory, and some do not have any parameters.

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 the options in any order.
- You must include a backup level and a tape file in the options.
- *parameters* can be one parameter or a list of parameters, each of which is associated with an option.
- List all parameters in the same order as their corresponding options.
- Separate each parameter with one or more spaces.
- If the parameter 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 could have changed the root volume.

Example of a dump command

```
dump 0fb rst0a 63 /vol/vol0/
```

The following list describes the elements of the command line:

- | | |
|-------------------|---|
| 0 | Does a full backup. |
| f | Specifies that a tape device is supplied in the command line. Its parameter is <code>rst0a</code> . |
| b | Specifies that a blocking factor is supplied in the command line. |
| 63 | The blocking factor. |
| /vol/vol0/ | The dump path. This command backs up to tape all files and directories in the <code>vol0</code> volume. |

Related concepts

[What increment chains are](#) on page 67

[How to specify tape devices for the backup](#) on page 69

Related tasks

- [Specifying the backup level](#)* on page 97
- [Omitting ACLs from a backup](#)* on page 105
- [Specifying a blocking factor](#)* on page 106
- [Specifying the tape file size](#)* on page 107
- [Specifying a list of files for backup](#)* on page 102
- [Specifying a name for a backup](#)* on page 106
- [Backing up all data that is not in a qtree](#)* on page 103
- [Updating the /etc/dumpdates file](#)* on page 98
- [Excluding specified files and directories](#)* on page 104

Where to enter 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.

Note: 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 using 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.

Specifying the backup level

You can specify a backup level for your `dump` command, based on which all files or only the most recently changed files are to be backed up to tape.

About this task

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.

Step

1. To specify the backup level, include the level number as an option. The range is 0 through 9.

Example

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

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

The following list describes the elements of the command line:

0	Does a full backup.
u	Records the backup in the <code>/etc/dumpdates</code> file.
f	Specifies that a tape device is supplied in the command line.
rst0a	The tape drive rewinds the tape.
/vol/vol1/users/ tom/specs	The directory to be backed up.

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

Related tasks

[Updating the `/etc/dumpdates` file](#) on page 98

[Backing up all data that is not in a `qtree`](#) on page 103

Improving incremental dump performance

Data ONTAP 7.3 and later provide an improved incremental dump performance, if you enable the `i2p` option on the volume to be backed up. You can accomplish this by setting the volume option `no_i2p` to `off`.

Step

1. To enable the `i2p` option on a particular volume, enter the following command:

```
vol options volume_name no_i2p off
```

volume_name is the name of the volume being backed up.

Note: By default, `i2p` is enabled.

Updating the `/etc/dumpdates` file

To keep track of the backups, update the `/etc/dumpdates` file.

Step

1. To update the `/etc/dumpdates` file, include the `u` option in the `dump` command line.

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
```

The following list describes the elements of the command line:

0	Does a full backup.
f	Specifies that a tape device is supplied in the command line.
u	Updates the <code>/etc/dumpdates</code> file.
rst0a	The tape drive rewinds the tape.
/vol/vol0	The directory to be backed up.

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
```

Related references

What the `/etc/dumpdates` file is on page 69

Specifying a local tape device

You can use a local tape device to back up the data.

Step

1. To specify local tape devices for a backup, use the `f` option and provide one or more tape devices, separated by commas, as a parameter to the `f` option.

Note: You cannot combine local and remote tape devices in a single command, and you can write to only one remote machine in a command.

Example

The following command specifies to write one tape file with one device:

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

The following list describes the elements of the command line:

0	Does a full backup.
f	Specifies that a tape device is supplied in the command line.
rst0a	The tape device.
/vol/vol0	The dump path.

Specifying a tape device on a remote storage system

You can use tape devices attached to a remote storage systems for a backup.

Step

1. To use a tape device on a remote storage system for the backup, use the `f` option and provide one or more tape devices, separated by commas, as a parameter to the `f` option.

Do not repeat the remote machine name.

Note: You cannot combine local and remote tape devices in a single command, and you can write to only one remote machine in a command.

Example

The following command performs a backup to a tape drive attached to a remote storage system named sales1. The tape drive does not rewind the tape.

```
dump 0f sales1:nrst0a /vol/vol1
```

The following list describes the elements of the command line:

0	Does a full backup.
f	Specifies that a tape device is supplied in the command line.
sales1	The name of the storage system that the tape drive is attached to.
nrst0a	The tape drive does not rewind the tape.
/vol/vol1	The volume to be backed up.

Next topics

Example: Tape drive attached to a remote storage system having an IPv6 address on page 100

Examples: Tape drive attached to a Solaris system on page 100

Example: Tape drive attached to a remote storage system having an IPv6 address

Starting with Data ONTAP 7.3.3, you can back up data to a tape device attached to a remote storage system having an IPv6 address.

This sample command performs a level 0 dump of the *voltest* volume to a remote tape device using an IPv6 address:

```
dump 0f [2001:0db8::10]:nrst01 /vol/voltest
```

In this example, *2001:0db8::10* indicates the IPv6 address of the storage system to which the remote tape device is attached.

Examples: Tape drive attached to a Solaris system

You can perform a backup to a 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 following list describes the elements of the command line:

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

The following command performs a backup to a tape drive on a Solaris system with a 2-GB limit. 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 following list describes the elements of the command line:

0	Does a full backup.
f	A tape device is supplied in the command line.
B	Specifies that the maximum tape file size allowed is supplied in the command line.
thompson	The name of the Solaris machine to which the tape drive is connected.
/dev/rmt/0n	The name of the remote tape drive.
2097151	The maximum tape file size allowed. This is equal to 2 GB.
/vol/vol1	The volume to be backed up.

Specifying the dump path

The dump path specifies one volume, qtree, or subtree to back up. (A subtree is a directory in a volume or qtree.)

About this task

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`

Step

1. To specify a single dump path, put the path name of the volume, qtree, or subtree that you want to back up at the end of the `dump` command.

Example

The following command contains the dump path `/vol/vol10`:

```
dump 0f rst0a /vol/vol10
```

The following list describes the elements of the command line:

0	Does a full backup.
f	Specifies that a tape device is supplied in the command line.

rst0a	The tape drive rewinds the tape.
/vol/vol0	The dump path.

Specifying a list of files for backup

You can back up some, but not all, subdirectories or files in the dump path using a single `dump` command.

About this task

You can specify for backup a list of one or more files. However, the files must all be in the same dump path. It is easier to specify a list rather than using a `dump` command for each subdirectory or file. It also helps you avoid starting multiple `dump` commands.

Steps

1. Use the `n` and `l` options in the command line.
2. Include a name for the backup as a parameter to the `n` option.
3. Put the path name of the volume, `qtree`, or `subtree` that you want to back up at the end of the `dump` command.
4. Enter the `dump` command line.
5. In response to prompts, enter each name as a path name relative to the dump path in the `dump` command.

Note: Do not specify a parent directory (`..`) or a directory that is a symbolic link.

6. To end the list, press the Enter key.

Example

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
```

The following list describes the elements of the command line:

0	Does a full backup.
u	Records the backup in the <code>/etc/dumpdates</code> file.
f	Specifies that a tape device is supplied in the command line.
n	Specifies that a name for the backup is supplied.

l	Specifies that the names of individual files and directories to be backed up will be entered interactively.
rst0a	The tape drive rewinds the tape.
user.1.3.5	The name of the backup.
/vol/vol1/home	The directory that contains the files to be backed up.

The output of the preceding dump command is as follows:

```
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
```

Backing up all data that is not in a qtree

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.

About this task

You cannot do incremental backups using this method.

Step

1. To back up all non-qtree data in a specified volume, use the **Q** option in the command line.

Example

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

```
dump 0fQ rst0a /vol/vol0
```

The following list describes the elements of the command line:

0	Does a full backup.
f	Specifies that a tape device is supplied in the command line.
Q	Excludes items in qtrees.
rst0a	The tape drive rewinds the tape.

<code>/vol/vol0</code>	The dump path.
------------------------	----------------

Excluding specified files and directories

You can exclude a list of files and directories from a backup. You can also specify a pattern based on which you can exclude files and directories from a backup. For example, you can exclude the files that end with `.core`.

About this task

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.

Steps

1. To exclude files from a backup, use the `x` option in the command line.
2. Include a string or comma-separated list of strings as a parameter for the `x` option.

Example

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 following list describes the elements of the command line:

0	Does a full backup.
u	Records the backup in the <code>/etc/dumpdates</code> file.
f	Specifies that a tape device is supplied in the command line.
X	Specifies that an exclude list is specified.
rst0a	The tape drive rewinds the tape.
tmp,*.o,core*,*backup*	The exclude list specifies files as follows: <ul style="list-style-type: none"> • <code>tmp</code> specifies that the file name is <code>tmp</code>.

- *.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

The volume to be backed up.

Omitting ACLs from a backup

You can omit ACLs from a backup. This provides a slight performance enhancement.

About this task

You omit ACLs in two situations:

- You plan to restore to a volume in an environment that does not support ACLs.
- You are backing up files or directories that do not contain ACLs.

Step

1. To omit ACLs from a backup, include the **A** option in the `dump` command line.

Note: This option does not take a parameter.

Example

The following command performs a level-0 backup of the `/vol/vol1` volume. The **A** option means that the backup does not include any ACL information.

```
dump 0Af rst0a /vol/vol1
```

The following list describes the elements of the command line:

0	Does a full backup.
A	Specifies not to back up ACLs.
f	Specifies that a tape device is supplied in the command line.
rst0a	The tape drive rewinds the tape.
/vol/vol1	The volume to be backed up.

Specifying a name for a backup

You can name a backup using the **n**. You can record this backup name in the `/etc/dumpdates` using the **u** option.

About this task

You specify a name for a backup 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.

Steps

1. To specify a name for the backup, include the **n** option in the `dump` command line.
2. Include a name for the backup as a parameter to the **n** option.

Example

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

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

The following list describes the elements of the command line:

0	Does a full backup.
f	Specifies that a tape device is supplied in the command line.
n	Specifies to name this backup.
rst0a	The tape drive rewinds the tape.
thisbackup	The name of the backup.
/vol/vol0	The dump path.

An output similar to the following appears in the `/etc/dumpdates` file:

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

Specifying a blocking factor

You can specify a blocking factor using the **b** option in the `dump` command line.

Step

1. To specify a blocking factor for a backup, include the **b** option in the `dump` command line.

Example

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 following list describes the elements of the command line:

0	Does a full backup.
u	Records the backup in the <code>/etc/dumpdates</code> file.
f	Specifies that a tape device is supplied in the command line.
b	Specifies that a blocking factor is provided.
rst0a	The tape drive rewinds the tape.
32	The blocking factor is 32, so writes 32 KB of data at a time.
/vol/vol1	The volume to be backed up.

Specifying the tape file size

You can specify the maximum size of the tape file in terms of tape blocks in a `dump` command. If you do a remote backup or plan to restore the backup on a system other than the storage system that was backed up, you might need to specify a tape file size.

About this task

Suppose you want the maximum tape file to be 2 GB; you must specify 2,097,151. This implies 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.

Some systems support only tape files of a limited size; for example, some Solaris systems do not support tape files larger than 2 GB.

Steps

1. To specify a tape file size, include the `B` option in the `dump` command line.
2. Include the tape file size, in KB, in the `dump` command as a parameter to the `B` 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/vol1
```

The following list describes the elements of the command line:

0	Does a full backup.
f	Specifies that a tape device is supplied in the command line.
B	Specifies that the file size is given in the command line.
rst0a	The tape drive rewinds the tape.
2097151	The file size is 2,097,151 KB.
/vol/vol1	The volume to be backed up.

Appending backups to tapes

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

Steps

1. To append a backup to a tape, move the tape to the desired location using the `mt` command.
2. Execute the `dump` command.

Attention: Use no-rewind device names to ensure that the tape is not rewound and that previous backups are not overwritten.

Related references

[Controlling tape drives](#) on page 35

Verifying the files backed up by a dump command backup

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

Steps

1. From your client, preserve the output to the console by using a utility such as a script.
2. List all the files in a backup by entering the following command:

```
restore tf rst0a
```

3. Compare the list to what you intended to back up.
4. For more detailed verification, use the `N` option of the `restore` command.

Checking the status of a dump backup

During a lengthy dump session, you are advised to monitor the progress and check the status of the session. This helps you to determine if the backup is proceeding as expected.

Step

1. To check the status of a dump command, enter the following command:

```
stat show dump
```

The output of the `stat show dump` command displays the following statistics about the data set and progress:

- The number of directories that will be dumped
- The number of files that will be dumped
- The number of NT STREAMS
- The number of ACLs
- The average directory size
- The average file size

The following are the progress statistics listed by the command:

- The number of directories dumped in Phase 3
- The amount of directory data, in KB, currently written to tape in Phase 3
- The number of inodes dumped in Phase 4
- The amount of inode data, in KB, currently written to tape in Phase 4

Example

The following is an example of the `stat show dump` command output:

```
filer1>stat show dump
dump:id_0:p1-ino:6097
dump:id_0:p1-dir:412
dump:id_0:p1-str-ino:0
dump:id_0:p1-str-dir:0
dump:id_0:p1-acl:0
dump:id_0:p3-dir:413
dump:id_0:p3-write:487
dump:id_0:p4-ino:1962
dump:id_0:p4-write:135043
```

Statistics shown in the preceding example are as follows:

- `id_0` is the instance name for dump statistics. The number part of the instance name specifies the dump ID.
- `p1-ino` shows the total number of regular inodes that will be dumped.
- `p1-dir` shows the total number of directory inodes that will be dumped.
- `p1-str-ino` shows the total number of NT stream inodes that will be dumped.

- p1-str-dir shows the total number of NT stream directories that will be dumped.
- p1-acl shows the total number of ACL inodes that will be dumped.
- p3-dir shows the total number of directory inodes that have been written in Phase 3.
- p3-write shows the total number of kilobytes (KB) of directory tape data that have been written in Phase 3.
- p4-ino shows the total number of inodes that have been dumped in Phase 4.
- p4-write shows the total number of kilobytes (KB) of inode tape data that have been written in Phase 4.

The following is an example of statistics shown in the backup log:

```
dmp ... /vol/compat/(3) Start (Level 0)
...
dmp ... /vol/compat/(3) End (126 MB)
dmp ... /vol/compat/(3) Log_msg (reg inodes: 1574 other inodes: 1061 dirs: 200 nt dirs: 54 nt inodes:
204 acis: 49)
dmp ... /vol/compat/(3) Log_msg (Phase 1 time: 261)
dmp ... /vol/compat/(3) Log_msg (Phase 3: directories dumped: 255)
dmp ... /vol/compat/(3) Log_msg (Phase 3: waf1 directory blocks read: 291)
dmp ... /vol/compat/(3) Log_msg (Phase 3: average waf1 directory blocks per inode: 1)
dmp ... /vol/compat/(3) Log_msg (Phase 3: average tape blocks per inode: 2)
dmp ... /vol/compat/(3) Log_msg (Phase 3 throughput (MB sec): read 0 write 0)
dmp ... /vol/compat/(3) Log_msg (Percent of phase3 time spent for: reading inos 0% dumping ino 93%)
dmp ... /vol/compat/(3) Log_msg (Percent of phase3 dump time spent for: convert-waf1-dirs 4% lev0-ra 1%)
dmp ... /vol/compat/(3) Log_msg (Phase 3 averages (usec): waf1 load buf time 27 level 0 ra time 62)
dmp ... /vol/compat/(3) Log_msg (Phase 4: inodes dumped: 2839)
dmp ... /vol/compat/(3) Log_msg (Phase 4: waf1 data blocks read: 55502)
dmp ... /vol/compat/(3) Log_msg (Phase 4: average waf1 data blocks per inode: 19)
dmp ... /vol/compat/(3) Log_msg (Phase 4: average tape data blocks per inode: 75)
dmp ... /vol/compat/(3) Log_msg (Phase 4 throughput (MB sec): read 51 write 50)
dmp ... /vol/compat/(3) Log_msg (Percent of phase4 time spent for: reading inos 3% dumping inos 94%)
dmp ... /vol/compat/(3) Log_msg (Tape write times (msec): average: 0 max: 1863)
dmp ... /vol/compat/(3) Log_msg (Tape changes: 1)
```

Statistics shown in the backup log example are as follows:

- reg inodes, other inodes, dirs, nt dirs, nt inodes, acis—The total number of regular inodes, other inodes such as symlinks or char devices, directory inodes, NT STREAMS inodes, and ACL inodes that will be dumped.
- Phase 3: directories dumped
—The total number of directory inodes dumped in Phase 3.
- Phase 3: waf1 directory blocks read
—The total number of WAFL directory blocks read.
- Phase 3: average waf1 directory block per inode
—The average size of directories that were dumped.
- Phase 3: average tape blocks per inode
—The average number of dump tape blocks (1K) for each directory inode.
- Phase 3 throughput (MB sec)
—The read and write throughputs, in MBps, for Phase 3.
- Percent of phase3 time spent for: reading inos and dumping inos
—An indication of where time is spent in Phase 3.
- Percent of phase3 dump time spent for: convert-waf1-dirs and lev0-ra
—An indication of where time is spent in Phase 3.

- Phase 3 averages (usec): waf1 load buf time and level 0 ra time
—An indication of how long it takes to read a WAFL directory block and how long it took to read ahead for these blocks.
- Phase 4: inodes dumped
—The total number of inodes dumped in Phase 4.
- Phase 4: waf1 data blocks read
—The total number of WAFL data blocks read.
- Phase 4: average waf1 data blocks per inode
—An indication of the average size of files that were dumped.
- Phase 4: average tape data blocks per inode
—The average number of dump tape blocks (1K) for each inode.
- Phase 4 throughput (MB sec)
—The read and write throughputs, in MBps, for Phase 4.
- Percent of phase4 time spent for: reading inos and dumping inos
—An indication of where time is spent in Phase 4.
- Percent of phase4 dump time spent for:waf1 read iovec and lev0-ra
—An indication of where time is spent in Phase 4.
- Phase 4 averages (usec): waf1 read iovec time and level 0 ra time
—An indication of how long it takes to read a file block and how long it took to read ahead for these blocks.
- Tape write times (msec): average and max
—An indication of how long it took to write out a tape block.
- Tape changes
—The number of tape changes.

Finding out whether a backup has to be restarted

To find out whether a backup initiated by the `dump` command is proceeding as expected or has aborted, you can run the `backup status` command.

Step

1. To know the status of a backup, enter the following command:

backup status

Following is an example of the `backup status` command's output:

```
filer1>backup status
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	urstla	Nov 28 00:22	1	/vol/vol0
6	RESTARTABLE	dump		Nov 27 00:22	3	/vol/vol1

The following list describes the elements of the dump table:

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 invocation of dump: CLI or NDMP.
Device	The current device to which the dump is writing.
Start Date	The 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, you can supply the dump ID at the end of the backup status command.

Following are the examples of the backup status command.

Example 1

```
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 following list describes the output of the command:

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

Example 2

```
filer> backup status 2
State:      RESTARTABLE                           Type:      ndmp
Path:       /vol/vol1                             Level:     0
Snapshot:   filer(0101184236)_vol1_filer_svp-dst.0
```



```
Snapshot:      snapshot_for_backup.9 [Dec 27 00:41]
Options:       b=63, X
Devices:       [none]
Completed:     1 tapefile(s)
Last Update:   Thu Dec 27 00:41:23 2007
```

The preceding example displays the following additional information:

Snapshot The Snapshot copies of the path that is being backed up.

Restarting a dump command backup

To restart an aborted backup, you must use the `R` option in the `dump` command.

Step

1. To restart a dump process that has been shown to be restartable, enter the following command from the storage system:

```
dump R[comma-separated-device-list] {path | ID}
```

`f` is an option that enables you to supply a device list.

`comma-separated-device-list` 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 in 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 `rst2a`: `rst0a,rst1a,rst2a,rst3a,rst4a`.

The command will use `rst3a,rst4a` to complete the backup. However, if the original device list contained any non-rewinding (`nrst`) devices or any devices not supported, users are required to supply a new device list at the restart of the dump.

`path` is the path that is listed in the dump table (the output of the `backup status` 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 backup.

`ID` is the unique ID displayed by the `backup status` command.

You can use either `path` or `ID` 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.

Related tasks

[Checking the status of a dump backup](#) on page 109

Deleting restartable dump command backups

You can delete a restartable dump using the dump ID.

Step

1. To delete a restartable backup, enter the following command:

```
backup terminate ID
```

ID is the unique ID in the dump table that the `backup status` command displays.

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

How to perform a dump restore using the CLI

You can use the `restore` command to restore data backed up to tape using the dump backup.

Next topics

[Restore command syntax](#) on page 115

[What restore types are](#) on page 115

[What modifiers are](#) on page 116

[Where to enter the restore command](#) on page 117

[Executing a restore command](#) on page 117

[Restoring incremental backups](#) on page 118

[Restoring each volume backed up as separate subtrees or qtrees](#) on page 118

[Restoring individual files and directories](#) on page 119

[Specifying a full restore](#) on page 119

[What a table-of-contents restore is](#) on page 120

[Specifying a resume restore](#) on page 121

[Specifying tape devices in the restore command](#) on page 122

[Specifying a single tape file on a multifile tape](#) on page 123

[Specifying the restore destination](#) on page 124

[Specifying the blocking factor during restore](#) on page 124

[Displaying detailed status output](#) on page 125

[Ignoring inode limitations](#) on page 126

[Specifying automatic confirmations](#) on page 127

[Specifying no ACLs to be restored](#) on page 127

[Specifying not to restore qtree information](#) on page 128

[Specifying a test restore](#) on page 129

[Restore examples: Restoring using a remote tape drive](#) on page 129

[Restore examples: Multiple tape restores](#) on page 131

Restore command syntax

The `restore` command consists of a set of options that include the restore types and the modifiers.

There are a set of rules that you have to follow when you enter the `restore` command:

- Specify only one restore type.
- Specify multiple options without intervening spaces.
- Enter the parameters for each option in the order that you specify the options. Separate each parameter 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.

The `restore` command syntax is as follows:

```
restore options [parameters] [files ...]
```

options can be one restore type with modifiers.

What restore types are

A restore type specifies the type of restore you are performing.

For a restore from tape, you must specify only one restore type. The following table summarizes the restore types.

Restore type	Description	Option
Restart	Restarts data recovery after an interruption.	R
Qtree table of contents	Lists qtree names and qtree information in a restore.	T
Full	Rebuilds the file system or subtree. If you are applying incrementals, you must specify this option.	r
File table of contents	Lists file names in a restore.	t
File	Extracts an individual file or subtree from the backup.	x

Related tasks

[Specifying a resume restore](#) on page 121

[Specifying table-of-contents restores](#) on page 121

[Specifying a full restore](#) on page 119

[Restoring individual files and directories](#) on page 119

What modifiers are

Modifiers specify optional actions.

The following list describes the various modifiers:

A Specifies not to restore ACLs.

D Specifies the directory into which the files are restored.

Parameter: The directory into which you are restoring files. Without a parameter, the files are restored to the directory from which they were backed up.

F Forces restore to continue regardless of inode limitations.

N Reads backup tapes without writing to the storage system.

Q Ignores qtree information.

b Specifies the blocking factor.

Parameter: The blocking factor that you used in the backup that you are restoring

f Specifies the tape device for each tape file.

Parameter: The name of one or more tape devices, separated by commas

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.

Parameter: The tape file number

v Specifies that the restore will display the inode number of each file restored.

y Specifies that the restore will not prompt the user if it encounters an error.

Related tasks

[Specifying no ACLs to be restored](#) on page 127

[Specifying the restore destination](#) on page 124

[Ignoring inode limitations](#) on page 126

[Specifying a test restore](#) on page 129

[Specifying not to restore qtree information](#) on page 128

[Specifying the blocking factor during restore](#) on page 124

[Specifying tape devices in the restore command](#) on page 122

[Specifying a single tape file on a multifile tape](#) on page 123

[Displaying detailed status output](#) on page 125

[Specifying automatic confirmations](#) on page 127

Where to enter the restore 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 the following 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. However, if you enter the `restore` command on the console, it could be terminated by pressing Ctrl-C 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.

Executing a restore command

You have to perform a series of steps to execute a `restore` command.

Steps

1. Place the tape containing the first tape file of the backup in the tape drive that you specify.
2. Enter the `restore` 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.

Restoring incremental backups

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.

About this task

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 Data ONTAP 6.2, the restore will fail. This is because there is a formatting code change between the two code releases. You need to run the full backup again after you have upgraded to Data ONTAP 6.2 or later.

Steps

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.

Attention: During an incremental restore operation, a temporary directory labeled `.restore_do_not_touch_XXXXXXX` 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. After all the incremental restores are completed, delete the `restore_symboltable` file from the root of the destination directory.

Related concepts

What increment chains are on page 67

Related tasks

Specifying the backup level on page 97

Restoring each volume backed up as separate subtrees or qtrees

You can restore an entire storage system even if you used separate `dump` commands to back up files, directories, and qtrees that make up each volume.

Steps

1. To restore each volume backed up as separate subtrees or qtrees, create the desired volumes.
2. Restore each backup to the appropriate volume.

Restoring individual files and directories

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

Steps

1. Use the `x` option in the `restore` 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.

Note: If you do not have a path in the command line, the `restore` command restores all data 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
```

The following list describes the elements of the command line:

X	Restores a specified file.
f	Specifies that a tape device is supplied in the command line.
rst0a	The tape device.
/src	The directory to be restored.

Specifying a full restore

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

Step

1. To specify a full restore, use the `r` option in the `restore` command line.

Example

The following command performs a full restore to the original location.

```
restore rf rst0a
```

The following list describes the elements of the command line:

r	Performs a full restore.
----------	--------------------------

f	Specifies that a tape device is supplied in the command line.
rst0a	The tape device.

What a table-of-contents restore is

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 the 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.

Why Remote Shell is preferred 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. Also, client consoles are more flexible and enable you to save the output.

Also, 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 types of tables of contents: file and qtree. These are explained in the following table.

Type	Description	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

You cannot combine the two types in a single command.

Specifying table-of-contents restores

Use the `t` or `T` option in the restore command to specify a table-of-contents restore.

Step

1. To specify a table-of-contents restore, use the `T` or `t` option in the `restore` command line, with files as parameter. If there is no parameter, the entire content of a backup is listed.

Example

The following command lists all files in a backup:

```
restore tf rst0a
```

The following list describes the elements of this command line:

t Lists all the files.

Note: Option `T` lists all qtree names.

f Specifies that a tape device is supplied in the command line.

rst0a The tape device.

Specifying a resume restore

If an entire tape file restore is stopped, you can resume the restore and avoid restoring again what has already been restored. However, there are some restrictions on this operation.

About this task

You must consider the following restrictions on resuming a restore:

- 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.
```

Steps

1. In the `restore` command line, use the `R` option first instead of the `r` option. It does not take a parameter.

2. Enter the rest of the same `restore` command that was interrupted. However, include only the tape files that were not restored.
3. Follow the prompts.

Example

The following command resumes a restore:

```
restore Rf rst0a
```

The following list describes the elements of the command line:

R	Resumes a restore.
f	Specifies that a tape device is supplied in the command line.
rst0a	The tape device.

Specifying tape devices in the restore command

When you perform a backup, you specified one or more tape devices. The files written by these devices can be on one or more tapes. When restoring, you have to list the tape devices in the same order that you used in the backup.

About this task

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.

Steps

1. To specify the tape devices for restores, use the `f` option in the `restore` command line.
2. List the tape devices as a parameter to the `f` option in the same order that you used in the backup. Separate multiple tape devices with a comma.

Note: If you do not specify at least one tape device, the `restore` command terminates.

The `restore` command restores from tape files consecutively, using the tape devices in the order that they appear in the command line.

Example

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

```
restore rf rst0a
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
rst0a	The tape device.

Specifying a single tape file on a multifile tape

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. You do this by moving the tape to the beginning of the file that is to be restored.

Steps

1. Use the **f** option in the `restore` command line.
2. Use the same tape compression type as a parameter to the **f** option that you used in the backup.
3. Use the **s** option in the `restore` command line to select the appropriate backup.
4. Include the relative position of the tape file that you are restoring as a parameter to the **s** option in the command line.

Note: Count the relative position from the current tape position. It is best to rewind the tape and start from its beginning.

Example

From a tape that has been rewound, the following command restores the third tape file from the beginning of that tape. It then rewinds the tape.

```
restore rfs rst0a 3
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
s	Selects a tape file.
rst0a	The tape device.
3	Specifies to use the third tape file.

Specifying 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 restoring the backed up data to a different location.

About this task

For example, if you created a backup and then installed multiple volumes on the storage system, you might specify a different volume or directory when you perform a restore.

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

Note: You should specify a restore destination even if you are restoring to the same destination from which you backed up. This ensures the files are restored where you want them to go and are traceable to that location.

Steps

1. To specify the restore destination, use the `D` option in the `restore` command line.
2. Include the absolute path name of the restore destination as a parameter to the `D` option.

Example

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

```
restore rfd rst0a /vol/destination
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
D	Specifies that a destination is supplied in the command line.
rst0a	The tape device.
/vol/destination	The destination is the <code>/vol/destination</code> volume.

Specifying the blocking factor during restore

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.

Steps

1. To specify the blocking factor, use the `b` option in the `restore` command line.

2. Include the blocking factor as a parameter to that option.

Example

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

```
restore rfb rst0a 63 /vol/destination
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
b	Specifies that a blocking factor is supplied in the command line.
rst0a	The tape device.
63	The blocking factor.
/vol/destination	The restore destination.

Displaying 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 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.

Step

1. To get status information about each file recovered, use the `v` option in the `restore` command line.

Note: This option does not take a parameter.

Example

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

```
restore rfv rst0a
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
v	Produces information about each file recovered.

rst0a The tape device.

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

Ignoring inode limitations

If you are sure that the restore consists mostly of files to be updated rather than new files, you can instruct the storage system to ignore the inode limitations.

About this task

What inodes are: 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 *Data ONTAP Storage Management Guide*.

How the `restore` command handles inodes: 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 instruct the storage system to ignore the calculations of the `restore` command.

Note: During a restore, if the inode count exceeds the maximum number of inodes allowed, the restore is terminated.

Step

1. To specify a restore to ignore inode limitations, use the `F` option in the `restore` command line.

Note: This option does not take a parameter.

Example

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

```
restore rfF rst0a
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
F	Specifies to ignore inode limitations.
rst0a	The tape device.

Specifying automatic confirmations

An automatic confirmation automatically answers all restore questions with a "yes." You usually use this mode on restores that are run using a Remote Shell connection.

About this task

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 unclear drives.

Attention: This option is not advisable for critical restores because it can cause silent failure.

Step

1. To specify automatic confirmations, use the `y` option in the `restore` command line.

Note: This option does not take a parameter.

Example

The following command restores a backup with automatic confirmations:

```
restore rfy rst0a
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
y	Specifies automatic confirmations.
rst0a	The tape device.

Specifying no ACLs to be restored

You can exclude ACLs from a restore. This provides a slight performance enhancement.

About this task

You can exclude ACLs in two situations:

- You plan to restore to an environment that does not support ACLs.
- The backup has no files or directories that contain ACLs.

Step

1. To exclude ACLs from a restore, include the **A** option in the `restore` command line.

Note: This option does not take a parameter.

Example

The following command restores a backup, but does not restore ACLs:

```
restore rfA rst0a
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
A	Specifies not to restore ACLs.
rst0a	The tape device.

Specifying not to restore qtree information

You can omit qtree information from a restore. In such cases, the qtrees are restored as ordinary directories.

Step

1. To omit qtree information from a restore, include the **Q** option in the `restore` command line.

Note: This option does not take a parameter.

Example

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

```
restore rfQ rst0a
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
Q	Specifies not to restore qtrees.
rst0a	The tape device.

Specifying a test restore

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

About this task

You can 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 ensure that a block size works, if the block size is unknown

Note: Because a test restore depends on the speed of reading from tape, it takes almost the same time as an actual restore.

Step

1. To specify a test restore, include the `N` option in the `restore` command line.

Note: This option does not take a parameter.

Example

The following command performs a test restore of a backup:

```
restore rfN rst0a
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
N	Specifies a test restore.
rst0a	The tape device.

Restore examples: Restoring using a remote tape drive

You can perform a storage system restore using a tape drive attached to a remote storage system or a tape drive attached to a Solaris system.

Example of a storage system restore using a tape drive attached to a remote storage system

Assume you have performed a backup using the following `dump` command:

```
dump 0f sales1:rst0a /vol/vol1
```

The following command performs a restore from a tape drive attached to a remote storage system named sales1. The tape drive then rewinds the tape.

```
restore rf sales1:rst0a
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
sales1	The name of the storage system.
rst0a	The restore is done using the rst0a tape device.

Example of a storage system restore using a tape drive attached to a Solaris system

Assume you have performed a backup using the following `dump` command:

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

The following command performs a restore from a tape drive on a Solaris system:

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

The following list describes the elements of the command line:

r	Performs a full restore.
f	Indicates that a tape device is supplied in the command line.
ritchie	The name of the Solaris machine to which the tape drive is connected.
/dev/rmt/0	The name of the tape device.

Example of restoring data from a tape drive attached to a remote storage system having an IPv6 address

The following sample command restores data from a tape device attached to a storage system having an IPv6 address. Data is restored to the `voltest` volume.

```
restore rFD [2001:0db8::10]:nrst01 /vol/voltest
```

Restore examples: Multiple tape restores

There are different types of multiple tape restores, such as multiple tapes on a single-tape drive, multiple tapes on two single-tape drives, and multiple tapes on a tape library.

Example of restore from multiple tapes on a single-tape drive

Assume you have performed a backup using the following `dump` command:

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

The following 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.

```
restore rf rst0a
```

The following list describes the elements of the command line:

- r** Performs a full restore.
- f** Specifies that 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.

Example of restore from multiple tapes on two single-tape drives

Assume you have performed a backup using the following `dump` command:

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

The first tape is in tape drive 0 and the second tape is in tape drive 1.

The following 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.

```
restore rf rst0a,rst1a
```

The following list describes the elements of the command line:

- r** Performs a full restore.
- f** Specifies that a tape device is supplied in the command line.
- rst0a** The restore is done using the `rst0a` tape device for the first tape.
- rst1a** The restore is done using the `rst1a` tape device for the second tape.

Example of a restore from multiple tapes on a tape library

Assume you have performed a backup using the following `dump` command:

```
dump of urst0a,urst0a /vol/vol1
```

The following command restores the /vol/vol1 volume from the two tapes used to back it up. It unloads the first tape and loads the second tape.

```
restore rf urst0a,urst0a /vol/vol
```

The following list describes the elements of the command line:

r	Performs a full restore.
f	Specifies that a tape device is supplied in the command line.
urst0a, urst0a	The tape drive unloads and loads each tape.

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. By default, event logging is set to `On`.

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 an active/active configuration, the set of backup log files for the takeover storage system remains separate from the backup log files for the failed storage system.

Next topics

[What the dump and restore event log message format is](#) on page 133

[Enabling or disabling event logging](#) on page 137

What the dump and restore event log message format is

For each dump and restore event, a message is written to the backup log file.

The format of the dump and restore event log message is as follows:

```
type timestamp identifier event (event_info)
```

The following list describes the fields in the event log message format.

- Each log message begins with one of the type indicators described in the following table.

Type	Description
log	Logging event

Type	Description
dmp	Dump event
rst	Restore event

- *timestamp* shows the date and time of the event.
- 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.

Next topics

[What logging events are](#) on page 134

[What dump events are](#) on page 134

[What restore events are](#) on page 136

What logging events are

The event field of a message that begins with a log specifies the beginning of a logging or the end of a logging.

It contains one of the events shown in the following table.

Event	Description
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.

What dump events are

The event field for a dump event contains an event type followed by event-specific information within parentheses.

The following table describes the events, their descriptions, and the related event information that might be recorded for a dump operation.

Event	Description	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

Event	Description	Event information
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	A dump is entering a new processing phase	The new phase name
Error	A dump has encountered an unexpected event	Error message
Snapshot	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 of a dump output

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/vol10/(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.

What restore events are

The event field for a restore event contains an event type followed by event-specific information in parentheses.

The following table provides information about the events, their descriptions, and the related event information that can be recorded for a restore operation.

Event	Description	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

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

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

You can turn the event logging on or off.

Step

1. To enable or disable event logging, enter the following command:

```
options backup.log.enable {on | off}
```

`on` turns event logging on.

`off` turns event logging off.

Note: Event logging is turned on by default.

Index

/etc/tape_config files 40

A

ACLs (access control lists)
 excluding from tape restores 127
 including in tape backups 66
 aliases, tape
 on multiple storage systems 31
 appending backups to tapes 108

B

backup and restore using NDMP services
 dump 78
 backups
 copying from tape with restore command 117
 creating snapshot_for_backup file for 66
 nonconsecutive, contents of 67
 parallel 94
 backups to tape (dump command)
 benefits of entering at console 96
 benefits of using Remote Shell 96
 estimating tapes required for 73
 rules for excluding files from 104
 syntax 94
 unattended 73
 where to enter the command 96
 backups to tape (dump)
 simultaneous dump 66

C

commands
 storage alias (displays tape aliases) 29
 storage show (displays tape drive information) 29
 compression type
 specifying in restores from tape 122
 considerations
 before using the dump command 73

D

DAR functionality 87
 data backup to tape
 using the dump engine 65

dump backup
 CIFS attributes, not backed up by 66
 decreasing tape backup time 72
 leaving volumes online for 66
 minimizing backup time and data loss 71
 use of Snapshot copies 66
 using for backups
 minimizing tapes used for 72
 dump command
 backup levels, defined 97
 deleting a restartable dump 114
 maximum tape blocks per tape file 107
 order of tape devices specified by 69
 specifying a blocking factor 106
 specifying a dump path 101
 specifying backup names 106
 specifying files and directories 102
 specifying local tape device names 99
 specifying tape blocks per tape file 107
 specifying to omit ACLs 105
 using for backups
 labeling backup tapes 72
 dump engine
 Data ONTAP version compatibility 75
 data that can be restored 75
 dump events 134
 dump restart command (restarts interrupted backup) 113
 dumpdates file
 principles applying to dumpdates file 69
 purpose 69
 reasons to update 69

E

emulating a qualified tape drive 44
 environment variables
 ACL_START 78
 BASE_DATE 78
 DATA_BLOCK_SIZE 78
 DEBUG 78
 DIRECT 78
 DMP_NAME 78
 DUMP_DATE 78
 ENHANCED_DAR_ENABLED 78
 EXCLUDE 78
 EXTRACT 78

EXTRACT_ACL 78
 FILESYSTEM 78
 FORCE 78
 HIST 78
 IGNORE_CTIME 78
 IGNORE_QTREES 78
 LEVEL 78
 LIST 78
 LIST_QTREES 78
 MULTI_SUBTREE_NAMES 78
 NDMP_UNICODE_FH 78
 NDMP_VERSION 78
 NO_ACLS 78
 NON_QUOTA_TREE 78
 NOWRITE 78
 RECOVER_FULL_PATH 78
 RECURSIVE 78
 UPDATE 78
 uses 64

error messages related to, example 42

event log files

effect of takeover mode on 133
 viewing, reasons for 133

event log messages

dump and restore
 event field 133

format

dump and restore 133

identifier field

dump and restore 133

start and stop logging events 134

timestamp field

dump and restore 133

type field

dump and restore 133

event logging

enabling or disabling 137

examples

event log

dump 134

restore 136

examples of ndmcopy command

example of the ndmcopy command where the
 address modes are explicitly forced to
 IPv6 90

migrating data from a source path on a remote host
 to a destination path on another remote
 host 90

migrating data from a source path on remote host to
 a destination path on the local storage
 system 90

migrating data from a source path to a different
 destination path on a remote host 90

migrating data from a source path to a different
 destination path on the same storage
 system 90

overwriting the /etc directory during the root
 volume migration 90

F

files

backing up using dump 66

excluding data from backup 104

excluding from dump command 104

I

increment chains, of backups 67

inodes

ignoring limits when restoring files 126

L

levels of backups 97

log files

for backup and restore events 133

LUN (logical unit number) 28

M

manage NDMP

how to 52

maximum number of simultaneous tape drives 24

mt command syntax 35

N

NDMP

advantages of 47

copying with local tool (ndmcopy) 88

debug log file, displaying 57

debug log message, displaying 57

debug messages 56

defined 47

disabling preferred network interface 53

displaying file history performance 92

- enabling or disabling service (ndmpd on/off) 52
- killing sessions (ndmpd kill command) 56
- preparing a storage system for basic management 62
- session information
 - displaying detailed status (ndmpd probe) 54
 - displaying status (ndmpd status command) 54
- setting preferred network interface 53
- showing max version supported (ndmpd version command) 59
- tape devices used with 61, 62
- using with tape libraries 61, 62
- version, need to specify 58

NDMP commands

- ndmp on 52
- ndmpcopy (uses local copy tool) 88
- ndmpd debug (outputs debug log file) 57
- ndmpd kill (terminates NDMP session) 56
- ndmpd on/off (enabling or disabling service) 52
- ndmpd probe (displays detailed status) 54
- ndmpd status (displays status) 54
- ndmpd version (shows max version supported) 59

ndmpcopy command

- examples 90

O

options

- backup.log.enable (turns event logging on or off) 137
- ndmp.preferred_interface (sets preferred network) 53
- ndmpd.offset_map.enable 87

P

physical path names (PPNs)

- format 27

Q

qtrees

- excluding data from backup 103
- omitting data from dump command 103

- qualified tape drives, defined 40

R

remote hosts 22

Remote Shell

- using to display table of contents for restores from tape 120

restartable backups

- deleting automatically 114
- qualifications 74

restore

- incremental backups 118

restore command

- disk space required for 76, 77
- information required for using 76, 77
- options 116
- restoring individual files 119
- specifying a full restore 119
- specifying a resume restore 121
- specifying a single tape file on a multifile tape 123
- specifying a test restore 129
- specifying automatic confirmations 127
- specifying no qtree information 128
- specifying table-of-contents restore 121
- specifying tape devices 122
- specifying the blocking factor 124
- specifying to exclude ACLs 127
- specifying to ignore inode limitations 126
- syntax 115
- types of restores 115
- using with Remote Shell 117

restore command, executing 117

restore events 136

restoring data from tapes 75

rewind type, specifying for tape devices 22

rules

- for restore command 115
- for specifying a resume restore 121

S

storage (aliasing) commands

- storage alias (assigns tape alias) 30
- storage unalias (removes tape alias) 31

storage systems

- adding Fiber Channel-attached drives dynamically 32
- displaying information about tape drive connections to 34

subtrees, defined 101

- sysconfig -m command (shows information about tape medium changers) 34

- sysconfig -v command (shows tape drive connections to storage system) 34

T

- tape aliases
 - definition 26
- tape configuration files
 - how the storage system uses 42
 - what are 40
- tape devices
 - local, defined 21
 - on remote Solaris systems 21
 - remote, defined 21
 - specifying compression type of 22
 - what are 21
- tape drives
 - in tape libraries, listing qualified 40
 - nonqualified
 - displaying information 43
 - using 42
 - showing status (mt -status) 39

- tape medium changers, displaying information
 - about 34
 - unloading tape after rewind (mt -offline) 38
- tape libraries
 - showing names assigned to 62
- tape reservations
 - what are 45
- tape restores
 - displaying a table of contents (files) 120
 - displaying detailed status output 125
 - running a test restore 129
 - specifying a restore destination 124
 - specifying automatic confirmations 127
 - specifying tape devices 122

W

- worldwide names (WWNs) 28