



Data ONTAP[®] 8.2

Data Protection Tape Backup and Recovery Guide

For 7-Mode



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1(408) 822-6000
Fax: +1(408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-07983_A0
May 2013

Contents

| | |
|--|-----------|
| Data protection using tape | 12 |
| Advantages and disadvantages of tape backup | 12 |
| Types of tape backup supported by Data ONTAP | 12 |
| How to initiate a dump or SMTape backup | 13 |
| Differences between dump backup and SMTape backup | 13 |
| Considerations before choosing a tape backup method | 14 |
| How online migration affects tape backup | 14 |
| How volume move operations affect tape backup | 15 |
| Tape drive management | 16 |
| What tape devices are | 16 |
| Types of tape devices | 16 |
| Tape device name format | 17 |
| Supported number of simultaneous tape devices | 18 |
| Displaying tape device statistics | 19 |
| Displaying supported tape devices | 19 |
| What assigning tape aliases is | 21 |
| What physical path names are | 22 |
| What serial numbers are | 23 |
| Displaying existing aliases of tape drives | 23 |
| Displaying information about tape drives and medium changers | 24 |
| Assigning tape aliases | 25 |
| Removing tape aliases | 25 |
| Propagating tape aliases to multiple storage systems | 26 |
| Considerations when configuring multipath tape access | 26 |
| How to add tape drives and libraries to storage systems | 27 |
| How to display tape drive and tape library information | 27 |
| Displaying information about tape drives | 27 |
| Displaying information about tape medium changers | 28 |
| Displaying information about tape drive connections to the storage system | 28 |
| Controlling tape drives | 29 |
| Moving a tape to the end of data | 30 |

| | |
|--|-----------|
| Moving forward to a file | 31 |
| Moving backward to the beginning of a file | 31 |
| Rewinding a tape | 32 |
| Taking a tape drive offline | 32 |
| Displaying status information | 33 |
| What qualified tape drives are | 34 |
| Format of the tape configuration file | 34 |
| How the storage system qualifies a new tape drive dynamically | 36 |
| How to use a nonqualified tape drive | 37 |
| Displaying information about nonqualified tape drives | 37 |
| Tape drive information required for emulation | 38 |
| Emulating a qualified tape drive | 38 |
| What tape reservations are | 39 |
| Enabling tape reservations | 40 |
| Disabling tape reservations | 40 |
| NDMP management | 41 |
| Advantages of NDMP | 41 |
| What NDMP security is | 41 |
| Specifying NDMP access by host or interface | 42 |
| Specifying the NDMP authentication type | 43 |
| Enabling or disabling NDMP connection logging | 43 |
| Specifying the NDMP password length | 44 |
| Generating an NDMP-specific password for non-root administrators | 45 |
| How to manage NDMP | 45 |
| Enabling and disabling NDMP services | 45 |
| Specifying a preferred network interface | 46 |
| Designating the range of ports for NDMP data connections | 47 |
| Turning off a data connection specification | 48 |
| Displaying the general status information about NDMP sessions | 48 |
| Displaying detailed NDMP session information | 48 |
| Optimizing NDMP communication performance | 50 |
| Terminating an NDMP session | 51 |
| Displaying the NDMP version | 51 |
| NDMP options | 51 |
| NDMP extensions supported by Data ONTAP | 53 |
| Tape backup using NDMP services | 53 |

| | |
|---|-----------|
| Common NDMP tape backup topologies | 53 |
| Considerations when using NDMP | 54 |
| Tape devices and configurations you can use with the storage system | 55 |
| Preparing for basic NDMP backup application management | 56 |
| What environment variables do | 57 |
| Data backup to tape using the dump engine | 58 |
| How a dump backup works | 58 |
| What the dump engine backs up | 59 |
| What increment chains are | 60 |
| How to specify tape devices for the backup | 61 |
| What the /etc/dumpdates file is | 61 |
| What the blocking factor is | 62 |
| Considerations before using the dump backup | 63 |
| Determining the amount of backup data | 63 |
| Estimating the number of tapes for the backup | 63 |
| Scalability limits for dump backup and restore sessions | 64 |
| When to restart a dump backup | 64 |
| How a dump restore works | 65 |
| What the dump engine restores | 66 |
| Considerations before restoring data | 67 |
| How to perform a dump backup and restore using NDMP services | 68 |
| Environment variables supported for dump | 68 |
| Enabling or disabling enhanced DAR functionality | 79 |
| What the ndmpcopy command does | 80 |
| Displaying file history statistics | 85 |
| How to perform a dump backup using the CLI | 87 |
| What the dump command syntax is | 87 |
| Where to enter the dump command | 89 |
| Specifying the backup level | 90 |
| Improving incremental dump performance | 91 |
| Updating the /etc/dumpdates file | 91 |
| Specifying a local tape device | 92 |
| Specifying a tape device on a remote storage system | 92 |
| Specifying the dump path | 94 |
| Specifying a list of files for backup | 95 |
| Backing up all data that is not in a qtree | 96 |

| | |
|--|-----|
| Excluding specified files and directories | 97 |
| Omitting ACLs from a backup | 98 |
| Specifying a name for a backup | 99 |
| Specifying a blocking factor | 99 |
| Specifying the tape file size | 100 |
| Appending backups to tapes | 101 |
| Verifying the files backed up by a dump command backup | 101 |
| Checking the status of a dump backup | 102 |
| Finding out whether a backup has to be restarted | 104 |
| How to get details about a specific backup | 105 |
| Restarting a dump command backup | 106 |
| Deleting restartable dump command backups | 107 |
| How to perform a dump restore using the CLI | 107 |
| Restore command syntax | 107 |
| What restore types are | 108 |
| What modifiers are | 108 |
| Where to enter the restore command | 109 |
| Executing a restore command | 110 |
| Restoring incremental backups | 110 |
| Restoring each volume backed up as separate subtrees or qtrees | 111 |
| Restoring individual files and directories | 111 |
| Specifying a full restore | 112 |
| What a table-of-contents restore is | 112 |
| Specifying a resume restore | 114 |
| Specifying tape devices in the restore command | 115 |
| Specifying a single tape file on a multfile tape | 115 |
| Specifying the restore destination | 116 |
| Specifying the blocking factor during restore | 117 |
| Displaying detailed status output | 118 |
| Ignoring inode limitations | 118 |
| Specifying automatic confirmations | 119 |
| Specifying no ACLs to be restored | 120 |
| Specifying not to restore qtree information | 121 |
| Specifying a test restore | 121 |
| Restore examples: Restoring using a remote tape drive | 122 |
| Restore examples: Multiple tape restores | 123 |

| | |
|---|------------|
| Data backup to tape using the SMTape engine | 125 |
| How SMTape backup works | 125 |
| What tape seeding is | 126 |
| Features of SMTape | 126 |
| Features not supported in SMTape | 127 |
| How to perform an SMTape backup and restore using NDMP services | 127 |
| Environment variables supported for SMTape | 127 |
| How to back up and restore using the SMTape commands | 130 |
| Backing up data to tape using SMTape | 130 |
| Displaying the volume geometry of a traditional volume | 131 |
| Displaying the image header of a tape | 132 |
| Restoring data from tape using SMTape | 133 |
| Aborting a backup or restore operation using <code>smtape abort</code> command | 134 |
| Continuing a backup or restore after reaching the end of tape | 134 |
| Displaying the status of SMTape backup and restore operations | 135 |
| Removing the SnapMirror status entries | 136 |
| Enabling or disabling concurrent volume SnapMirror and SMTape backup operations | 138 |
| Performing SMTape restores | 139 |
| Performing a baseline restore | 139 |
| Performing an incremental restore | 140 |
| What event logging is | 141 |
| Accessing the event log files | 141 |
| What the dump and restore event log message format is | 141 |
| What logging events are | 142 |
| What dump events are | 142 |
| What restore events are | 144 |
| What the SMTape event log message format is | 145 |
| What SMTape CLI backup and restore events are | 146 |
| What SMTape backup events are | 146 |
| What SMTape restore events are | 147 |
| Enabling or disabling event logging | 148 |
| Error messages for tape backup and restore | 150 |
| Backup and restore error messages | 150 |
| Resource limitation: no available thread | 150 |
| Duplicated tape drive (<code>tape_drive</code>) specified in the tape argument list | 150 |

| | |
|--|-----|
| Invalid tape drive tape_drive in tape argument list | 150 |
| Tape reservation preempted | 151 |
| Could not initialize media | 151 |
| Too many active dumps/restores currently in progress | 151 |
| Media error on tape write | 151 |
| Tape write failed | 152 |
| Tape write failed - new tape encountered media error | 152 |
| Tape write failed - new tape is broken or write protected | 152 |
| Tape write failed - new tape is already at the end of media | 152 |
| Tape write error | 152 |
| Media error on tape read | 152 |
| Tape read error | 153 |
| Already at the end of tape | 153 |
| Tape record size is too small. Try a larger size. | 153 |
| Tape record size should be block_size1 and not block_size2 | 153 |
| Tape record size must be in the range between 4KB and 256KB | 153 |
| NDMP error messages | 154 |
| Network communication error | 154 |
| Message from Read Socket: error_string | 154 |
| Message from Write Dirnet: error_string | 154 |
| Read Socket received EOF | 154 |
| ndmpd invalid version number: version_number | 155 |
| Error: Unable to retrieve session information | 155 |
| ndmpd session session_ID not active | 155 |
| No such user user_name | 155 |
| Cannot generate NDMP password | 155 |
| The specified operation could not be completed as the volume is moving . | 155 |
| Could not obtain vol ref for Volume volume_name | 156 |
| ndmpcopy error messages | 156 |
| Ndmcopy: Socket connection to host_name failed | 156 |
| Ndmcopy: Error opening NDMP connection | 156 |
| Ndmcopy: Client authentication request failed | 156 |
| Ndmcopy: Authentication failed for source | 157 |
| Ndmcopy: Authentication failed for destination | 157 |
| Ndmcopy: Failed to start dump on source | 157 |
| Ndmcopy: Failed to start restore on destination | 157 |

| | |
|--|-----|
| Ndmcopy: Error in getting extension list | 157 |
| Error getting local hostname | 158 |
| Ndmcopy: Connection setup for transfer failed | 158 |
| CONNECT: Connection refused | 158 |
| Invalid name. Source filer name does not resolve to the specified address mode | 158 |
| Invalid name. Destination filer name does not resolve to the specified address mode | 159 |
| Dump error messages | 159 |
| No default tape device list | 159 |
| Invalid/offline volume | 159 |
| Unable to lock a snapshot needed by dump | 159 |
| Failed to determine snapshot type | 159 |
| Volume is temporarily in a transitional state | 160 |
| Unable to locate bitmap files | 160 |
| Failed to locate the specified restartable dump | 160 |
| Dump context created from NDMP. Cannot restart dump | 160 |
| Unable to locate snapshot | 160 |
| Invalid inode specified on restart | 161 |
| Invalid restart context. Cannot restart dump | 161 |
| Failed to retrieve saved info for the restartable dump | 161 |
| Destination volume is read-only | 161 |
| Destination qtree is read-only | 161 |
| IB restore in progress | 161 |
| Could not access volume in path: volume_name | 162 |
| No files were created | 162 |
| SMTape error messages | 162 |
| Internal assertion | 162 |
| Job aborted due to shutdown | 162 |
| Job not found | 162 |
| Job aborted due to Snapshot autodelete | 163 |
| Invalid volume path | 163 |
| UNIX style RMT tape drive is not supported | 163 |
| Volume is currently in use by other operations | 163 |
| Volume offline | 163 |
| Volume not restricted | 163 |

| | |
|--|-----|
| Tape is currently in use by other operations | 164 |
| Invalid input tape | 164 |
| Too many active jobs | 164 |
| Failed to allocate memory | 164 |
| Failed to get data buffer | 164 |
| Failed to create job UUID | 164 |
| Failed to create snapshot | 165 |
| Failed to find snapshot | 165 |
| Failed to lock snapshot | 165 |
| Failed to access the named snapshot | 165 |
| Failed to softlock qtree snapshots | 165 |
| Failed to delete softlock | 166 |
| Failed to delete snapshot | 166 |
| Image header missing or corrupted | 166 |
| Chunks out of order | 166 |
| Tapes out of order | 166 |
| Already read volume_name tape_number | 166 |
| Mismatch in backup set ID | 167 |
| Aborting: Destination volume, volume_name, is too small | 167 |
| Aborting: Destination volume, volume_name, is a clone | 167 |
| Source volume size is greater than maximum supported SIS volume size on this platform. Aborting | 167 |
| Incompatible SnapMirror or copy source Version. Aborting | 168 |
| Transfers from volume volume_name are temporarily disabled | 168 |
| Too many active transfers at once, aborting | 168 |
| Invalid contents in destination volume geometry string volume_geometry_string, aborting | 168 |
| Cannot init input, aborting | 168 |
| Source volume is not a flexible volume. Aborting | 169 |
| Source volume is a flexible volume. Aborting | 169 |
| Destination is not an aggregate. Aborting | 169 |
| Source is not an aggregate. Aborting | 169 |
| Source is not a hybrid aggregate. Aborting | 169 |
| Invalid checksum for the chunk descriptor | 170 |
| Received VBN header with invalid checksum error_string, aborting transfer on volume volume_name | 170 |

| | |
|---|------------|
| Duplicate VBN VBN_number received for volume volume_name, aborting transfer | 170 |
| Bad block in read stream. VBN = VBN_number, max_VBN = max_VBN_number | 171 |
| Invalid checksum found for one of the data block, where VBN number is VBN_number | 171 |
| Block for VBN VBN_number failed checksum verification, aborting the current transfer on volume volume_name | 172 |
| Language setting for the Snapshot is not found | 172 |
| Volume is currently under migration | 172 |
| Failed to get latest snapshot | 172 |
| Failed to load new tape | 172 |
| Remote tape not supported | 173 |
| Failed to initialize tape | 173 |
| Failed to initialize restore stream | 173 |
| Failed to read backup image | 174 |
| Invalid backup image magic number | 174 |
| Chunk format not supported | 174 |
| Invalid backup image checksum | 174 |
| Mismatch in backup level number | 174 |
| Mismatch in backup time stamp | 174 |
| Volume read-only | 175 |
| Invalid source path: /vol/newvol/ | 175 |
| Copyright information | 176 |
| Trademark information | 177 |
| How to send your comments | 178 |
| Index | 179 |

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.

Types of tape backup supported by Data ONTAP

Data ONTAP supports two types of tape backup: the dump backup and the SMTape backup.

Tape backup using dump

Dump is a Snapshot copy-based backup to tape, in which your file system data is backed up to tape. The Data ONTAP dump engine backs up files, directories, and the applicable ACL information to tape. Dump supports level-0, differential, and incremental backups.

Tape backup using SMTape

SMTape is a Snapshot copy-based high performance disaster recovery solution that backs up blocks of data to tape. You can use SMTape to perform volume backups to tapes. However, you cannot

perform a backup at the qtree or subtree level. SMTape supports level-0, differential, and incremental backups.

Related concepts

[Data backup to tape using the dump engine](#) on page 58

[Data backup to tape using the SMTape engine](#) on page 125

[Differences between dump backup and SMTape backup](#) on page 13

How to initiate a dump or SMTape backup

You can initiate a dump or SMTape backup by using the Data ONTAP CLI commands or through NDMP-compliant backup applications.

When you use a backup application to back up your data, you must choose the backup type when initiating a backup.

You can perform a CLI-based dump backup or restore using the Data ONTAP `dump` and `restore` commands.

Similarly, you can perform a CLI-based SMTape backup or restore using the Data ONTAP `smtape backup` and `smtape restore` commands.

Related concepts

[Data backup to tape using the dump engine](#) on page 58

[Data backup to tape using the SMTape engine](#) on page 125

Differences between dump backup and SMTape backup

There are certain differences between dump and SMTape backup engines such as type of data backed up, support of single file restore, and preservation of deduplication. The SMTape backup provides faster backup performance when compared to a dump backup.

The following table lists the differences between SMTape backup and dump backup:

| SMTape backup | Dump backup |
|--|--|
| Backs up blocks of data to tape. | Backs up files and directories to tape. |
| Does not support single file restore. | Supports single file restore. |
| Capable of backing up multiple Snapshot copies in a volume. | Capable of backing up only the base Snapshot copy. |
| Preserves deduplication while backing up and restoring data. | Does not preserve deduplication while backing up data. |

Considerations before choosing a tape backup method

You must consider your business requirements before choosing a tape backup method. Data ONTAP supports dump backup and SMTape backup methods.

You should use dump backup and restore if you want the following features:

- A backup and recovery solution that helps you perform the following tasks:
 - Perform Direct Access Recovery (DAR) of files and directories.
 - Back up some, but not all, subdirectories or files in a specific path.
 - Exclude specific files and directories during a backup.
- Preserve your backups for several years.

You should use SMTape backup and restore if you want the following features:

- A disaster recovery solution that provides high performance.
- To use tape backup to perform an initial full-volume transfer of a source SnapMirror volume to a remote destination storage system and then perform incremental transfers over the network.
In such cases, you can perform an SMTape backup of the SnapMirror volume to a tape, ship the tape to the remote location and restore the contents to a target volume, and set up a SnapMirror relationship. After the SnapMirror relationship is established, the incremental backups are performed over the network. You can also use this method to establish a SnapMirror relationship between source and destination storage systems over a low-bandwidth connection.
- To preserve the deduplication on the backed up data during the restore operation.
- To back up large volumes.

If you use the dump engine to back up volumes with a large number of small files, your backup performance might be affected. This is because, the dump engine performs a file system level backup and has to traverse through the files and directories to back up the volumes. These volumes can be more efficiently backed up to tape by using SMTape.

How online migration affects tape backup

You cannot perform a tape backup or restore of a vFiler volume during the cutover phase of online migration.

Online migration affects tape backup in the following ways:

- Backup or restore of a vFiler volume during the cutover phase of online migration results in the following message: `volume is currently under migration`
- During the cutover phase of online migration, transfer of file system data using the `ndmpcopy` command results in a failure.
- After online migration, incremental backup of a vFiler volume is possible depending on whether the backup is made from the vFiler unit or `vfiler0`.

| Before migration, if a backup is made from... | After migration, if an incremental backup is made from... | Is incremental backup of the vFiler volume possible after migration? |
|---|---|--|
| vFiler unit | vFiler unit | yes |
| vFiler unit | vfiler0 | No |
| vfiler0 | vFiler unit | No |
| vfiler0 | vfiler0 | No |

For more information about online migration, see the *Data ONTAP MultiStore Management Guide for 7-Mode*.

How volume move operations affect tape backup

You cannot perform a tape backup or restore operation while a volume move operation is in cutover phase. Similarly, the cutover phase of a volume move operation cannot start while a tape backup or restore is in progress. You must wait until one of the operations is complete before initiating the other.

For more information about volume move operations, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

Tape drive management

You need to manage tape drives when you back up data from the storage system to tape and 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 by 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 that tape file.

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, SAS, and parallel SCSI tape drives and libraries when they are connected to the interface ports. Data ONTAP detects these drives when their interfaces are enabled.

Types of tape devices

There are two types of tape devices: local and remote. A local tape device is on a storage system that performs the tape operation. A remote tape device is connected through the network to a host or storage system that is performing the tape operation.

The remote tape device has a trust relationship with the storage system that performs the tape operation. The remote magnetic tape (RMT) protocol, which is a bundled component of Data ONTAP runs on the remote tape device.

Note: SMTape does not support remote tape backups and restores.

Note: You cannot use tape devices associated with tape libraries (media changers) on remote hosts.

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, rewind type, 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 (available through the nodeshell). 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
- m** Medium compression
- l** Low 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 `urst` 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 25

Supported number of simultaneous tape devices

Data ONTAP supports a maximum of 64 simultaneous tape drive connections, 16 medium changers, and 16 bridge or router devices for each storage system in any mix of Fibre Channel, SCSI, or SAS attachments.

Tape drives or medium changers can be devices in physical or virtual tape libraries or stand-alone devices.

Note: Although a storage system can detect 64 tape drive connections, the maximum number of backup and restore sessions that can be performed simultaneously depends upon the scalability limits of the backup engine.

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 Compression
Setting      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 Compression
Setting      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 simplifies the process of device identification. Aliasing binds a physical path name (PPN) or a serial number (SN) of a tape or a medium changer to a persistent, but modifiable alias name.

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. |
| 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`, `restore`, `smtape backup`, and `smtape restore` commands.

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

There are two types of names available for aliasing: physical path name and serial number.

Related tasks

[Assigning tape aliases](#) on page 25

[Removing tape aliases](#) on page 25

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

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.

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.

SAS tape devices are also direct-attached devices. For example, the PPN 5c.4 indicates that in a storage system, the SAS HBA is connected in slot 5, SAS tape is connected to port C of the SAS HBA, and the device ID is 4.

PPNs of Fibre Channel 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 device ID 3 and has the LUN 2.

The LUN (logical unit number) 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 device ID and LUN is connected to port 5 of the switch MY_SWITCH, the new tape drive would be accessible by using MY_SWITCH:5.3L2.

What serial numbers are

A serial number (SN) is a unique identifier for a tape drive or a medium changer. Starting with Data ONTAP 8.2, Data ONTAP generates aliases based on SN instead of the WWN.

Since the SN is a unique identifier for a tape drive or a medium changer, the alias remains the same regardless of the multiple connection paths to the tape drive or medium changer. This helps storage systems to track the same tape drive or medium changer in a tape library configuration.

The SN of a tape drive or a medium changer does not change even if you rename the Fibre Channel switch to which the tape drive or medium changer is connected. However, in a tape library if you replace an existing tape drive with a new one, then Data ONTAP generates new aliases because the SN of the tape drive changes. Also, if you move an existing tape drive to a new slot in a tape library or remap the tape drive's LUN, Data ONTAP generates a new alias for that tape drive.

Attention: You must update the backup applications with the newly generated aliases.

The SN of a tape device uses the following format: SN[xxxxxxxxxx]L[X]

x is an alphanumeric character and Lx is the LUN of the tape device. If the LUN is 0, the Lx part of the string is not displayed.

Each SN consists of up to 32 characters; the format for the SN is not case-sensitive.

Displaying existing aliases of tape drives

You can determine the existing aliases of tape drives and medium changers by using the `storage alias` command.

Step

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

```
storage alias
```

Example

The following example displays two tape drive aliases, `st0` and `st2`, and two medium changer aliases, `mc0` and `mc1` mapped to serial numbers:

```
STSW-3070-2_cluster-01>storage alias
Alias                Mapping
-----
st0                  SN[HU1008922R]
st2                  SN[1068060730]
mc0                  SN[c940abe8b0c3a0980248c8]
mc1                  SN[2B13078413]L1
```

Displaying information about tape drives and medium changers

You can display information about tape drives and medium changers that helps you assign tape aliases.

Step

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

```
storage show {tape | mc} [name]
```

name is the name of the tape device or medium changer.

Examples

The following examples show detailed information about tape device, nrst0a, and medium changer, 3d.0L1:

```
filer1>storage show tape nrst0a
Tape Drive:          3c.0
Description:         Hewlett-Packard LTO-5
Serial Number:      HU1008922R
WWNN:               5:001:10a001:389194
WWPN:               5:001:10a001:389194
Alias Name(s):      st0
Device State:       available
```

```
filer1>storage show mc 3d.0L1
Medium Changer:     3d.0L1
Description:        OVERLAND NEO Series
Serial Number:      2B13078413
WWNN:               5:005:076312:4b4d6c
WWPN:               5:005:076312:4b4d6c
Alias Name(s):      mc1
Device State:       available
```


Assigning tape aliases

You can assign tape aliases to provide a correspondence between the logical names of backup devices and a name permanently assigned to a port, a tape drive, or a medium changer.

Step

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

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

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.

SN is the unique identifier of a 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 SN[HU106150D4]
```

The medium changer alias `mc80` is mapped to its serial number `SN[HU106150D4]` on LUN 0.

Removing tape aliases

You can remove aliases from tape drives, 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.

Considerations when configuring multipath tape access

You can configure multiple paths from the storage system to access tape drives in a tape library. If one path fails, then the storage system can use the other paths to access tape drives without having to immediately repair the failed path. This ensures that tape operations can be restarted.

You must take into account a list of considerations when configuring multipath tape access from your storage system:

- In tape libraries that support LUN mapping, for multipath access to a LUN group, LUN mapping must be symmetrical on each path.

Tape drives and media changers are assigned to LUN groups (set of LUNs that share the same initiator path set) in a tape library. All tape drives of a LUN group must be available for backup and restore operations on all multiple paths.

- Maximum of two paths can be configured from the storage system to access tape drives in a tape library.
- Multipath tape access does not support load balancing.

How to add tape drives and libraries to storage systems

You can add 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.

In a tape library configuration, you must configure a tape drive or medium changer on LUN 0 of a target port for Data ONTAP to discover all medium changers and tape drives on that target port.

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.

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
rst01 - 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 autoload 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 move and position the tape drives by using the `mt` command.

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> | Indicates 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. |

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

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

Attention: When you use an 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.

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 17

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 17

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. |

What qualified tape drives are

A qualified tape drive is a tape drive that has been tested and found to work properly on storage systems.

You can add support for tape drives to existing Data ONTAP releases by using the tape configuration file. To download the tape configuration file, go to the NetApp Support Site. You can view the instructions required to download the tape configuration file, add support to Data ONTAP for a tape drive that was qualified after the release of the Data ONTAP version, and view the current list of supported tape drives at the NetApp Support Site.

Only qualified tape drives are listed in the tape qualification list. 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 view information about qualified and nonqualified tape drives, tape libraries, and tape drive connections to the storage system.

Related information

[Data Protection - Supported Devices: netapp.com/us/solutions/a-z/data-protection-devices.html](http://netapp.com/us/solutions/a-z/data-protection-devices.html)

[Tape Device Configuration file: support.netapp.com/NOW/download/tools/tape_config/](http://support.netapp.com/NOW/download/tools/tape_config/)

Format of the tape configuration file

The tape configuration file format consists of fields such as vendor ID, product ID, and details of compression types for a tape drive. This file also consists of optional fields for enabling the autoloader feature of a tape drive and changing the command timeout values of a tape drive.

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

| Item | Size | Description |
|-------------------------------|----------------|--|
| <i>vendor_id</i> (string) | up to 8 bytes | The vendor ID as reported by the SCSI Inquiry command. |
| <i>product_id</i> (string) | up to 16 bytes | The product ID as reported by the SCSI Inquiry command. |
| <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 the <code>sysconfig -v</code> or <code>sysconfig -t</code> command; otherwise, <code>INQ_VENDOR_ID</code> is displayed. |

| Item | Size | Description |
|---|-------------------|---|
| <code>product_pretty</code> (string) | up to 16 bytes | If this parameter is present, it is specified by the string displayed by the <code>sysconfig -v</code> or <code>sysconfig -t</code> command; otherwise, <code>INQ_PRODUCT_ID</code> 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 |
|--|-------------------|---|
| <code>{l m h a}_description=(string)</code> | up to 16 bytes | The string to print for the <code>sysconfig -t</code> command that describes characteristics of the particular density setting. |
| <code>{l m h a}_density=(hex codes)</code> | | 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. |
| <code>{l m h a}_algorithm=(hex codes)</code> | | 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=(Boolean yes/no)</code> | This field is set to <code>yes</code> if the tape drive has an automatic loading feature; that is, after tape cartridge is inserted, 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> . |
| <code>cmd_timeout_0x</code> | 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 drive. The timeout value can be expressed in minutes (m), seconds (s), or milliseconds (ms). Note: You should change this field only with guidance from technical support. |

To download and view the tape configuration file, go to the NetApp Support Site.

Example of a tape configuration file format

The tape configuration file format for the HP LTO5 ULTRIUM tape drive is as follows:

```

vendor_id="HP"
product_id="Ultrium 5-SCSI"
id_match_size=9
vendor_pretty="Hewlett-Packard"
product_pretty="LTO-5"
l_description="LTO-3(ro)/4 4/800GB"
l_density=0x00
l_algorithm=0x00
m_description="LTO-3(ro)/4 8/1600GB cmp"
m_density=0x00
m_algorithm=0x01
h_description="LTO-5 1600GB"
h_density=0x58
h_algorithm=0x00
a_description="LTO-5 3200GB cmp"
a_density=0x58
a_algorithm=0x01
autoload="yes"

```

Related information

Tape Device Configuration file: support.netapp.com/NOW/download/tools/tape_config

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.

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.

Displaying information about nonqualified tape drives

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

About this task

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.

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 nrst1a 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 to 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. 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

[What qualified tape drives are](#) on page 34

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 tape drives, medium changers, bridges, and tape libraries.

Note: All the systems that share devices in a 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 or target reset; however, not all devices implement SCSI Persistent Reservations correctly.

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 types of data transfer 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 communicate with NDMP-enabled network-attached backup applications (also called *Data Management Applications* or *DMAs*), data servers, and tape servers participating in backup or recovery operations. All network communications occur over TCP/IP or TCP/IPv6 network. NDMP also provides low-level control of tape drives and medium changers.

Advantages of NDMP

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 by using the `ndmpd password userid` command.

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 perform NDMP operations through an NDMP-compliant backup application. For the NDMP password to be generated, the NDMP user must have the `login-ndmp` capability.

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 value 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

```
Thu Apr 15 09:27:00 GMT Apr 15 09:27:00
[host1:ndmp.connection.accept:info]: ndmpd.access allowed for version
= 4, sessionId = 2922, from src ip = 192.0.2.68, dst ip = 192.0.2.100,
src port = 41855, 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>ndmpd password barbaraD
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.

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.

Designating the range of ports for NDMP data connections

Data ONTAP supports a designated range of TCP/IP ports that can be used for NDMP data connections in response to NDMP_DATA_LISTEN and NDMP_MOVER_LISTEN operations.

About this task

Data ONTAP 8.0 7-Mode and earlier versions do not support data migration by using the `ndmcopy` command and three-way tape backups in environments where the source and destination networks are separated by a firewall. This is because the data or mover port that is used in a data transfer is unpredictable.

Starting with Data ONTAP 8.0.1, administrators can designate range of ports that can be used for NDMP data connections in response to NDMP_DATA_LISTEN and NDMP_MOVER_LISTEN operations. Therefore, Data ONTAP enables you to perform data migration by using `ndmcopy` command and 3-way tape backups even in environments where the source and destination networks are separated by a firewall.

Step

1. To enable the data port range, enter the following command:

```
options ndmpd.data_port_range {start_port-end_port}
```

The `ndmpd.data_port_range` option allows administrators to specify a port range on which the NDMP server can listen for data connections.

The `start_port` and `end_port` indicate the range of ports designated for data connection and can have values between 1024 and 65535; `start_port` must be less than or equal to `end_port`.

If a valid range is specified, NDMP uses a port within that range to listen for incoming data connections. A listen request fails if no ports in the specified range are free.

The default value for `ndmpd.data_port_range` option is `all`. The `all` implies that any available port can be used to listen for data connections.

Note: The `ndmpd.data_port_range` option is persistent across reboots.

Example

```
Filer1> options ndmpd.data_port_range 1024-2048
```

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
```

Example

```
filerA> ndmpd status
ndmpd ON.
Session: 12923
Active
version:                4
Operating on behalf of primary host.
tape device:           nrst0a
mover state:           Active
data state:             Connected
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

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:             192.0.2.65
  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

filerA> ndmpd probe
ndmpd ON.
Session: 12923
  isActive:                TRUE
  protocol version:        4
  effHost:                 Local
  authorized:              TRUE
  client addr:             192.0.2.100:1591
```

```

spt.device_id:      none
spt.ha:             -1
spt.scsi_id:        -1
spt.scsi_lun:       -1
tape.device:        nrst0a
tape.mode:          Read/Write
mover.state:        Active
mover.mode:         Read
mover.pauseReason: N/A
mover.haltReason:   N/A
mover.recordSize:   64512
mover.recordNum:    0
mover.bytesMoved:   0
mover.seekPosition: 0
mover.bytesLeftToRead: 0
mover.windowOffset: 0
mover.windowLength: 18446744073709551615
mover.position:     0
mover.setRecordSizeFlag: true
mover.setWindowFlag: true
mover.connect.addr_type: LOCAL
data.operation:     None
data.state:         Connected
data.haltReason:    N/A
data.connect.addr_type: LOCAL
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. 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 ndmpd.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.

Displaying the NDMP version

Starting with Data ONTAP 8.2, the storage system supports only NDMP version 4. You can view the latest NDMP version that the storage system is currently set to use.

Step

1. To view the NDMP version, enter the following command:

```
ndmpd version
```

The latest version that NDMP currently allows you to use is displayed.

NDMP options

You can use NDMP options to manage NDMP on your storage system.

The following table lists the NDMP options that you can use with the `options` command:

| Option | Description | Default value |
|---|--|---------------|
| <code>ndmpd.access {all legacy host[!]=hosts if[!]=interfaces}</code> | Specifies the hosts or interfaces through which NDMP sessions are permitted or blocked | all |
| <code>ndmpd.authtype {challenge plaintext plaintext,challenge}</code> | Specifies whether a storage system accepts plaintext, challenge, or both to authenticate NDMP session requests | challenge |
| <code>ndmpd.connectlog.enabled {on off}</code> | Enables or disables the monitoring of NDMP connections | off |

| Option | Description | Default value |
|---|--|---------------|
| ndmpd.enable {on off} | Enables or disables NDMP service on your storage system | off |
| ndmpd.maxversion | Specifies the highest NDMP version Note: Starting with Data ONTAP 8.2, only NDMP version 4 is supported. | 4 |
| ndmpd.ignore_ctime.enabled {on off} | Enables or disables incremental backup of files that have their ctime changed since the previous backup | off |
| ndmpd.offset_map.enable {on off} | Enables or disables offset map generation during backup | on |
| ndmpd.password_length {8 16} | Specifies the length of NDMP password | 16 |
| ndmpd.preferred_interface {interface disable} | Specifies the preferred network interface to be used for NDMP data connections | disable |
| ndmpd.tcpnodelay.enable {on off} | Optimizes the performance of the NDMP socket through which the storage system communicates with the DMA | off |
| ndmpd.tcpwinsize {tcp_window_size} | Specifies the TCP window size for data connection Note: The valid TCP window size range is 8192 - 7631441. | 32768 |

| Option | Description | Default value |
|---|---|---------------|
| <code>ndmpd.data_port_range</code> { <i>start_port</i> - <i>end_port</i> all} | Specifies a port range on which the NDMP server can listen for data connections Note: <ul style="list-style-type: none"> • <i>start_port</i> and <i>end_port</i> can have values between 1024 and 65535; <i>start_port</i> must be less than or equal to <i>end_port</i>. • It is best to use <i>start_port</i> and <i>end_port</i> values between 18600 and 18699. | all |

NDMP extensions supported by Data ONTAP

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

The following are the NDMP v4 extensions supported by Data ONTAP:

- Restartable backup
This extension is not supported by SMTape.
- SnapVault management
- SnapMirror management
- Snapshot management
- CAE (Connection Address Extension) for IPv6 support

Tape backup using NDMP services

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

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.

- A storage system supports a maximum of 32 SMTape sessions over a TCP/IP network.
- 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.

Note: SMTape does not support file history.

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.

Note: When your data protection needs are limited to disaster recovery, it is recommended that you use SMTape to back up data.

- When a SnapMirror destination is backed up to tape, only the data on the volume is backed up.

The SnapMirror relationships and the associated metadata are not backed up to tape. Therefore, during restore, only the data on that volume is restored and the associated SnapMirror relationships are not restored.

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.

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.

Example

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 perform 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 view the NDMP version supported by the storage system, enter the following command:

```
ndmpd version
```

Note: Your backup application must support the same NDMP version that is supported by your storage system.

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.

Example

The following example shows the authentication type specified is plaintext and challenge:

```
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 performing 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 perform NDMP backup and restore operations.

Related tasks

[Enabling and disabling NDMP services](#) on page 45

[Specifying a preferred network interface](#) on page 46

[Specifying NDMP access by host or interface](#) on page 42

[Specifying the NDMP authentication type](#) on page 43

[Generating an NDMP-specific password for non-root administrators](#) on page 45

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 68

Related references

[Environment variables supported for SMTape](#) on page 127

Data backup to tape 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 back up 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 Data ONTAP operating in 7-Mode 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, the dump engine creates a Snapshot copy for the backup. After the backup operation is completed, the dump engine deletes this Snapshot copy.

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

How a dump backup works

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

You can back up a volume, a 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. After the backup operation is completed, the dump engine deletes this Snapshot copy.

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.

The dump engine does not back up inconsistent LUN clones. Inconsistent LUN clones are LUN clones whose backing Snapshot copies are missing and therefore have missing data blocks.

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) with 1024 Access Control Entries (ACEs)

Note: If you restore ACLs backed up from storage systems running Data ONTAP 8.2 to storage systems running Data ONTAP 8.1.x and earlier that have an ACE limit lower than 1024, a default ACL is restored.

- 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 vault destination volume or a volume SnapMirror destination to tape, only the data on the volume is backed up. The associated metadata is not backed up. Therefore, when you try

to restore the volume, only the data on that volume is restored. Information about the volume 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. You can perform nine levels of incremental backup operations.

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 at each subsequent backup.
- 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 |

| Back-up order | Increment level | Increment chain | Base | Files backed up |
|---------------|-----------------|-----------------|---|--|
| 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 |

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 8.1 supports a blocking factor between the range of 4 KB to 256 KB. The default blocking factor is 63 KB. 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`.

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.

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 63

Scalability limits for dump backup and restore sessions

You must be aware of the maximum number of dump backup and restore sessions that can be performed simultaneously on storage systems of different system memory capacities. This maximum number depends on the system memory of a storage system.

| System memory of a storage system | Total number of dump backup and restore sessions |
|--|--|
| Less than 16 GB | 4 |
| Greater than or equal to 16 GB but less than 24 GB | 16 |
| Greater than or equal to 24 GB | 32 |

Note: If you use `ndmptcopy` command to copy data within storage systems, two sessions are established: dump backup and dump restore.

Note: The number of dump backup and restore sessions indicate backup and restore operations initiated by NDMP as well as the `dump` and `restore` commands.

You can obtain the system memory of your storage system by using the `sysconfig -a` command (available through the nodeshell). For more information about using this command, see the man pages.

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.

You can restart an aborted backup to tape only if the following conditions are true:

- The aborted backup is in phase IV.
- All the associated Snapshot copies that were locked by the `dump` command are available.

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/voll                   Level:     0
Snapshot: filer(0101184236)_voll_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 106

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.

Note: If you restore ACLs backed up from storage systems running Data ONTAP 8.2 to storage systems running Data ONTAP 8.1.x and earlier that have an ACE limit lower than 1024, a default ACL is restored.

- 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 `/vs1/vol1/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.
- Clustered Data ONTAP volume can be restored to a 7-Mode volume.
- SnapLock volumes are restored as normal read/write volumes.

- The destination volume for a restore operation might have files with mandatory or advisory locks. While performing restore operation to such a destination volume, the dump engine ignores these locks.

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:

- You must perform a level 0 restore and incremental restores before you upgrade to Data ONTAP 6.2 or later.
- You must 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

You can restore the backed up data to its original path or to a different destination. If you are restoring the backed up data to a different destination, you must prepare the destination for the restore.

Before restoring data either to its original path or to a different destination, you must have the following information and meet the following requirements:

- The level of the restore
- The path into which you are restoring the data
- The blocking factor used during the backup
- If you are doing an incremental restore, all tapes must be in the backup chain.
- A tape drive that is available and compatible with the tape to be restored from.

Before restoring data to a different destination, you must perform the following operations:

- 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.

To rename a file, directory, or qtree during restore without using DAR, you must set the `EXTRACT` environment variable to `E`.

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: Restore checks for volume space and inode availability on the destination volume when the restore operation starts. Setting the `FORCE` environment variable to `Y` causes the restore operation to skip checks for volume space and inode availability on the destination path. If there is not enough volume space or inodes available on the destination volume, the restore operation recovers as much data allowed by the destination volume space and inode availability. The restore operation stops when there is no more volume space or inodes left.

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.

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 <code>dump</code> command. When set to -1, the BASE_DATE incremental specifier is disabled. When set to 0 on a level 0 backup, incremental backups are enabled. Subsequent to the initial backup, the value of the DUMP_DATE variable from the previous incremental backup is assigned to the BASE_DATE variable. These variables are an alternative to the <code>/etc/dumpdates</code> file for controlling incremental backups. These variables are an alternative to the LEVEL/UPDATE based 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 n option of the dump 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 prefixing 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 in 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 |

| Environment variable | Valid values | Default | Description |
|----------------------|-----------------------|---------|--|
| 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. |
| EXTRACT | Y, N, or E | N | <p>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. To rename a file, directory, or qtree during restore without using DAR, you must set the EXTRACT environment variable to E.</p> |

| Environment variable | Valid values | Default | Description |
|----------------------|---------------|---------|--|
| 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 A option of the restore 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, /vol/vol0/etc. |
| FORCE | Y or N | N | <p>The FORCE variable is equivalent to the F option of the restore command.</p> <p>Determines if the restore operation must check for volume space and inode availability on the destination volume.</p> <p>Setting this variable to Y causes the restore operation to skip checks for volume space and inode availability on the destination path.</p> <p>If there is not enough volume space or inodes available on the destination volume, the restore operation recovers as much data allowed by the destination volume space and inode availability. The restore operation stops when there is no more volume space or inodes left.</p> |

| 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> |

| Environment variable | Valid values | Default | Description |
|----------------------|---------------|---------|--|
| LEVEL | 0-9 | 0 | 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. |
| LIST | Y or N | N | Lists the backed-up file names and inode numbers without actually restoring the data. The LIST variable is equivalent to the <code>t</code> option of the <code>restore</code> command. |
| LIST_QTREES | Y or N | N | Lists the backed-up qtrees without actually restoring the data. 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. |

| Environment variable | Valid values | Default | Description |
|----------------------|--------------------|---------|---|
| NDMP_UNICODE_FH | Y or N | N | Specifies that a Unicode name is included in addition to the NFS name of the file in the file history information. This option is not used by most backup applications and should not be set unless the backup application is designed to receive these additional file names. The HIST variable must also be set. |
| NDMP_VERSION | <i>return_only</i> | none | You should not modify the NDMP_VERSION variable. Created by the backup operation, the NDMP_VERSION variable returns the NDMP version. Data ONTAP sets the NDMP_VERSION variable during a backup for internal use and to pass to a backup application for informational purposes. The NDMP version of an NDMP session is not set with this variable. |
| NO_ACLS | Y or N | N | Specifies that ACLs not be copied when backing up data. The NO_ACLS variable is equivalent to the A option of the dump command. Ordinarily a backup using the dump command writes out metadata related to Windows ACLs. The NO_ACLS variable stops this information from being backed up. |

| Environment variable | Valid values | Default | Description |
|----------------------|---------------------|---------|--|
| 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 dump 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 restore command. This variable is used for debugging. |
| PATHNAME_SEPARATOR | <i>return_value</i> | none | Specifies the pathname separator character. This character depends upon the file system being backed up. For Data ONTAP, the character "/" is assigned to this variable. NDMP server sets this variable prior to starting a tape backup operation. |

| 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. DIRECT and ENHANCED_DAR_ENABLED must be enabled (set to Y) as well. If RECOVER_FULL_PATHS 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. |
| TYPE | dump or smtape | dump | Specifies the type of backup you can choose to perform tape backup and restore operations. Data ONTAP supports two types of backup: the dump backup and SMTape backup. |
| UPDATE | Y or N | Y | Updates the metadata information to enable LEVEL based incremental backup. |
| VERBOSE | Y or N | N | Increases the log messages while performing a tape backup or restore operation. |

Enabling or disabling enhanced DAR functionality

Enhanced direct access recovery (DAR) functionality provides support for directory DAR and DAR of files with NT Streams. You can enable or disable enhanced DAR functionality by using the `options ndmpd.offset_map.enable` command. By default, enhanced DAR functionality is enabled.

About this task

Enhanced direct access recovery (DAR) functionality is supported only for the NDMP-initiated dump backup and restore and provides improved restore performance. This feature is not supported in SMTape backup and restore.

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

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

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 54

What the ndmpcopy command does

The `ndmpcopy` command enables you to transfer file system data between storage systems that support NDMP v3 or v4.

The `ndmpcopy` 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 IPv4 addresses of source and destination storage systems in the `ndmpcopy` command.

Starting with Data ONTAP 7.3.3, the `ndmpcopy` 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 `ndmpcopy` command to use an IPv6 address mode to establish the data connection.

Using the `ndmpcopy` 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 `ndmpcopy` command.

Copying data using ndmpcopy

You can run the `ndmpcopy` 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 run `ndmpcopy` on a single storage system that is both the source and the destination of the data transfer.

Step

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

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


- You can specify one or more options in the `ndmpcopy` command. The following table lists the available options:

| Option | Description |
|--------------------------------------|--|
| <code>-sa username:[password]</code> | 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. |
| <code>-da username:[password]</code> | Destination authorization that specifies the user name and password for connecting to the destination storage system |
| <code>-st {md5 text}</code> | Sets the source authentication type to be used when connecting to the source storage system Note: <code>md5</code> is the default authentication type used. The <code>md5</code> authentication exchanges the user name and password in encrypted form. The <code>text</code> authentication exchanges the user name and password in clear text. |
| <code>-dt {md5 text}</code> | Sets the destination authentication type to be used when connecting to the destination storage system |
| <code>-l</code> | 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. |
| <code>-d</code> | Enables generation of <code>ndmpcopy</code> debug log messages <code>ndmpcopy</code> debug log files appear in the root volume <code>/etc/log</code> directory. The <code>ndmpcopy</code> debug log file names are in the form <code>ndmpcopy.yyyymmdd</code> . |

| Option | Description |
|----------|---|
| -f | Enables forced mode This mode enables overwriting system files in the <code>/etc</code> directory on the root volume. |
| -h | Prints the help message |
| -p | Prompts you to enter the password for source and destination authorization This password overrides the password specified for <code>-sa</code> and <code>-da</code> options. Note: You can use this option only when the command is running in an interactive console. |
| -exclude | Excludes specified files or directories from the path specified for data transfer The value can be a comma-separated list of directory or file names such as <code>"*.pst"</code> or <code>"*.txt"</code> . |

- *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 systems.
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 address mode for the data connection 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 address mode for the data connection 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 45

Related references

[Examples of the `ndmcopy` command](#) on page 83

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>ndmcopy -da username:password -f /vol/rootvol
remotehost1:/vol/rootvol
```

Example of the ndmcopy 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>ndmcopy -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. SMTape does not support file history and therefore SMTape initiated backups do not have any file history statistics associated with them.

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.

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 256, 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 <code>dump</code> 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 <code>qtree</code> . |
| u | Updates the <code>/etc/dumpdates</code> 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 60

[How to specify tape devices for the backup](#) on page 61

Related tasks

[Specifying the backup level](#) on page 90

- Omitting ACLs from a backup* on page 98
- Specifying a blocking factor* on page 99
- Specifying the tape file size* on page 100
- Specifying a list of files for backup* on page 95
- Specifying a name for a backup* on page 99
- Backing up all data that is not in a qtree* on page 96
- Updating the /etc/dumpdates file* on page 91
- Excluding specified files and directories* on page 97

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 pressing 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 91

[Backing up all data that is not in a `qtree`](#) on page 96

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 61

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. |

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

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/vol10 | 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. |
| /vol/vol0 | 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>. • <code>*.o</code> specifies that the file name ends in <code>.o</code> (for example, <code>program.o</code>). |

| | |
|------------------------|--|
| | <ul style="list-style-type: none"> • <code>core*</code> specifies that the file name begins with the <code>core</code> string (for example, <code>core.small</code>). • <code>*backup*</code> specifies that the file name contains the backup string (for example, <code>spec.backup.1</code>). |
| <code>/vol/vol1</code> | 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/vol10` volume using a tape file size of 2,097,151, so that a tape file is no larger than 2 GB:

```
dump 0fB rst0a 2097151 /vol/vol0
```

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 29

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 acfs: 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, acfs—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-wafl-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 | urst1a | 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/voll             Level:        0
Snapshot:      filer(0101184236)_voll_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[f 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 102

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.

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 114

[Specifying table-of-contents restores](#) on page 113

[Specifying a full restore](#) on page 112

[Restoring individual files and directories](#) on page 111

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 120

Specifying the restore destination on page 116

Ignoring inode limitations on page 118

Specifying a test restore on page 121

Specifying not to restore qtree information on page 121

Specifying the blocking factor during restore on page 117

Specifying tape devices in the restore command on page 115

Specifying a single tape file on a multifile tape on page 115

Displaying detailed status output on page 118

Specifying automatic confirmations on page 119

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.

Note: The `restore_symboltable` file contains information required to perform incremental restores. If you intend to perform incremental restores, ensure that this file is not deleted until all the incremental restores are completed.

Related concepts

What increment chains are on page 60

Related tasks

Specifying the backup level on page 90

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 for a tape restore. This is useful in determining what files or qtrees exist 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 restore

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 performed 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 multiframe 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 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 for 7-Mode*.

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 rffF rst0a
```

`r` performs a full restore.

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

`F` specifies that you can ignore inode limitations.

`rst0a` is 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 unclean 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 <code>rst0a</code> 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 <code>rst0a</code> 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/vol1
```

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.

Data backup to tape using the SMTape engine

SMTape is a high performance disaster recovery solution from Data ONTAP that backs up blocks of data to tape. It is a Snapshot copy-based backup to tape feature. In releases earlier than Data ONTAP 8.0 operating in 7-Mode, SMTape is referred to as SM2T.

You can use SMTape to perform volume backups to tapes. However, you cannot perform a backup at the qtree or subtree level. SMTape supports level-0, differential, and incremental backups. Using SMTape, you can back up 255 Snapshot copies. For subsequent baseline, incremental, or differential backups, you must delete older backed up Snapshot copies.

When you perform an SMTape level-0 backup, you can specify the name of the Snapshot copy to be backed up to tape. When you specify a Snapshot copy for the backup, all the Snapshot copies older than the specified Snapshot copy are also backed up to tape.

If you do not specify a Snapshot copy for the backup, the following happens:

- If the volume is read-writeable, a Snapshot copy is created automatically. The Snapshot copy that is created and all older Snapshot copies are backed up to tape.
- If the volume is read-only, all the Snapshot copies till the latest Snapshot copy are backed up to tape. Any new Snapshot copies created after the backup has started is not backed up.

When you perform an SMTape incremental or differential backup, the NDMP-compliant backup applications create and manage Snapshot copies.

You can perform an SMTape backup and restore by using NDMP-compliant backup applications or by using the `smtape backup` and `smtape restore` Data ONTAP CLI commands operating in 7-Mode. However, if you want to perform an incremental backup, you must perform both baseline and incremental backups by using only the NDMP-compliant backup applications.

These commands replace the `snapmirror store` and `snapmirror retrieve` commands in releases earlier than Data ONTAP 8.0 operating in 7-Mode.

How SMTape backup works

SMTape backup writes blocks of data to tape in a predefined process.

The following table describes the process that SMTape uses to back up data to tape.

| Stage | Action |
|-------|---|
| 1 | Data ONTAP creates a base Snapshot copy for the backup. If a Snapshot copy name is provided, Data ONTAP uses this Snapshot copy as the base Snapshot copy. |

| Stage | Action |
|-------|--|
| 2 | Data ONTAP begins transferring blocks of data to tape. |

What tape seeding is

Tape seeding is an SMTape functionality that helps you initialize the destination storage system in a volume SnapMirror relationship.

Tape seeding enables you to establish a SnapMirror relationship between a storage system and a destination system over a low-bandwidth connection. Incremental mirroring of Snapshot copies from the source to the destination is feasible over a low bandwidth connection. However, an initial mirroring of the base Snapshot copy would take a long time over a low-bandwidth connection. In such a case, you can perform an SMTape backup of the source volume to a tape and use the tape to transfer the initial base Snapshot copy to the destination. You can then set up incremental SnapMirror updates to the destination system using the low-bandwidth connection.

Features of SMTape

SMTape features such as tape seeding, backup of Snapshot copies, incremental and differential backups, and preservation of deduplication and compression features on restored volumes help you optimize your tape backup and restore operations.

The following are the features of SMTape:

- Provides a high performance disaster recovery solution.
- Does not require a license.
- Storage systems support only 32 concurrent backup and restore sessions.
Even if another node is taken over, the storage system allows only 32 sessions instead of 64 sessions.
- Supported only on NDMP v4.
- Supports restore of backup images that were created across up to two major Data ONTAP releases.

For example, on a storage system running Data ONTAP 8.1.x, you can restore data backed up from Data ONTAP 7.3.x and Data ONTAP 8.0.x.

- Supports restore of an incremental backup in Data ONTAP 8.1.1.
- Supports restore of data backed up in Data ONTAP 8.0.x from 32-bit or 64-bit aggregates to volumes in 64-bit aggregates in Data ONTAP 8.1.x and later.

Note: You can restore data to volumes created across up to two major consecutive Data ONTAP releases only.

For example, if you back up data in Data ONTAP 8.0.x from either a 32-bit or a 64-bit aggregate, then you can restore this data to volumes in 64-bit aggregates in Data ONTAP 8.1.x and the next major Data ONTAP release.

- Supports tape seeding.
- Supports concurrent volume SnapMirror and SMTape backup operations when backing up SnapMirror destination to tape.
- Supports backup of Snapshot copies.
- Supports deduplicated volumes and preserves deduplication on the restored volumes.
- Supports blocking factor in multiples of 4 KB, in the range of 4 KB to 256 KB.
- Supports backup of large aggregate volumes.
- Supports backup of compressed volumes and preserves compression on the restored volumes.
- Supports incremental and differential backups.

Features not supported in SMTape

There are certain features that are not supported in SMTape, such as SnapLock and FlexCache volume, remote tape backup using CLI commands, and restartable backup.

The following features are not supported in SMTape:

- Remote tape backup using the CLI
- SnapLock volume and FlexCache volume
- Restartable backup
- Multiple backups on a single tape
- Backup or restore of selected files or directories
- Verification of files backed up

How to perform an SMTape backup and restore using NDMP services

You can perform an SMTape-based backup and restore by using NDMP-compliant backup applications.

Data ONTAP provides a set of environment variables that enable you to perform a block-level tape backup and restore using NDMP services. However, SMTape does not support DAR and file system data transfer between storage systems.

Environment variables supported for SMTape

Data ONTAP supports a set of environment variables for SMTape. These variables are used to communicate information about a SMTape backup or restore operation between an NDMP-enabled backup application and a storage system.

The following table lists the environment variables supported by Data ONTAP for SMTape backup and recovery, their valid values, default value, and description:

| Environment variable | Valid values | Default value | Description |
|----------------------|--------------|---------------|---|
| DUMP_DATE | return_value | none | At the end of an SMTape backup, DUMP_DATE contains a string identifier that identifies the reference Snapshot copy for the next incremental backup. The resulting value of DUMP_DATE is used as the BASE_DATE value for subsequent incremental backups. |
| BASE_DATE | DUMP_DATE | none | BASE_DATE specifies the start date for incremental backups. BASE_DATE is a string representation of the reference Snapshot identifiers. Using the BASE_DATE string, SMTape locates the reference Snapshot copy. Subsequent to the initial backup, the value of the DUMP_DATE variable from the previous incremental backup is assigned to the BASE_DATE variable. |
| FILESYSTEM | string | none | Specifies the path name of the root of the data that is being backed up. For example, /vol/vol0/etc. |
| SMTAPE_BACKUP_SET_ID | string | none | Specifies the backup set ID for the baseline backup and subsequent incremental backups. Backup set ID is a 128-bit unique ID that identifies the sequence of incremental backups with respect to baseline backup. |

| Environment variable | Valid values | Default value | Description |
|------------------------|---|---------------|---|
| SMTAPE_SNAPSHOT_NAME | Any valid Snapshot copy that is available in the volume | Invalid | When the SMTAPE_SNAPSHOT_NAME variable is set to a Snapshot copy, that Snapshot copy and its older Snapshot copies are backed up to tape. This variable is available only in the SMTape backup context. |
| SMTAPE_DELETE_SNAPSHOT | Y or N | N | <p>When the SMTAPE_DELETE_SNAPSHOT variable is set to Y, SMTape deletes the auto-Snapshot copy after the backup operation is complete.</p> <p>If you specify a Snapshot copy name for the backup, this Snapshot copy is not deleted, but the softlock on this Snapshot copy is removed.</p> <p>Note: When the variable is set to N, SMTape does not delete the auto-Snapshot copy and leaves the softlock on this Snapshot copy.</p> |

| Environment variable | Valid values | Default value | Description |
|----------------------|--------------|---------------|---|
| SMTAPE_BREAK_MIRROR | Y or N | N | <p>When the SMTAPE_BREAK_MIRROR variable is set to Y, it ensures that the SnapMirror relationship established by the restore operation is broken after the operation is complete. This variable is available only in the SMTape restore context.</p> <p>Note: After a successful restore, the restored volume is in the restricted state and does not become writable unless the SMTAPE_BREAK_MIRROR variable is set to Y.</p> |

How to back up and restore using the SMTape commands

You can perform an SMTape backup and restore by using the Data ONTAP CLI commands. You can also manage your SMTape-initiated backup and restore by using the CLI commands.

You can back up and restore data by using the `smtape backup` and `smtape restore` commands. You can also display the volume geometry of a traditional volume and the image header of a tape, abort or continue a backup or restore operation, and display the status of a backup or restore operation by using the SMTape CLI commands.

Backing up data to tape using SMTape

You can perform an SMTape backup by using the `smtape backup` command. You can specify a Snapshot copy name for the backup, in which case the specified Snapshot copy is used as the base Snapshot copy for the backup. When you do not specify a Snapshot copy for the backup, a base Snapshot copy is created and backed up.

Step

1. Enter the following command:

```
smtape backup [-g volume_geometry] [-b block_size] [-s snapshot_name]
path tape_device
```

`g` specifies that the geometry of the backup image is supplied in the command line. This option optimizes the tape for a particular traditional volume destination and increases the restore performance. This option is applicable only to traditional volumes.

Note: The geometry of a FlexVol volume is always 1.

volume geometry is the volume geometry of the traditional volume. You can determine the geometry by using the `smtape restore -g` command on that traditional volume.

b specifies that a blocking factor for the backup is supplied in the command line. It can be in multiples of 4 KB, in the range of 4 KB to 256 KB. The default tape record size is 240 KB.

block_size is the blocking factor for the backup.

s specifies that the base Snapshot copy is supplied in the command line.

snapshot_name is the base Snapshot copy that must be used for the backup.

path is the path of the data to be backed up.

tape_device is the tape device to be used for the backup.

If the storage system crashes during the backup session, the auto-Snapshot copy of the volume being backed up continues to remain in the storage system. This stale Snapshot copy is deleted when you perform an SMTape backup of the volume again.

A unique job ID in the range of 1 to 99999 is assigned to this backup operation. You can subsequently use this job ID to check the backup status or to abort the backup operation. Also, an entry is made in the `/etc/log/backup` file.

In the following example, the data in `/vol/testdata` is backed up to the `rst0a` tape device in blocks of 256 KB:

```
filer>smtape backup -b 256 /vol/testdata rst0a
Job 9 started
```

Related tasks

[Displaying the volume geometry of a traditional volume](#) on page 131

Displaying the volume geometry of a traditional volume

You can view the volume geometry of a specific traditional volume by using the `smtape restore` command.

Step

1. To display the volume geometry of a traditional volume, enter the following command:

```
smtape restore -g path
```

g is the volume geometry of the traditional volume that is to be displayed.

path is the path of the traditional volume.

You get an improved restore performance if you use the output of this command while backing up the data by using the `smtape backup -g` command.

Displaying the image header of a tape

You can display the image header of a tape in a specific tape device by using the `smtape restore` command.

Step

1. To display the image header of a tape in a tape device, enter the following command:

```
smtape restore -h tape_device
```

`h` is the image header of a tape that is to be displayed.

`tape_device` is the tape device that has the tape for which the image header is to be displayed.

Note: The image header of tape backups created by using the `snapmirror store` command in releases earlier than Data ONTAP 7.3 can be read by using the `smtape restore -h` command.

The following example displays the header image of a tape in the tape drive `rst1a`:

```
filer> smtape restore -h rst1a
Tape Number           : 1
WAFL Version          : 21054
BareMetal Version     : 9
Source Filer          : filer
Source Volume         : testdata
Source Volume Capacity : 51200MB
Source Volume Used Size : 3407MB
Source Snapshot       :
snapshot_for_smtape.db6bb83a-0b99-11de-
a2dc-00a              980de1c2.0
Volume Type           : Flexible
Is Aggregate          : no
Is SIS Volume         : no
Backup Set ID         : d7b1812a-0f90-11de-a2dc-00a0980de1c2
Backup Version        : 0:0
Backup Sequence No.   : 0
Backup Mode           : dw-data
Time of Backup        : Wed Mar 11 05:36:12 GMT 2009
Time of Previous Backup : None
Volume Total Inodes   : 1638399
Volume Used Inodes    : 102
Volume Attributes     : /
Number of Snapshots   : 1
Snapshot ID           : 76
Snapshot Time         : Fri Mar 6 04:30:31 GMT 2009
Snapshot Name         :
snapshot_for_smtape.db6bb83a-0b99-11de-
a2dc-00a
```

```
980de1c2.0
```

Restoring data from tape using SMTape

You can perform a level-0 restore of a backup image in a specific tape device to a destination volume using the `smtape restore` command.

About this task

The `smtape restore` command works the same way as the SnapMirror to Tape restore and provides users with the ability to initialize a volume SnapMirror destination volume by using backup images from tapes. After the restore, a volume SnapMirror relationship can be established between the source volume and the destination volume through the `snapmirror` commands. Prior to a restore operation, the volume must be in restricted mode. Any existing data on the volume is overwritten during the restore. The volume stays restricted during the restore operation and the restored volume is in the read-only state after a successful restore.

Note: Tape backups created by using the `snapmirror store` command in releases earlier than Data ONTAP 8.0 can be restored by using the `smtape restore` command.

Step

1. To restore data from tape to a destination volume, enter the following command:

```
smtape restore [-b block_size] path tape_device
```

b specifies that the tape record size to be used is supplied in the command line.

block_size is the blocking factor that was used during the backup.

path is the path to which the data has to be restored to.

tape_device contains the data to be restored.

A unique job ID in the range of 1 to 99999 is assigned to this restore operation. You can subsequently use this job ID to check the restore status or to abort the restore operation. Also, an entry is made in the `/etc/log/backup` file.

The following example shows how to restore the data in `rst1a` tape drive to the `/vol/testdata` volume

```
filer>smtape restore /vol/testdata rst1a
Job 10 started
```

Related tasks

[Backing up data to tape using SMTape](#) on page 130

Aborting a backup or restore operation using smtape abort command

You can abort a backup or restore operation using the `smtape abort` command. To abort a backup or restore operation, you must know its job ID.

Step

1. To abort a backup or restore operation, enter the following command:

```
smtape abort job_id
```

Note: To abort an SMTape backup or restore operation initiated through NDMP, you must also terminate the associated NDMP session.

The specified job is aborted and an entry is made in the `/etc/log/backup` file.

The following example aborts the SMTape operation with job ID 9.

```
filer>smtape abort 9
Job 9 aborted
```

Related tasks

[Terminating an NDMP session](#) on page 51

Continuing a backup or restore after reaching the end of tape

You can continue a backup or restore operation after it has reached the end of current tape and is in the wait state to write output to or accept input from a new tape.

About this task

When an SMTape backup or restore operation reaches the end of tape, and the backup or restore operation requires more than one tape to complete, one of the following messages is displayed on the console:

```
Change tape for smtape backup with job id <job ID>
```

```
Change tape for smtape restore with job id <job ID>
```

To continue your backup or restore operation, you must change the tape and use the `smtape continue` command.

Step

1. To continue your backup or restore operation after changing the tape, enter the following command:

```
smtape continue job_id [tape_device]
```

job_id is the job ID of the backup or restore operation to be continued.

tape_device is the tape device to be used to continue with the backup or restore operation. If you do not specify a tape device, the current tape device is used.

Displaying the status of SMTape backup and restore operations

You can display the status of backup and restore operations using the `smtape status` command. You can display the status for a specific job ID or for a specific backup or restore path.

Step

1. To display the status of backup and restore operations, enter the following command:

```
smtape status [-l] [[-p path] | [job_id]]
```

`l` displays a detailed status.

`p` displays the status of a specific path.

path is the path for which the status must be displayed.

job_id is the job ID for which the status must be displayed.

The following example displays the status of current backup and restore jobs:

```
filer>smtape status
Job ID  Seq No  Type      Status Path          Device
Progress
      1      0  Backup  Active /vol/vol0/  urst0a      240 MB
      2      0  Restore Active /vol/vol1/  urst1a      201 MB
```

The following example displays a detailed status for the backup job ID 3:

```
filer>smtape status -l 3
Job ID:          3
Sequence No:     0
Type:            Backup
Status:          Active
Path:            /vol/testdata
Device:          rstla
Progress:        1243360 KB
Job Begin:       Wed Mar 11 06:08:01 GMT 2009
Job End:         -
Last Update Time: Wed Mar 11 06:08:14 GMT 2009
```

Removing the SnapMirror status entries

When you perform an SMTape backup or restore, a SnapMirror status entry is created in the storage system. If you do not want to use SMTape backup for tape seeding, you must remove the SnapMirror status entries after the backup or restore completes.

How to remove the SnapMirror status entries created during backup

After a successful SMTape backup, the base Snapshot copy and the SnapMirror status entry are retained in the volume, which can be used to initialize a volume SnapMirror relationship.

Note: In case of a failed backup, the base Snapshot copy is automatically deleted. However, the SnapMirror status entry is retained.

You can delete a SnapMirror status entry by using the `snapmirror release` command or by deleting the Snapshot copy.

How to remove the SnapMirror status entries created during restore

After a successful SMTape restore, a SnapMirror status entry is created in the storage system. This SnapMirror status entry lists the restored volume and the base Snapshot copy name that is used to synchronize the SnapMirror source and destination volumes during tape seeding. To remove this SnapMirror status entry, you must make the restored volume writable, and then delete the base Snapshot copy. After the Snapshot copy is deleted, the SnapMirror status entry is automatically removed.

Removing the SnapMirror status entry after an SMTape backup

You can remove the SnapMirror status entry corresponding to the volume you backed up.

Steps

1. To list the SnapMirror status entries, enter the following command:

```
snapmirror status vol_name
```

vol_name is the name of the volume that you backed up.

The SnapMirror status of the volume is displayed. In case of a successful SMTape backup, the source is the volume being backed up and the destination is a Snapshot copy. This Snapshot copy has a name in the `snapmirror_tape_hexchar` format, in which *hexchar* is a set of hexadecimal characters specific to the Snapshot copy. In case of a failed SMTape backup, the source is the volume being backed up and the destination is a tape name.

2. To remove the SnapMirror status entry by releasing the SnapMirror relationship, enter one of the following commands:

If the backup... Then enter the following commands:

Is successful **snapmirror release vol_name**
 snapmirror_tape_hexadecimal_char

fails **snapmirror release vol_name filer_name:tape_device**

vol_name is the volume being backed up.

filer_name is the name of the storage system to which the tape device is attached.

tape_device is the tape device to which the volume is backed up.

When the backup is successful and the SnapMirror status entry is removed by using the `snapmirror release` command, SMTape deletes the auto-Snapshot copy. If you specify a Snapshot copy name for the backup, this Snapshot copy is not deleted, but the softlock on this Snapshot copy is removed.

The following example removes the SnapMirror status entry for the testdata volume that was successfully backed up:

```
filer1>snapmirror status testdata

Snapmirror is on.
Source           Destination                               State  Lag   Status
filer1:testdata  snapmirror_tape_2b8da4a4-1fa9-11de-842e-000c29d658dc Source  0:02:31 Idle

filer1>snapmirror release testdata snapmirror_tape_2b8da4a4-1fa9-11de-842e-000c29d658dc
```

The following example removes the SnapMirror status entry for the testdata2 volume that failed during the backup:

```
filer1>snapmirror status testdata2

Snapmirror is on.
Source           Destination           State  Lag   Status
filer1:testdata2  filer1:rst1a         Source  -    Idle

filer1>snapmirror release testdata2 filer1:rst1a
```

In case of a failed backup, though the SnapMirror status entry is deleted, the SnapMirror release command displays an error message as follows:

```
snapmirror release: testdata2 filer1:rst1a: No release-able destination found that matches those
parameters.
Use 'snapmirror destinations' to see a list of release-able destinations.
```

Removing the SnapMirror status entry after an SMTape restore

You can remove the SnapMirror status entry corresponding to the volume you restored.

Steps

1. To list the SnapMirror status entries, enter the following command:
snapmirror status vol_name

vol_name is the name of the volume that you restored.

The SnapMirror status of the volume is displayed. In case of a successful SMTape restore, the source is a Snapshot copy with a name in the format *snapshot_for_smtape.hexchar* and the destination is the restored volume. The *hexchar* in the Snapshot copy name is a set of hexadecimal characters specific to that Snapshot copy. In case of a failed SMTape restore, the source is a tape device name and the destination is the volume that failed to restore.

- To remove the SnapMirror status entry by releasing the SnapMirror relationship, enter one of the following commands:

| If the restore... | Then... |
|-------------------|---|
| Succeeded | <ol style="list-style-type: none"> To break the SnapMirror relationship, enter the following command: <code>snapmirror break vol_name</code> To remove the SnapMirror status entry, enter the following command: <code>snap delete vol_name snapshot_for_smtape.hexadecimal_char</code> |
| Failed | <p>Destroy the volume.</p> <p>Note: It is safe to destroy the volume because it is not useful due to the failed restore.</p> |

The following example removes the SnapMirror status entry for the testdata volume that was successfully restored.

```
filer1>snapmirror status testdata
Snapmirror is on.
Source                                     Destination      State           Lag
Status
snapshot_for_smtape.3fde069c-2639-11de-90f6-00a0980c225b.0  filer1:testdata  Snapmirrored   00:15:12
Idle

filer1>snapmirror break testdata
snapmirror break: Destination testdata is now writable.
Volume size is being retained for potential snapmirror resync.
If you would like to grow the volume and do not expect to resync,
set vol option fs_size_fixed to off.

filer1>snap delete testdata snapshot_for_smtape.3fde069c-2639-11de-90f6-00a0980c225b.0
Wed Apr  8 18:57:41 PDT [fsr-u29: waf1.snap.delete:info]: Snapshot copy snapshot_for_smtape.
3fde069c-2639-11de-90f6-00a0980c225b.0
on volume testdata was deleted by the Data ONTAP function snapcmd_delete. The unique ID for this
Snapshot copy is (1, 11).
```

Enabling or disabling concurrent volume SnapMirror and SMTape backup operations

Starting with Data ONTAP 8.1, you can make an SMTape backup of a volume SnapMirror destination when SnapMirror transfers are in progress. You can run volume SnapMirror and SMTape

backup operations concurrently by enabling the `vsm.smtape.concurrent.cascade.support` option on a volume SnapMirror destination system.

About this task

- The default value for the `vsm.smtape.concurrent.cascade.support` option is `off`. Any change in the option takes effect in the next volume SnapMirror or SMTape backup operation. It does not affect the operations that are in progress.
- When the `vsm.smtape.concurrent.cascade.support` option is enabled, SMTape backup locks only the base and incremental Snapshot copies. Any of the intermediate Snapshot copies (Snapshot copies between the base Snapshot copy and incremental Snapshot copy) can be deleted and deleted Snapshot copies are not backed up to tape.

Note: The `smtape restore -h` command lists the Snapshot copies that are present at the start of SMTape backup. Therefore, the list might include the Snapshot copies that are deleted by the SnapMirror update.
- The first SnapMirror update after a SnapMirror resync operation and SMTape backup operation cannot run concurrently.
- Concurrent volume SnapMirror and SMTape backup operations are supported only on FlexVol volumes and not on traditional volumes.

Step

1. To enable or disable concurrent volume SnapMirror and SMTape backup operations, enter the following command:

```
options vsm.smtape.concurrent.cascade.support {on|off}
```

`on` enables volume SnapMirror and SMTape backup operations to run concurrently.

`off` disables volume SnapMirror and SMTape backup operations to run concurrently.

Performing SMTape restores

You can perform baseline and incremental restores by using the SMTape engine to restore data on a volume from tape. You must follow a certain workflow to perform baseline and incremental restores.

Performing a baseline restore

You can perform a baseline restore if you want to restore the entire data set that has been backed up. You must follow a workflow to perform a baseline restore of data on a volume from tape.

Steps

1. Restrict the volume.
2. Restore the volume from tape.

After restore is complete, the volume is in a read-only state.

3. Make the volume writeable by either using the `snapmirror break` command or Data ONTAP APIs.

Performing an incremental restore

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

Steps

1. Restrict the volume.
2. Perform a baseline restore of the volume from tape.

The volume is in a read-only state.

Note: You must not change the state of the volume or use the `snapmirror break` command on the volume.

3. Perform incremental restores of the volume from tape in a chronological order.
4. After completing the required number of incremental restores, you can make the volume writeable by either using the `snapmirror break` command or Data ONTAP APIs.

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

Once every week, the event log files are rotated. The `/etc/log/backup` file is renamed to `/etc/log/backup.0`, the `/etc/log/backup.0` file is renamed 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.[0-5]` and the current `/etc/log/backup` file).

If a takeover occurs in an HA pair, the set of backup log files for the takeover storage system remains separate from the backup log files for the failed storage system.

Accessing the event log files

You can access the event log files for tape backup and restore operations at the `/etc/log/` directory by using the `rdfile` command. You can view these event log files to monitor tape backup and restore operations.

Step

1. You can access the event log files for tape backup and restore operations by entering the following command:

```
rdfile /etc/log/backup
```

With additional configurations, you can also use a web browser to access these log files. For more information about accessing a node's log files by using a web browser, see the *Data ONTAP System Administration Guide for 7-Mode*.

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 |
| 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.

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 is started | Dump level and the type of dump |
| Restart | A dump restarts | Dump level |
| End | Dumps completed successfully | Amount of data processed |
| Abort | The operation is cancelled | Amount of data processed |
| Options | Specified options are listed | All options and their associated values, including NDMP options |

| Event | Description | Event information |
|--------------|--|--|
| 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 <code>/etc/dumpdates</code> file 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/vol0/(1) End (1224 MB)
```

There are five phases in a dump operation (map files, map directories, dump directories, dump files, and dump ACLs).

The log file for a dump operation begins with either a Start or Restart event and ends with either an End or Abort event.

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 is started | 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 is cancelled | 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 a cancelled 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)
```

What the SMTape event log message format is

For each SMTape event, a message is written to the backup log file in a specified format.

The format of the SMTape event log message is as follows:

```
job_id time_stamp vol_path event(event_info)
```

The following list describes the fields in the event log message format.

- The *job_id* field shows the unique ID allocated to the SMTape backup or restore job.
- The *time_stamp* field shows the date and time at which SMTape backup or restore event occurred.
- The *vol_path* is the volume path associated with the SMTape backup or restore job.
- The *event* field shows the event name.
- The *event_info* field shows the event specific information.

What SMTape CLI backup and restore events are

The event field for an SMTape backup or restore event begins with a CLI event type followed by event-specific information within parentheses.

The following table describes the CLI events and their descriptions recorded for an SMTape backup and restore operation initiated from the CLI. The event information for these events is the tape device name.

| Event | Description |
|--------------|--|
| CLI-Backup | The SMTape backup operation is initiated by using the <code>smtape backup</code> command. |
| CLI-Restore | The SMTape restore operation is initiated by using the <code>smtape restore</code> command. |
| CLI-Abort | The SMTape backup or restore operation is aborted by using the <code>smtape abort</code> command. |
| CLI-Continue | The SMTape backup or restore operation is continued after a tape change by using the <code>smtape continue</code> command. |

What SMTape backup events are

The event field for an SMTape backup 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 are recorded for an SMTape backup operation.

| Event | Description | Event information |
|----------------|--|---|
| BKP-Start | An SMTape CLI or NDMP backup begins | The level of backup and the backup set ID that identifies the backup session. |
| BKP-Params | The parameters for the backup job | Parameters of the backup operation, such as the origin of the command that specifies whether the command was initiated from NDMP or CLI, the tape record size used in the backup, and the tape device name. |
| BKP-DW-Start | Data warehouse begins for the backup job | Does not have any event information. |
| BKP-DW-End | Data warehouse ends for the backup job | Time taken to complete the backup job and the number of blocks backed up to tape. |
| BKP-Tape-Stats | The tape statistics for the backup job | The backup statistics, such as the wait time, the wait count, total count, and the available count. |

| Event | Description | Event information |
|--------------|--|--|
| BKP-End | The backup job ends | The amount of data backed up to tape, the time taken for the backup, and the performance in GB/hour. |
| BKP-Abort | The backup job aborts | A message indicating the reason for aborting the backup job. |
| BKP-Tape-Chg | The backup job is waiting for a tape change | The job ID of the backup operation that waits for a tape change. |
| BKP-Continue | The backup operation continues after a tape change | The job ID of the backup operation that continues after a tape change. |
| BKP-Warning | The backup operation has encountered an unexpected event | The reason for the unexpected event. |

Example of an SMTape backup output

The following is an example of the output for an SMTape backup operation:

```
(null) Tue May 5 11:15:00 PDT /vol/testdata CLI-Backup (rst9a)
1 Tue May 5 11:15:00 PDT /vol/testdata BKP-Start (level 0 backup of Backup Set ID
f99f17ac-3b32-11de-9682-00a0980c225b)
1 Tue May 5 11:15:00 PDT /vol/testdata BKP-Params (originator=CLI mode=dw-data
tape_record_size=240KB tape=rst9a all_snapshots tape_seeding)
1 Tue May 5 11:15:00 PDT /vol/testdata BKP-DW-Start
1 Tue May 5 11:27:04 PDT /vol/testdata BKP-DW-End (phase completed in 0:12:04; 9214285 blocks moved)
1 Tue May 5 11:27:04 PDT /vol/testdata BKP-Tape-Stats (wait_time=684s wait_count=128990
total_count=153951 avail_count=130809/17974/4906/262)
1 Tue May 5 11:27:04 PDT /vol/testdata BKP-End (backed up 36.857 GB bytes in 0:12:04;
performance=183.267 GB/hour)
```

What SMTape restore events are

The event field for an SMTape restore event contains an event type followed by event-specific information within parentheses.

The following table describes the SMTape restore events, their descriptions, and the related event information that are recorded for an SMTape restore operation:

| Event | Description | Event information |
|-----------|--------------------------------------|---|
| RST-Start | An SMTape CLI or NDMP restore begins | The restore set ID that identifies the restore session. |

| Event | Description | Event information |
|--------------|---|---|
| RST-Params | The parameters for the restore job | Parameters of the restore operation, such as the origin of the command that specifies whether the command was initiated from NDMP or CLI, the tape record size for the restore, and the tape device name. |
| RST-End | A restore job completed successfully | The amount of data restored from tape, the time taken for the restore, and the performance in GB/hour. |
| RST-Tape-Chg | The restore job is waiting for a tape change | The job ID of the restore operation that waits for a tape change. |
| RST-Continue | The restore job continues after a tape change | The job ID of the restore operation that continues after a tape change. |
| RST-Abort | The restore job aborts | A message indicating the reason for aborting the restore job. |

Example of an SMTape restore output

The following is an example of the output for an SMTape restore operation:

```
(null) Thu May  7 18:41:52 PDT /vol/testdata CLI-Restore (rst8a)
29 Thu May  7 18:41:52 PDT /vol/testdata RST-Start (Restore Set ID bc24cbb0-3d03-11de-
bef3-00a0980c225b)
29 Thu May  7 18:41:52 PDT /vol/testdata RST-Params (originator=CLI mode=image tape_record_size=240KB
tape=rst8a)
29 Thu May  7 18:42:01 PDT /vol/testdata RST-End (restored 399.840 MB bytes in 0:00:09;
performance=159.936 GB/hour)
```

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.

Error messages for tape backup and restore

You might encounter an error message when performing a dump or SMTape-based backup or restore operation due to various reasons.

Backup and restore error messages

You might encounter an error message while performing a tape backup or restore using SMTape or dump.

Resource limitation: no available thread

| | |
|--------------------------|---|
| Message | Resource limitation: no available thread |
| Cause | The maximum number of active local tape I/O threads is currently in use. You can have a maximum of 16 active local tape drives. |
| Corrective action | Wait for some tape jobs to finish before starting a new backup or restore job. |

Duplicated tape drive (*tape_drive*) specified in the tape argument list

| | |
|--------------------------|--|
| Message | Duplicated tape drive (<i>tape_drive</i>) specified in the tape argument list |
| Cause | You have specified a tape drive name twice in the argument list of the backup or restore command. If a tape drive name is duplicated in the <code>dump</code> or <code>smtape backup</code> command, data is backed up twice to the tape attached to that tape drive. If the tape drive name is duplicated in the <code>restore</code> or <code>smtape restore</code> command, data is restored twice to the destination. |
| Corrective action | Retry the job without specifying the same tape drive more than once in the tape argument list. |

Invalid tape drive *tape_drive* in tape argument list

| | |
|--------------------------|--|
| Message | Invalid tape drive <i>tape_drive</i> in tape argument list |
| Cause | The tape drive specified for the backup or restore operation is not valid. |
| Corrective action | Use a valid tape drive and retry the operation. Use the <code>sysconfig -t</code> command to get a list of valid tape drives. |

Tape reservation preempted

| | |
|--------------------------|--|
| Message | Tape reservation preempted |
| Cause | The tape drive is in use by another operation or the tape has been closed prematurely. |
| Corrective action | Ensure that the tape drive is not in use by another operation and that the DMA application has not aborted the job and then retry. |

Could not initialize media

| | |
|--------------------------|--|
| Message | Could not initialize media |
| Cause | You might get this error for one of the following reasons: <ul style="list-style-type: none"> • The tape drive used for the backup is corrupt or damaged. • The tape does not contain the complete backup or is corrupt. • The maximum number of active local tape I/O threads is currently in use. You can have a maximum of 16 active local tape drives. |
| Corrective action | <ul style="list-style-type: none"> • If the tape drive is corrupt or damaged, retry the operation with a valid tape drive. • If the tape does not contain the complete backup or is corrupt, you cannot perform the restore operation. • If tape resources are not available, wait for some of the backup or restore jobs to finish and then retry the operation. |

Too many active dumps/restores currently in progress

| | |
|--------------------------|---|
| Message | Too many active dumps/restores currently in progress |
| Cause | A maximum number of backup and/or restore jobs are already running. |
| Corrective action | Retry the operation after some of the currently running jobs have finished. |

Media error on tape write

| | |
|--------------------------|--|
| Message | Media error on tape write |
| Cause | The tape used for the backup is corrupted. |
| Corrective action | Replace the tape and retry the backup job. |

Tape write failed

| | |
|--------------------------|--|
| Message | Tape write failed |
| Cause | The tape used for the backup is corrupted. |
| Corrective action | Replace the tape and retry the backup job. |

Tape write failed - new tape encountered media error

| | |
|--------------------------|--|
| Message | Tape write failed - new tape encountered media error |
| Cause | The tape used for the backup is corrupted. |
| Corrective action | Replace the tape and retry the backup. |

Tape write failed - new tape is broken or write protected

| | |
|--------------------------|---|
| Message | Tape write failed - new tape is broken or write protected |
| Cause | The tape used for the backup is corrupted or write-protected. |
| Corrective action | Replace the tape and retry the backup. |

Tape write failed - new tape is already at the end of media

| | |
|--------------------------|---|
| Message | Tape write failed - new tape is already at the end of media |
| Cause | There is not enough space on the tape to complete the backup. |
| Corrective action | Replace the tape and retry the backup. |

Tape write error

| | |
|--------------------------|---|
| Message | Tape write error - The previous tape had less than the required minimum capacity, <i>size</i> MB, for this tape operation, The operation should be restarted from the beginning |
| Cause | The tape capacity is insufficient to contain the backup data. |
| Corrective action | Use tapes with larger capacity and retry the backup job. |

Media error on tape read

| | |
|----------------|--------------------------|
| Message | Media error on tape read |
|----------------|--------------------------|

| | |
|--------------------------|---|
| Cause | The tape from which data is being restored is corrupted and might not contain the complete backup data. |
| Corrective action | If you are sure that the tape has the complete backup, retry the restore operation. If the tape does not contain the complete backup, you cannot perform the restore operation. |

Tape read error

| | |
|--------------------------|--|
| Message | Tape read error |
| Cause | The tape drive is damaged or the tape does not contain the complete backup. |
| Corrective action | If the tape drive is damaged, use another tape drive. If the tape does not contain the complete backup, you cannot restore the data. |

Already at the end of tape

| | |
|--------------------------|---|
| Message | Already at the end of tape |
| Cause | The tape does not contain any data or must be rewound. |
| Corrective action | If the tape does not contain data, use the tape that contains the backup and retry the restore job. Otherwise, rewind the tape and retry the restore job. |

Tape record size is too small. Try a larger size.

| | |
|--------------------------|---|
| Message | Tape record size is too small. Try a larger size. |
| Cause | The blocking factor specified for the restore operation is smaller than the blocking factor that was used during the backup. |
| Corrective action | Use the same blocking factor that was specified during the backup. In case of an SMTape restore operation, use the <code>smtape restore -h tape_drive</code> command to determine the correct blocking factor. |

Tape record size should be `block_size1` and not `block_size2`

| | |
|--------------------------|--|
| Message | Tape record size should be <code>block_size1</code> and not <code>block_size2</code> |
| Cause | The blocking factor specified for the local restore is incorrect. |
| Corrective action | Retry the restore job with <code>block_size1</code> as the blocking factor. |

Tape record size must be in the range between 4KB and 256KB

| | |
|----------------|---|
| Message | Tape record size must be in the range between 4KB and 256KB |
|----------------|---|

Cause The blocking factor specified for the backup or restore operation is not within the permitted range.

Corrective action Specify a blocking factor in the range of 4 KB to 256 KB.

NDMP error messages

You might encounter an error message while performing a tape backup or restore using NDMP-enabled commercial backup applications.

Network communication error

Message `Network communication error`

Cause Communication to a remote tape in an NDMP three-way connection has failed.

Corrective action Check the network connection to the remote mover.

Message from Read Socket: `error_string`

Message `Message from Read Socket: error_string`

Cause Restore communication from the remote tape in NDMP 3-way connection has errors.

Corrective action Check the network connection to the remote mover.

Message from Write Dirnet: `error_string`

Message `Message from Write Dirnet: error_string`

Cause Backup communication to a remote tape in an NDMP three-way connection has an error.

Corrective action Check the network connection to the remote mover.

Read Socket received EOF

Message `Read Socket received EOF`

Cause Attempt to communicate with a remote tape in an NDMP three-way connection has reached the End Of File mark. You might be attempting a three-way restore from a backup image with a larger block size.

Corrective action Specify the correct block size and retry the restore operation.

ndmpd invalid version number: *version_number*

| | |
|--------------------------|--|
| Message | <code>ndmpd invalid version number: <i>version_number</i></code> |
| Cause | The NDMP version specified is not supported by the storage system. |
| Corrective action | Specify NDMP version 4. |

Error: Unable to retrieve session information

| | |
|--------------------------|--|
| Message | <code>Error: Unable to retrieve session information</code> |
| Cause | The system is probably overloaded. |
| Corrective action | Retry the operation. |

ndmpd session *session_ID* not active

| | |
|--------------------------|---|
| Message | <code>ndmpd session <i>session_ID</i> not active</code> |
| Cause | The NDMP session might not exist. |
| Corrective action | Use the <code>ndmpd status</code> command to view the active NDMP sessions. |

No such user *user_name*

| | |
|--------------------------|---|
| Message | <code>No such user <i>user_name</i></code> |
| Cause | The specified user might not exist. |
| Corrective action | Use the <code>useradmin</code> command to list the valid users of the system. |

Cannot generate NDMP password

| | |
|--------------------------|--|
| Message | <code>Cannot generate NDMP password</code> |
| Cause | The user might not have the <code>login-ndmp</code> capability. |
| Corrective action | Ensure that the user has the proper capabilities to access NDMP. |

The specified operation could not be completed as the volume is moving

| | |
|--------------------------|---|
| Message | <code>The specified operation could not be completed as the volume is moving</code> |
| Cause | The volume is under migration. |
| Corrective action | Retry the operation after the volume migration is complete. |

Could not obtain vol ref for Volume `volume_name`

| | |
|--------------------------|--|
| Message | Could not obtain vol ref for Volume <code>vol_name</code> |
| Cause | The volume reference could not be obtained because the volume might be in use by other operations. |
| Corrective action | Retry the operation later. |

ndmcopy error messages

You might encounter an error message while transferring data between storage systems using the `ndmcopy` command.

Ndmcopy: Socket connection to `host_name` failed

| | |
|--------------------------|---|
| Message | Ndmcopy: Socket connection to <code>host_name</code> failed |
| Cause | Unable to create a socket connection from the storage system to the <code>host_name</code> . |
| Corrective action | Ensure that you can ping the <code>host_name</code> from the storage system. Ensure that the NDMP server is up and running by using the <code>ndmpd status</code> command. If the NDMP server is not running, enable NDMP service. Ensure that the firewall settings for NDMP are configured on both the network and the storage system. Check whether the option <code>ndmpd.access</code> is set correctly. |

Ndmcopy: Error opening NDMP connection

| | |
|--------------------------|---|
| Message | Ndmcopy: Error opening NDMP connection |
| Cause | Unable to create a socket connection from the storage system to the host name. |
| Corrective action | Ensure that you can ping the <code>host_name</code> from the storage system. Ensure that the NDMP server is up and running by using the <code>ndmpd status</code> command. If the NDMP server is not running, enable NDMP service. Ensure that the firewall settings for NDMP are configured on both the network and the storage system. Check whether the option <code>ndmpd.access</code> is set correctly. |

Ndmcopy: Client authentication request failed

| | |
|----------------|---|
| Message | Ndmcopy: Client authentication request failed |
| Cause | Authentication parameters might be incorrect. |

Corrective action Ensure that the authentication parameters are correct. Verify the authentication type (plaintext or md5) by using the `options.ndmpd.authtype` command on the source and destination storage systems.

Ndmpcopy: Authentication failed for source

Message `Ndmpcopy: Authentication failed for source`

Cause Authentication parameters might be incorrect for the source.

Corrective action Ensure that the authentication parameters are correct. Verify the authentication type (plaintext or md5) by using the `options.ndmpd.authtype` command on the source storage system.

Ndmpcopy: Authentication failed for destination

Message `Ndmpcopy: Authentication failed for destination`

Cause Authentication parameters might be incorrect for the destination.

Corrective action Ensure that the authentication parameters are correct. Verify the authentication type (plaintext or md5) by using the `options.ndmpd.authtype` command on the destination storage system.

Ndmpcopy: Failed to start dump on source

Message `Ndmpcopy: Failed to start dump on source`

Cause Could not establish a data connection to the NDMP server.

Corrective action Ensure that the network connectivity between the source and the destination is appropriate. Also, ensure that the firewall settings for NDMP are configured on both the network and the storage system.

Ndmpcopy: Failed to start restore on destination

Message `Ndmpcopy: Failed to start restore on destination`

Cause Could not establish a data connection to the NDMP server.

Corrective action Ensure that the network connectivity between the source and the destination is appropriate. Also, ensure that the firewall settings for NDMP are configured on both the network and the storage system.

Ndmpcopy: Error in getting extension list

Message `Ndmpcopy: Error in getting extension list`

| | |
|--------------------------|---|
| Cause | Either the source or the destination does not support IPv6. |
| Corrective action | Retry the operation using IPv4. |

Error getting local hostname

| | |
|--------------------------|--|
| Message | Error getting local hostname |
| Cause | Local machine might not have a valid host name. |
| Corrective action | Ensure that the local machine has a valid host name. |

Ndmpcopy: Connection setup for transfer failed

| | |
|--------------------------|---|
| Message | Ndmpcopy: Connection setup for transfer failed |
| Cause | The ndmpcopy command failed to establish a data connection between the source and the destination. |
| Corrective action | Ensure that the network connectivity between the source and the destination is appropriate. Also, ensure that the firewall settings for NDMP are configured on both the network and the storage system. |

CONNECT: Connection refused

| | |
|--------------------------|---|
| Message | CONNECT: Connection refused |
| Cause | The NDMP server refuses connections in the following scenarios: <ul style="list-style-type: none">• The NDMP connections running on the server has reached the maximum limit.• The NDMP server is shutting down. |
| Corrective action | Retry the NDMP connection later. |

Invalid name. Source filer name does not resolve to the specified address mode

| | |
|--------------------------|--|
| Message | Invalid name. Source filer name does not resolve to the specified address mode |
| Cause | The storage system name does not resolve to the specified address mode (IPv4 or IPv6). |
| Corrective action | Retry the operation using the appropriate address mode. |

Invalid name. Destination filer name does not resolve to the specified address mode

- Message** Invalid name. Destination filer name does not resolve to the specified address mode
- Cause** The storage system name does not resolve to the specified address mode (IPv4 or IPv6).
- Corrective action** Retry the operation using the appropriate address mode.

Dump error messages

You might encounter an error message while performing a tape backup or restore using the dump engine.

No default tape device list

- Message** No default tape device list
- Cause** The tape device list specified in the `dump` command is incorrect.
- Corrective action** Specify a valid tape device list in the `dump` command and retry the backup.

Invalid/offline volume

- Message** Invalid/offline volume
- Cause** The volume specified in the `dump` command is offline or has been deleted.
- Corrective action** If the volume is offline, bring the volume back online and make the volume writable and then perform the backup. If the volume has been deleted, you cannot perform the backup.

Unable to lock a snapshot needed by dump

- Message** Unable to lock a snapshot needed by dump
- Cause** The Snapshot copy specified for the backup is not available.
- Corrective action** Retry the backup with a different Snapshot copy.
Use the `snap list` command to see the list of available Snapshot copies.

Failed to determine snapshot type

- Message** Failed to determine snapshot type

- Cause** The Snapshot copy specified for the backup is not available.
- Corrective action** Retry the backup with a different Snapshot copy.
Use the `snap list` command to see the list of available Snapshot copies.

Volume is temporarily in a transitional state

- Message** `Volume is temporarily in a transitional state`
- Cause** The volume being backed up is temporarily in an unmounted state.
- Corrective action** Wait for some time and perform the backup again.

Unable to locate bitmap files

- Message** `Unable to locate bitmap files`
- Cause** The bitmap files required for the backup operation might have been deleted. In this case, the backup cannot be restarted.
- Corrective action** Perform the backup again.

Failed to locate the specified restartable dump

- Message** `Failed to locate the specified restartable dump`
- Cause** The dump ID specified for restarting the failed backup is invalid.
- Corrective action** Restart the backup with the correct dump ID.
Use the `backup status` command to determine the dump ID of the failed backup that you are trying to restart.

Dump context created from NDMP. Cannot restart dump

- Message** `Dump context created from NDMP. Cannot restart dump.`
- Cause** The dump operation was initiated through NDMP, but you are attempting to restart it from CLI.
- Corrective action** Restart the dump operation through NDMP.

Unable to locate snapshot

- Message** `Unable to locate snapshot`
- Cause** The Snapshot copies required for restarting the backup are not available.
- Corrective action** Backup cannot be restarted. Perform the backup again.

Invalid inode specified on restart

| | |
|--------------------------|---|
| Message | <code>Invalid inode specified on restart</code> |
| Cause | The inode specified for the NDMP-initiated backup is invalid. |
| Corrective action | Try to restart the backup with a valid inode number and offset. |

Invalid restart context. Cannot restart dump

| | |
|--------------------------|--|
| Message | <code>Invalid restart context. Cannot restart dump.</code> |
| Cause | The registry might be corrupt. |
| Corrective action | Restart the backup again. If it fails, you must redo the backup. |

Failed to retrieve saved info for the restartable dump

| | |
|--------------------------|--|
| Message | <code>Failed to retrieve saved info for the restartable dump.</code> |
| Cause | The registry might be corrupt. |
| Corrective action | Restart the backup again. If it fails, you must redo the backup. |

Destination volume is read-only

| | |
|--------------------------|---|
| Message | <code>Destination volume is read-only</code> |
| Cause | The path to which the restore operation is attempted to is read-only. |
| Corrective action | Try restoring the data to a different location. |

Destination qtree is read-only

| | |
|--------------------------|--|
| Message | <code>Destination qtree is read-only</code> |
| Cause | The qtree to which the restore is attempted to is read-only. |
| Corrective action | Try restoring the data to a different location. |

IB restore in progress

| | |
|--------------------------|--|
| Message | <code>IB restore in progress</code> |
| Cause | An SMTape restore is currently running. You cannot perform a dump-based restore when an SMTape restore is running. |
| Corrective action | Retry the restore operation after the SMTape restore operation finishes. |

Could not access volume in path: volume_name

- Message** `Could not access volume in path: volume_name`
- Cause** The destination volume specified in the `restore` command does not exist.
- Corrective action** Try to restore the data to a different volume or create a new volume with the specified name.

No files were created

- Message** `No files were created`
- Cause** A directory DAR was attempted without enabling the enhanced DAR functionality.
- Corrective action** Enable the enhanced DAR functionality and retry the DAR.

SMTape error messages

You might encounter an error message while performing a tape backup or restore using SMTape.

Internal assertion

- Message** `Internal assertion`
- Cause** There is an internal SMTape error.
- Corrective action** Report the error and send the `etc/log/backup` file to technical support.

Job aborted due to shutdown

- Message** `Job aborted due to shutdown`
- Cause** The storage system is being rebooted.
- Corrective action** Retry the job after the storage system reboots.

Job not found

- Message** `Job not found`
- Cause** The backup or restore job is not active.
- Corrective action** Check the job number and retry the job.

Job aborted due to Snapshot autodelete

- Message** Job aborted due to Snapshot autodelete
- Cause** The volume does not have enough space and has triggered the autodeletion of Snapshot copies.
- Corrective action** Free up space in the volume and retry the job.

Invalid volume path

- Message** Invalid volume path
- Cause** The specified volume for the backup or restore operation is not found.
- Corrective action** Retry the job with a valid volume path and volume name.

UNIX style RMT tape drive is not supported

- Message** UNIX style RMT tape drive is not supported
- Cause** A remote tape drive was specified for the backup or restore job.
- Corrective action** SMTape does not support remote tapes. Use a local tape drive for the job.

Volume is currently in use by other operations

- Message** Volume is currently in use by other operations
- Cause** The volume is currently in use by another SnapMirror operation. You cannot perform an SMTape operation when another SnapMirror operation is using the volume.
- Corrective action** Retry the job after the SnapMirror operation finishes.

Volume offline

- Message** Volume offline
- Cause** The volume being backed up is offline.
- Corrective action** Bring the volume online and retry the backup.

Volume not restricted

- Message** Volume not restricted
- Cause** The destination volume to which data is being restored is not restricted.

Corrective action Restrict the volume and retry the restore operation.

Tape is currently in use by other operations

Message Tape is currently in use by other operations

Cause The tape drive is in use by another job.

Corrective action Retry the backup after the currently active job is finished.

Invalid input tape

Message Invalid input tape

Cause The signature of the backup image is not valid in the tape header. The tape has corrupted data or does not contain a valid backup image.

Corrective action Retry the restore job with a valid backup image.

Too many active jobs

Message Too many active jobs

Cause A maximum number of SMTape jobs are already running. You can have a maximum of 32 SMTape jobs running simultaneously.

Corrective action Retry the operation after some of the SMTape jobs have finished.

Failed to allocate memory

Message Failed to allocate memory

Cause The system has run out of memory.

Corrective action Retry the job later when the system is not too busy.

Failed to get data buffer

Message Failed to get data buffer

Cause The storage system ran out of buffers.

Corrective action Wait for some storage system operations to finish and then retry the job.

Failed to create job UUID

Message Failed to create job UUID

Cause The storage system could not create an UUID because the system is too busy.

Corrective action Reduce the system load and then retry the job.

Failed to create snapshot

Message Failed to create snapshot

Cause The volume already contains the maximum number of Snapshot copies.

Corrective action Delete some Snapshot copies and then retry the backup operation.

Failed to find snapshot

Message Failed to find snapshot

Cause The Snapshot copy specified for the backup is unavailable.

Corrective action Check if the specified Snapshot copy is available. If not, retry with the correct Snapshot copy.

Failed to lock snapshot

Message Failed to lock snapshot

Cause The Snapshot copy is either in use or has been deleted.

Corrective action If the Snapshot copy is in use by another operation, wait for that operation to finish and then retry the backup. If the Snapshot copy has been deleted, you cannot perform the backup.

Failed to access the named snapshot

Message Failed to access the named snapshot

Cause The Snapshot copy might have been deleted.

Corrective action If the Snapshot copy was deleted, you cannot perform the backup operation. If the Snapshot copy exists, retry the job.

Failed to softlock qtree snapshots

Message Failed to softlock qtree snapshots

Cause The Snapshot copy is in use or the Snapshot copy has been deleted.

Corrective action If the Snapshot copy is in use by another operation, wait for that operation to finish and then retry the job. If the Snapshot copy has been deleted, you cannot perform the backup operation.

Failed to delete softlock

- Message** `Failed to delete softlock`
- Cause** The system could not remove the softlock for a Snapshot copy.
- Corrective action** If the Snapshot copy is no longer required, delete the softlock manually by using the `registry` command.

Failed to delete snapshot

- Message** `Failed to delete snapshot`
- Cause** The auto-Snapshot copy could not be deleted because it is in use by other operations.
- Corrective action** Use the `snap` command to determine the status of the Snapshot copy. If the Snapshot copy is not required, delete it manually.

Image header missing or corrupted

- Message** `Image header missing or corrupted`
- Cause** The tape does not contain a valid SMTape backup.
- Corrective action** Retry with a tape containing a valid backup.

Chunks out of order

- Message** `Chunks out of order`
- Cause** The backup tapes are not being restored in the correct sequence.
- Corrective action** Retry the restore operation and load the tapes in the correct sequence.

Tapes out of order

- Message** `Tapes out of order`
- Cause** The first tape of the tape sequence for the restore operation does not have the image header.
- Corrective action** Load the tape with the image header and retry the job.

Already read volume_name tape_number

- Message** `Already read volume_name tape_number`

| | |
|--------------------------|---|
| Cause | The tape has already been processed. |
| Corrective action | Be sure to load the correct tape when changing tapes. |

Mismatch in backup set ID

| | |
|--------------------------|---|
| Message | Mismatch in backup set ID |
| Cause | The tape loaded during a tape change is not a part of the backup set. |
| Corrective action | Load the correct tape and retry the job. |

Aborting: Destination volume, `volume_name`, is too small

| | |
|--------------------------|---|
| Message | Aborting: Destination volume, <code>volume_name</code> , is too small |
| Cause | The destination volume for the restore is not large enough for the backed up data. |
| Corrective action | Create a larger volume for the restore job. Use the <code>smtape restore -h <i>tape_drive</i></code> command to determine the volume size of the backup image. |

Aborting: Destination volume, `volume_name`, is a clone

| | |
|--------------------------|--|
| Message | Aborting: Destination volume, <code>volume_name</code> , is a clone |
| Cause | You might be trying to restore an SMTape backup to a FlexClone volume. SMTape does not support data restoration to a FlexClone volume. |
| Corrective action | Try to restore the data to a regular FlexVol volume. |

Source volume size is greater than maximum supported SIS volume size on this platform. Aborting

| | |
|--------------------------|---|
| Message | Source volume size is greater than maximum supported SIS volume size on this platform. Aborting |
| Cause | The backup image is from a SIS volume (deduplication-enabled volume) that is larger than the maximum size supported by the restore volume. The maximum volume size when deduplication is enabled depends on the platform that you are using. For more information about the maximum volume size supported for different storage systems when deduplication is enabled, see <i>Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode</i> . |
| Corrective action | Restore the backup image on a platform that allows larger deduplication-enabled volumes. |

Incompatible SnapMirror or copy source Version. Aborting

- Message** `Incompatible SnapMirror or copy source Version. Aborting`
- Cause** The tape contains an incompatible backup image. The backup image is generated from a newer version of Data ONTAP.
- Corrective action** Use the correct Data ONTAP version to restore the backup image.

Transfers from volume `volume_name` are temporarily disabled

- Message** `Transfers from volume volume_name are temporarily disabled`
- Cause** The volume is currently in use by other operations.
- Corrective action** Wait for other operations to finish and then retry the job.

Too many active transfers at once, aborting

- Message** `Too many active transfers at once, aborting`
- Cause** A maximum number of SMTape and SnapMirror transfers are already running.
- Corrective action** Retry the operation after some of the SMTape and SnapMirror transfers have finished.

Invalid contents in destination volume geometry string `volume_geometry_string`, aborting

- Message** `Invalid contents in destination volume geometry string
volume_geometry_string, aborting`
- Cause** The format of the destination volume geometry string might be incorrect.
- Corrective action** Use the `smtape restore -g destination_volume_path` command to identify the destination geometry.

Cannot init input, aborting

- Message** `Cannot init input, aborting`
- Cause** Tape read/write operation fails or the tape is not connected properly to the storage system.
- Corrective action** Ensure that you can perform tape read/write operation using the `dump` or `restore` command.

Source volume is not a flexible volume. Aborting

| | |
|--------------------------|--|
| Message | Source volume is not a flexible volume. Aborting |
| Cause | The tape contains the backup of a traditional volume and you might be trying to restore the backup image to a flexible volume. |
| Corrective action | Try restoring the backup image to a traditional volume. Use the <code>smtape restore -h <i>tape_device</i></code> command and see the Volume Type field to determine the type of volume. |

Source volume is a flexible volume. Aborting

| | |
|--------------------------|---|
| Message | Source volume is a flexible volume. Aborting |
| Cause | The tape contains the backup of a flexible volume and you might be trying to restore the backup image to a traditional volume. |
| Corrective action | Try restoring the backup image to a flexible volume. Use the <code>smtape restore -h <i>tape_device</i></code> command and see the Volume Type field to determine the type of volume. |

Destination is not an aggregate. Aborting

| | |
|--------------------------|--|
| Message | Destination is not an aggregate. Aborting |
| Cause | The tape contains the backup of an aggregate and you might be trying to restore the backup image to a volume. |
| Corrective action | Try restoring the backup image to an aggregate. Use the <code>smtape restore -h <i>tape_device</i></code> command and see the Volume Type field to determine the type of volume. |

Source is not an aggregate. Aborting

| | |
|--------------------------|--|
| Message | Source is not an aggregate. Aborting |
| Cause | The tape contains the backup of a volume and you might be trying to restore the backup image to an aggregate. |
| Corrective action | Try restoring the backup image to a volume. Use the <code>smtape restore -h <i>tape_device</i></code> command and see the Volume Type field to determine the type of volume. |

Source is not a hybrid aggregate. Aborting

| | |
|----------------|--|
| Message | Source is not a hybrid aggregate. Aborting |
|----------------|--|

- Cause** The backup image might not be of a hybrid aggregate and you are trying to restore to a hybrid aggregate.
- Corrective action** Use the `smtape restore -h tape_device` command and see the volume. Type field to determine the type of volume.

Invalid checksum for the chunk descriptor

- Message** `Invalid checksum for the chunk descriptor`
- Cause** You might get this error message for one of the following reasons:
- The tape is not positioned at the location of the backup image.
 - The tape is corrupt or damaged.
 - The wrong tape is loaded for restore.
- Corrective action**
- If the tape is not positioned correctly, position the tape at the location of the backup image and retry the operation.
 - If the tape is corrupt, you cannot perform the restore operation.
 - If the wrong tape is loaded, retry the operation with the correct tape.

Received VBN header with invalid checksum `error_string`, aborting transfer on volume `volume_name`

- Message** `Received VBN header with invalid checksum error_string, aborting transfer on volume volume_name`
- Cause** You might get this error message for one of the following reasons:
- The tape is not positioned at the location of the backup image.
 - The tape is corrupt or damaged.
 - The wrong tape is loaded for restore.
- Corrective action**
- If the tape is not positioned correctly, position the tape at the location of the backup image and retry the operation.
 - If the tape is corrupt, you cannot perform the restore operation.
 - If the wrong tape is loaded, retry the operation with the correct tape.

Duplicate VBN `VBN_number` received for volume `volume_name`, aborting transfer

- Message** `Duplicate VBN VBN_number received for volume volume_name, aborting transfer`
- Cause** You might get this error message for one of the following reasons:

- The tape is not positioned at the location of the backup image.
- The tape is corrupt or damaged.
- The wrong tape is loaded for restore.

Corrective action

- If the tape is not positioned correctly, position the tape at the location of the backup image and retry the operation.
- If the tape is corrupt, you cannot perform the restore operation.
- If the wrong tape is loaded, retry the operation with the correct tape.

Bad block in read stream. VBN = VBN_number, max_VBN = max_VBN_number

Message Bad block in read stream. VBN = *VBN_number*, max_VBN = *max_VBN_number*

Cause You might get this error message for one of the following reasons:

- The tape is not positioned at the location of the backup image.
- The tape is corrupt or damaged.
- The wrong tape is loaded for restore.

Corrective action

- If the tape is not positioned correctly, position the tape at the location of the backup image and retry the operation.
- If the tape is corrupt, you cannot perform the restore operation.
- If the wrong tape is loaded, retry the operation with the correct tape.

Invalid checksum found for one of the data block, where VBN number is VBN_number

Message Invalid checksum found for one of the data block, where VBN number is *VBN_number*

Cause You might get this error message for one of the following reasons:

- The tape is not positioned at the location of the backup image.
- The tape is corrupt or damaged.
- The wrong tape is loaded for restore.

Corrective action

- If the tape is not positioned correctly, position the tape at the location of the backup image and retry the operation.
- If the tape is corrupt, you cannot perform the restore operation.
- If the wrong tape is loaded, retry the operation with the correct tape.

Block for VBN *VBN_number* failed checksum verification, aborting the current transfer on volume *volume_name*

| | |
|--------------------------|---|
| Message | Block for VBN <i>VBN_number</i> failed checksum verification, aborting the current transfer on volume <i>volume_name</i> |
| Cause | You might get this error message for one of the following reasons: <ul style="list-style-type: none"> • The tape is not positioned at the location of the backup image. • The tape is corrupt or damaged. • The wrong tape is loaded for restore. |
| Corrective action | <ul style="list-style-type: none"> • If the tape is not positioned correctly, position the tape at the location of the backup image and retry the operation. • If the tape is corrupt, you cannot perform the restore operation. • If the wrong tape is loaded, retry the operation with the correct tape. |

Language setting for the Snapshot is not found

| | |
|--------------------------|---|
| Message | Language setting for the Snapshot is not found |
| Cause | Cannot get the language setting from the Snapshot ID. |
| Corrective action | It is a warning message; the backup or restore operation continues. |

Volume is currently under migration

| | |
|--------------------------|---|
| Message | Volume is currently under migration |
| Cause | Volume migration and SMTape backup cannot run simultaneously. |
| Corrective action | Retry the backup job after the volume migration is complete. |

Failed to get latest snapshot

| | |
|--------------------------|---|
| Message | Failed to get latest snapshot |
| Cause | The latest Snapshot copy might not exist because the volume is being initialized by SnapMirror. |
| Corrective action | Retry after initialization is complete. |

Failed to load new tape

| | |
|----------------|-------------------------------|
| Message | Failed to load new tape |
| Cause | Error in tape drive or media. |

Corrective action Replace the tape and retry the operation.

Remote tape not supported

Message Remote tape not supported

Cause A remote tape drive is specified for the backup or restore job.

Corrective action SMTape does not support remote tapes. Use a local tape drive for the job.

Failed to initialize tape

Message Failed to initialize tape

Cause You might get this error message for one of the following reasons:

- The backup image is not of SMTape.
- The tape blocking factor specified is incorrect.
- The tape is corrupt or damaged.
- The wrong tape is loaded for restore.

Corrective action

- If the backup image is not of SMTape, retry the operation with a tape that has SMTape backup.
- If the blocking factor is incorrect, specify the correct blocking factor and retry the operation.
- If the tape is corrupt, you cannot perform the restore operation.
- If the wrong tape is loaded, retry the operation with the correct tape.

Failed to initialize restore stream

Message Failed to initialize restore stream

Cause You might get this error message for one of the following reasons:

- The backup image is not of SMTape.
- The tape blocking factor specified is incorrect.
- The tape is corrupt or damaged.
- The wrong tape is loaded for restore.

Corrective action

- If the backup image is not of SMTape, retry the operation with a tape that has the SMTape backup.
- If the blocking factor is incorrect, specify the correct blocking factor and retry the operation.
- If the tape is corrupt, you cannot perform the restore operation.
- If the wrong tape is loaded, retry the operation with the correct tape.

Failed to read backup image

| | |
|--------------------------|---|
| Message | Failed to read backup image |
| Cause | The tape is corrupt. |
| Corrective action | If the tape is corrupt, you cannot perform the restore operation. |

Invalid backup image magic number

| | |
|--------------------------|---|
| Message | Invalid backup image magic number |
| Cause | The backup image is not of SMTape. |
| Corrective action | If the backup image is not of SMTape, retry the operation with a tape that has the SMTape backup. |

Chunk format not supported

| | |
|--------------------------|---|
| Message | Chunk format not supported |
| Cause | The backup image is not of SMTape. |
| Corrective action | If the backup image is not of SMTape, retry the operation with a tape that has the SMTape backup. |

Invalid backup image checksum

| | |
|--------------------------|---|
| Message | Invalid backup image checksum |
| Cause | The tape is corrupt. |
| Corrective action | If the tape is corrupt, you cannot perform the restore operation. |

Mismatch in backup level number

| | |
|--------------------------|--|
| Message | Mismatch in backup level number |
| Cause | The tape loaded during a tape change is not a part of the backup set. |
| Corrective action | Use the <code>smtape restore -h</code> command to verify the header information of a tape. |

Mismatch in backup time stamp

| | |
|----------------|---|
| Message | Mismatch in backup time stamp |
| Cause | The tape loaded during a tape change is not a part of the backup set. |

Corrective action Use the `smtape restore -h` command to verify the header information of a tape.

Volume read-only

Message Volume read-only
Cause You cannot restore to a read-only volume.
Corrective action Use a read/write volume for SMTape restore.

Invalid source path: /vol/newvol/

Message Invalid source path: `/vol/newvol/`
Cause The specified volume path is invalid.
Corrective action Volume path should not be terminating with the "/" character; use `/vol/newvol`.

Copyright information

Copyright © 1994–2013 NetApp, Inc. All rights reserved. Printed in the U.S.

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. patents, foreign patents, or pending applications.

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

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

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 www.ibm.com/legal/copytrade.shtml.

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

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

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

NetApp, Inc. NetCache is certified RealSystem compatible.

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 email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

- A**
- ACLs (access control lists)
 - excluding from tape restores [120](#)
 - aliases, tape
 - on multiple storage systems [26](#)
 - appending backups to tapes [101](#)
- B**
- backup and restore using NDMP services
 - dump [68](#)
 - SMTape [127](#)
 - backups
 - copying from tape with restore command [110](#)
 - parallel [87](#)
 - backups to tape (dump command)
 - benefits of entering at console [89](#)
 - benefits of using Remote Shell [89](#)
 - estimating tapes required for [63](#)
 - rules for excluding files from [97](#)
 - syntax [87](#)
 - unattended [63](#)
 - where to enter the command [89](#)
 - baseline and incremental restores
 - workflow of [139](#)
 - baseline restores
 - performing [139](#)
 - blocking factor
 - about [62](#)
- C**
- compression type
 - specifying in restores from tape [115](#)
 - considerations
 - before choosing a tape backup method [14](#)
 - before using the dump command [63](#)
- D**
- data backup
 - using SMTape [130](#)
 - data restore
 - using dump engine for [66](#)
 - using SMTape [133](#)
 - dump
 - about [58](#)
 - backing up directories using [58](#)
 - backing up files using [58](#)
 - dump and restore events
 - viewing log messages for [142](#)
 - dump and SMTape backup
 - differences [13](#)
 - dump backup
 - how it works [58](#)
 - dump backup and restore sessions
 - scalability limits for [64](#)
 - dump command
 - backup levels, defined [90](#)
 - deleting a restartable dump [107](#)
 - maximum tape blocks per tape file [100](#)
 - order of tape devices specified by [61](#)
 - specifying a blocking factor [99](#)
 - specifying a dump path [94](#)
 - specifying backup names [99](#)
 - specifying files and directories [95](#)
 - specifying local tape device names [92](#)
 - specifying tape blocks per tape file [100](#)
 - specifying to omit ACLs [98](#)
 - dump engine
 - using environment variables for [68](#)
 - See also* dump
 - dump error messages
 - could not access volume in path: volume_name [162](#)
 - destination qtree is read-only [161](#)
 - destination volume is read-only [161](#)
 - dump context created from NDMP. Cannot restart dump. [160](#)
 - failed to determine snapshot type [159](#)
 - failed to locate the specified restartable dump [160](#)
 - failed to retrieve saved info for the restartable dump. [161](#)
 - IB restore in progress [161](#)
 - invalid inode specified on restart [161](#)
 - invalid restart context. Cannot restart dump. [161](#)
 - invalid/offline volume [159](#)
 - no default tape device list [159](#)
 - no files were created [162](#)
 - unable to locate bitmap files [160](#)
 - unable to locate snapshot [160](#)

- unable to lock a snapshot needed by dump [159](#)
- volume is temporarily in a transitional state [160](#)

dump events

- about [142](#)

dump restart command (restarts interrupted backup) [106](#)

dump restores

- about [65](#)

dumpdates file

- principles applying to dumpdates file [61](#)

- purpose [61](#)

- reasons to update [61](#)

E

enhanced DAR functionality

- disabling [79](#)

- enabling [79](#)

environment variables

- BASE_DATE [127](#)

- DUMP_DATE [127](#)

- FILESYSTEM [127](#)

- SMTAPE_BACKUP_SET_ID [127](#)

- SMTAPE_BREAK_MIRROR [127](#)

- SMTAPE_DELETE_SNAPSHOT [127](#)

- SMTAPE_SNAPSHOT_NAME [127](#)

- uses [57](#)

environment variables for dump

- understanding [68](#)

event logging

- enabling or disabling [148](#)

events during dump and restore operations

- logging [141](#)

events in takeover mode

- log files [141](#)

examples

- event log

 - SMTape backup [146](#)

 - SMTape restore [147](#)

examples of ndmcopy command

- migrating data from a source path on a remote host to a destination path on another remote host [83](#)

- migrating data from a source path on remote host to a destination path on the local storage system [83](#)

- migrating data from a source path to a different destination path on a remote host [83](#)

- migrating data from a source path to a different destination path on the same storage system [83](#)

- overwriting the /etc directory during the root volume migration [83](#)

existing aliases of tape drives and medium changers

- determining [23](#)

F

files

- excluding data from backup [97](#)

- excluding from dump command [97](#)

files and data backup

- using dump engine for [59](#)

format of dump and restore event log messages

- about [141](#)

format of SMTape event log messages

- about [145](#)

I

ignore inode limitations

- during restore operation [118](#)

image header of a tape

- displaying [132](#)

increment chains

- understanding [60](#)

incremental restores

- performing [140](#)

information about tape drives

- information about medium changers

 - viewing

- viewing

information about tape drives and medium changers

- displaying [24](#)

initiate a tape backup

- how to [13](#)

inode limitations

- ignoring [118](#)

L

latest NDMP version

- displaying [51](#)

levels of backups [90](#)

levels of incremental backup

- specifying [60](#)

M

manage NDMP

- how to [45](#)

multipath tape access

- about [26](#)

considerations for [26](#)
 understanding [26](#)

N

NDMP

about [41](#)
 advantages of [41](#)
 considerations [54](#)
 disabling preferred network interface [48](#)
 displaying file history performance [85](#)
 enabling or disabling service (ndmpd on/off) [45](#)
 firewall policy [54](#)
 killing sessions (ndmpd kill command) [51](#)
 ndmpcopy command [54](#)
 options [51](#)
 range of ports for data connections, designating [47](#)
 session information
 displaying detailed status (ndmpd probe) [48](#)
 displaying status (ndmpd status command) [48](#)
 setting preferred network interface [46](#)
 tape backup topologies
 Storage system-to-data server-to-tape [53](#)
 Storage system-to-local-tape [53](#)
 Storage system-to-network attached tape library [53](#)
 Storage system-to-tape attached to another storage system [53](#)
 tape devices used with [55](#)
 using with tape libraries [55](#)

NDMP backup application management
 preparing a storage system for [56](#)

NDMP commands

ndmp on [45](#)
 ndmpcopy (uses local copy tool) [80](#)
 ndmpd kill (terminates NDMP session) [51](#)
 ndmpd on/off (enabling or disabling service) [45](#)
 ndmpd probe (displays detailed status) [48](#)
 ndmpd status (displays status) [48](#)

NDMP error messages

cannot generate NDMP password [155](#)
 could not obtain vol ref for Volume *volume_name* [156](#)
 error: Unable to retrieve session information [155](#)
 message from Read Socket: error_string [154](#)
 message from Write Dirnet: error_string [154](#)
 ndmpd invalid version number: *version_number* [155](#)
 ndmpd session *session_ID* not active. [155](#)
 network communication error [154](#)

no such user *user_name* [155](#)
 read Socket received EOF [154](#)
 the specified operation could not be completed as the volume is moving [155](#)

ndmpcopy command

about [80](#)
 examples [83](#)

ndmpcopy error messages

CONNECT: Connection refused [158](#)
 error getting local hostname [158](#)
 invalid name. Destination filer name does not resolve to the specified address mode [159](#)
 invalid name. Source filer name does not resolve to the specified address mode [158](#)
 ndmpcopy: Authentication failed for destination [157](#)
 ndmpcopy: Authentication failed for source [157](#)
 ndmpcopy: Client authentication request failed [156](#)
 ndmpcopy: Connection setup for transfer failed [158](#)
 ndmpcopy: Error in getting extension list [157](#)
 ndmpcopy: Error opening NDMP connection [156](#)
 ndmpcopy: Failed to start dump on source [157](#)
 ndmpcopy: Failed to start restore on destination [157](#)
 ndmpcopy: Socket connection to *host_name* failed [156](#)

nonqualified tape drives
 using [37, 38](#)

O

options

backup.log.enable (turns event logging on or off) [148](#)
 concurrent operations
 volume SnapMirror and SMTape backup enabling or disabling [138](#)
 ndmp.preferred_interface (sets preferred network) [46](#)
 ndmpd.data_port_range [47](#)
 vsm.smtape.concurrent.cascade.support [138](#)

P

physical path names (PPNs)
 about [22](#)

Q

qtrees

excluding data from backup [96](#)
 omitting data from dump command [96](#)

- omitting information from tape restores [121](#)
- qualified tape drives
 - about [34](#)
 - emulating [38](#)

R

- restartable backups
 - deleting automatically [107](#)
 - qualifications [64](#)
- restore
 - incremental backups [110](#)
- restore command
 - disk space required for [67](#)
 - information required for using [67](#)
 - options [108](#)
 - restoring individual files [111](#)
 - specifying a full restore [112](#)
 - specifying a resume restore [114](#)
 - specifying a single tape file on a multifile tape [115](#)
 - specifying a test restore [121](#)
 - specifying automatic confirmations [119](#)
 - specifying no qtree information [121](#)
 - specifying table-of-contents restore [113](#)
 - specifying tape devices [115](#)
 - specifying the blocking factor [117](#)
 - specifying to exclude ACLs [120](#)
 - syntax [107](#)
 - types of restores [108](#)
 - using with Remote Shell [109](#)
- restore command, executing [110](#)
- restore events
 - about [144](#)
- restores
 - performing baseline and incremental [139](#)
- rules
 - for restore command [107](#)
 - for specifying a resume restore [114](#)

S

- serial numbers
 - about [23](#)
- simultaneous backup or restore sessions
 - supported number of [18](#)
- SMTape
 - aborting a backup or restore job [134](#)
 - about [125](#)
 - backing up data using [130](#)
 - backing up files using [125](#)

- backup and restore using CLI commands [130](#)
- continuing a backup or restore [134](#)
- features not supported [127](#)
- features of [126](#)
- performing restores using [139](#)
- removing the snapmirror status entry
 - after a backup operation [136](#)
 - after a restore operation [137](#)
- restoring data [133](#)
- SMTape backup
 - how it works [125](#)
 - using Snapshot copies [130](#)
- SMTape backup and restore operations
 - displaying status of [135](#)
- smtape commands
 - smtape abort [134](#)
 - smtape backup [130](#)
 - smtape continue [134](#)
 - smtape restore [132](#), [133](#)
 - smtape status [135](#)
- SMTape engine
 - See* SMTape
- SMTape error messages
 - aborting: Destination volume, *volume_name*, is a clone [167](#)
 - aborting: Destination volume, *volume_name*, is too small [167](#)
 - already read *volume_name* tape_number [166](#)
 - bad block in read stream. VBN = *VCN_number*, max_VBN = *max_VCN_number* [171](#)
 - block for VBN *VCN_number* failed checksum verification, aborting the current transfer on volume *volume_name* [172](#)
 - cannot init input, aborting [168](#)
 - chunk format not supported [174](#)
 - chunks out of order [166](#)
 - destination is not an aggregate. Aborting [169](#)
 - duplicate VBN *VCN_number* received for volume *volume_name*, aborting transfer [170](#)
 - failed to access the named snapshot [165](#)
 - failed to allocate memory [164](#)
 - failed to create job UUID [164](#)
 - failed to create snapshot [165](#)
 - failed to delete snapshot [166](#)
 - failed to delete softlock [166](#)
 - failed to find snapshot [165](#)
 - failed to get data buffer [164](#)
 - failed to get latest snapshot [172](#)
 - failed to initialize restore stream [173](#)
 - failed to initialize tape [173](#)

- failed to load new tape [172](#)
 - failed to lock snapshot [165](#)
 - failed to read backup image [174](#)
 - failed to softlock qtree snapshots [165](#)
 - image header missing or corrupted [166](#)
 - incompatible SnapMirror or copy source Version. Aborting [168](#)
 - internal assertion [162](#)
 - invalid backup image checksum [174](#)
 - invalid backup image magic number [174](#)
 - invalid checksum for the chunk descriptor [170](#)
 - invalid checksum found for one of the data block, where VBN number is *VBN_number* [171](#)
 - invalid contents in destination volume geometry string *volume_geometry_string*, aborting [168](#)
 - invalid input tape [164](#)
 - invalid source path: */vol/newvol/* [175](#)
 - invalid volume path [163](#)
 - job aborted due to shutdown [162](#)
 - job aborted due to Snapshot autodelete [163](#)
 - job not found [162](#)
 - language setting for the Snapshot is not found [172](#)
 - mismatch in backup level number [174](#)
 - mismatch in backup set ID [167](#)
 - mismatch in backup time stamp [174](#)
 - received VBN header with invalid checksum *error_string*, aborting transfer on volume *volume_name* [170](#)
 - remote tape not supported [173](#)
 - source is not a hybrid aggregate. Aborting [169](#)
 - source is not an aggregate. Aborting [169](#)
 - source volume is a flexible volume. Aborting [169](#)
 - source volume is not a flexible volume. Aborting [169](#)
 - source volume size is greater than maximum supported SIS volume size on this platform. Aborting [167](#)
 - tape is currently in use by other operations [164](#)
 - tapes out of order [166](#)
 - too many active jobs [164](#)
 - too many active transfers at once, aborting [168](#)
 - transfers from volume *volume_name* are temporarily disabled [168](#)
 - UNIX style RMT tape drive is not supported [163](#)
 - volume is currently in use by other operations [163](#)
 - volume is currently under migration [172](#)
 - volume not restricted [163](#)
 - volume offline [163](#)
 - volume read-only [175](#)
 - SMTape events
 - backup events [146](#)
 - CLI backup and restore [146](#)
 - restore events [147](#)
 - SnapMirror status entries
 - removing [136](#)
 - removing entries
 - after a backup [136](#)
 - after a restore [137](#)
 - storage (aliasing) commands
 - storage unalias (removes tape alias) [25](#)
 - storage systems
 - adding Fiber Channel-attached drives dynamically to [27](#)
 - displaying information about tape drive connections to [28](#)
 - dynamically adding tape drives and libraries to [27](#)
 - subtrees, defined [94](#)
 - supported NDMP extensions
 - about [53](#)
 - sysconfig -m command (shows information about tape medium changers) [28](#)
 - sysconfig -v command (shows tape drive connections to storage system) [28](#)
- ## T
- table-of-contents for tape restores
 - about [112](#)
 - table-of-contents restore display
 - using remote shell connection for [112](#)
 - tape aliases
 - assigning [25](#)
 - definition [21](#)
 - using serial numbers for [23](#)
 - tape backup
 - using NDMP [41](#)
 - tape backup and recovery
 - using NDMP [41](#)
 - using the dump engine [58](#)
 - using the SMTape engine [125](#)
 - tape backup and restore error messages
 - already at the end of tape [153](#)
 - could not initialize media [151](#)
 - duplicated tape drive (tape_drive) specified in the tape argument list [150](#)
 - invalid tape drive tape_drive in tape argument list [150](#)
 - media error on tape read [152](#)
 - media error on tape write [151](#)
 - resource limitation: no available thread [150](#)

- tape read error [153](#)
 - tape record size is too small [153](#)
 - tape record size must be in the range between 4KB and 256KB [153](#)
 - tape record size should be block_size1 and not block_size2 [153](#)
 - tape reservation preempted [151](#)
 - tape write error [152](#)
 - tape write failed [152](#)
 - tape write failed - new tape encountered media error [152](#)
 - tape write failed - new tape is already at the end of media [152](#)
 - tape write failed - new tape is broken or write protected [152](#)
 - too many active dumps/restores currently in progress [151](#)
 - tape backup and restore event log files
 - accessing [141](#)
 - tape backup and restore operations
 - accessing the event log files for [141](#)
 - tape configuration files
 - accessing [34](#)
 - format of [34](#)
 - tape device name
 - format [17](#)
 - tape devices
 - about [16](#)
 - types of [16](#)
 - tape drive connections
 - supported number of [18](#)
 - tape drives
 - controlling [29](#)
 - managing [16](#)
 - nonqualified
 - using [37](#)
 - rewinding tape (mt -rewind) [32](#)
 - showing status (mt -status) [33](#)
 - tape medium changers, displaying information about [28](#)
 - unloading tape after rewind (mt -offline) [32](#)
 - tape drives and libraries to storage systems
 - dynamically adding [27](#)
 - tape drives dynamically
 - qualifying [36](#)
 - tape drives to storage systems
 - dynamically adding [27](#)
 - tape libraries
 - showing names assigned to [55](#)
 - tape libraries to storage systems
 - dynamically adding [27](#)
 - tape reservations
 - what are [39](#)
 - tape restores
 - displaying a table of contents for [112](#)
 - displaying detailed status output [118](#)
 - running a test restore [121](#)
 - specifying a restore destination [116](#)
 - specifying automatic confirmations [119](#)
 - specifying tape devices [115](#)
 - tape seeding [126](#)
 - tapes
 - displaying the image header of [132](#)
 - traditional volumes
 - displaying the volume geometry of [131](#)
 - types of tape backup [12](#)
- ## V
- volume move operations
 - performing tape operations during [15](#)