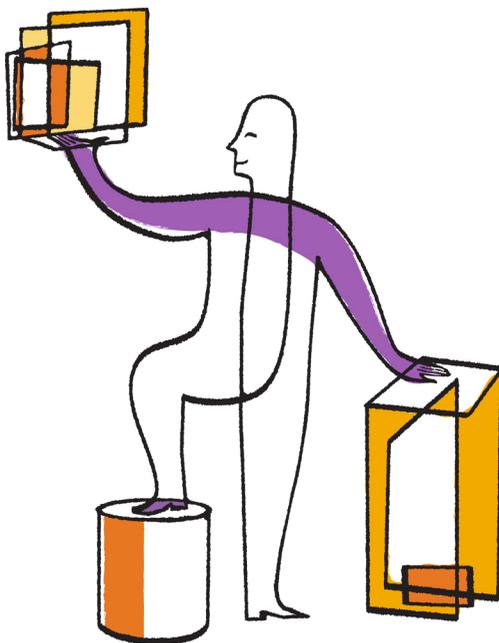




Clustered Data ONTAP[®] 8.2

Data Protection Tape Backup and Recovery Guide



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: 215-07965_A0
May 2013

Contents

Tape backup of FlexVol volumes and Infinite Volumes	8
Tape backup and restore workflow for FlexVol volumes	8
Tape backup and restore workflow for Infinite Volumes	9
Understanding tape drives	11
What qualified tape drives are	11
Format of the tape configuration file	12
How the storage system qualifies a new tape drive dynamically	14
What tape devices are	14
Tape device name format	15
Supported number of simultaneous tape devices	16
What tape aliasing is	17
What physical path names are	17
What serial numbers are	18
Considerations when configuring multipath tape access	19
How to add tape drives and libraries to storage systems	19
What tape reservations are	20
Managing tape drives	21
Commands for viewing tape drive information	21
Using a nonqualified tape drive	22
Assigning tape aliases	23
Removing tape aliases	24
Enabling or disabling tape reservations	24
Commands for verifying tape library connections	25
Understanding NDMP for FlexVol volumes	27
About NDMP modes of operation	27
Considerations when using NDMP	28
What environment variables do	29
Environment variables supported by Data ONTAP	29
Common NDMP tape backup topologies	39
Supported NDMP authentication methods	40
What enhanced DAR functionality is	40
Scalability limits for NDMP sessions	41

Managing node-scoped NDMP mode for FlexVol volumes	42
Commands for managing node-scoped NDMP mode	42
Node-scoped NDMP mode options	43
User authentication in a node-scoped NDMP mode	46
Managing Vserver aware NDMP mode for FlexVol volumes	47
Commands for managing Vserver aware NDMP mode	48
Vserver aware NDMP mode options	49
NDMP extensions supported by Data ONTAP	53
What Cluster Aware Backup extension does	53
Availability of volumes and tape devices for backup and restore on different LIF types	53
What affinity information is	54
What preferred-interface-role is	55
NDMP data connection types	56
User authentication in a Vserver aware NDMP mode	57
Generating an NDMP-specific password for NDMP users	57
Understanding dump engine for FlexVol volumes	59
How a dump backup works	59
What the dump engine backs up	60
What increment chains are	61
What the blocking factor is	62
How a dump restore works	62
What the dump engine restores	63
Considerations before restoring data	64
Scalability limits for dump backup and restore sessions	65
Tape backup and restore between Data ONTAP 7-Mode and clustered Data ONTAP	65
How dump backs up data from a vault destination volume	66
How dump works with SFO and ARL	66
How dump works with volume movement operation	67
Monitoring tape backup and restore operations for FlexVol volumes	68
Accessing the event log files	68
What the dump and restore event log message format is	69
What logging events are	69
What dump events are	69

What restore events are	70
Enabling or disabling event logging	71
Error messages for tape backup and restore of FlexVol volumes	72
Backup and restore error messages	72
Resource limitation: no available thread	72
Tape reservation preempted	72
Could not initialize media	72
Too many active dumps/restores currently in progress	73
Media error on tape write	73
Tape write failed	73
Tape write failed - new tape encountered media error	73
Tape write failed - new tape is broken or write protected	73
Tape write failed - new tape is already at the end of media	73
Tape write error	74
Media error on tape read	74
Tape read error	74
Already at the end of tape	74
Tape record size is too small. Try a larger size.	74
Tape record size should be block_size1 and not block_size2	75
Tape record size must be in the range between 4KB and 256KB	75
NDMP error messages	75
Network communication error	75
Message from Read Socket: error_string	75
Message from Write Dirnet: error_string	75
Read Socket received EOF	76
ndmpd invalid version number: version_number	76
ndmpd session session_ID not active	76
Could not obtain vol ref for Volume volume_name	76
Data connection type	
["NDMP4_ADDR_TCP" "NDMP4_ADDR_TCP_IPv6"] not	
supported for ["IPv6" "IPv4"] control connections	76
DATA LISTEN: CAB_DATA_CONN_PREPARE not sent or addr type	
mismatch	77
DATA CONNECT: CAB_DATA_CONN_PREPARE not sent or addr	
type mismatch	77
Error:show failed: Cannot get password for user '<username>'	77

Dump error messages	78
Destination volume is read-only	78
Destination qtree is read-only	78
Dumps temporarily disabled on volume, try again	78
No files were created	78
Unable to lock a snapshot needed by dump	78
Unable to locate bitmap files	79
Volume is temporarily in a transitional state	79
Tape backup and restore of Infinite Volumes	80
Protocols for tape backup and restore of Infinite Volumes	80
Tape backup topology to use for Infinite Volumes	81
How storage classes affect tape backup of Infinite Volumes	81
Understanding incremental tape backup of Infinite Volumes using SnapDiff and Snapshot copies	81
What SnapDiff is	81
SnapDiff support for backup applications	82
Guidelines for using LIFs with SnapDiff	82
What Snapshot copies you should back up to tape	82
What a namespace mirror constituent is	83
When namespace mirror constituents are created	83
How modifying an Infinite Volume affects namespace mirror constituents	84
How namespace mirror constituents affect the state of an Infinite Volume	85
How aggregates are selected for namespace mirror constituents	86
Space and node requirements for namespace mirror constituents	87
How the file system size setting for an Infinite Volume affects SnapDiff ...	88
Setting up incremental tape backup of Infinite Volumes	88
Verifying space requirements in an Infinite Volume for SnapDiff resources (cluster administrators only)	89
Enabling SnapDiff on an Infinite Volume	91
Selecting a Snapshot copy of an Infinite Volume for SnapDiff to use	93
Creating a baseline for incremental tape backup	93
Setting up incremental tape backup of destination Infinite Volumes	95
Transferring a Snapshot copy to the destination Infinite Volume	95
Creating a baseline for incremental tape backup	96

Incrementally backing up changed files from an Infinite Volume to tape	97
Restoring files from tape for an Infinite Volume	98
Reenabling SnapDiff after expanding a destination Infinite Volume onto new nodes	98
Disabling SnapDiff for an Infinite Volume	100
Copyright information	102
Trademark information	103
How to send your comments	104
Index	105

Tape backup of FlexVol volumes and Infinite Volumes

For FlexVol volumes, Data ONTAP supports tape backup and restore through the Network Data Management Protocol (NDMP). For Infinite Volumes, Data ONTAP supports tape backup and restore through a mounted volume. Infinite Volumes do not support NDMP. The type of volume determines what method to use for backup and recovery.

Tape backup of FlexVol volumes using NDMP

NDMP allows you to back up storage systems directly to tape, resulting in efficient use of network bandwidth. Clustered Data ONTAP supports dump engine for tape backup. 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 access control list (ACL) information to tape. You can back up an entire volume, an entire qtree, or a subtree that is neither an entire volume nor an entire qtree. Dump supports level-0, differential, and incremental backups.

You can perform a dump backup or restore by using NDMP-compliant backup applications. Starting with Data ONTAP 8.2, only NDMP version 4 is supported.

Tape backup of Infinite Volumes using a mounted volume

You can back up and restore Infinite Volumes using any data management application that can back up files over a volume mounted with the NFS or CIFS protocols and that supports SnapDiff. However, you cannot back up or restore Infinite Volumes with NDMP.

Related concepts

[Understanding NDMP for FlexVol volumes](#) on page 27

[Understanding dump engine for FlexVol volumes](#) on page 59

[Tape backup and restore of Infinite Volumes](#) on page 80

Tape backup and restore workflow for FlexVol volumes

You can perform tape backup and restore operations by using an NDMP-enabled backup application. The workflow provides an overview of tasks that are involved in performing tape backup and restore operations.

The high-level tasks that are required to perform a tape backup and restore operation are as follows:

1. Set up a tape library configuration choosing an NDMP-supported tape topology.
2. Enable NDMP services on your storage system.

You can enable the NDMP services either at a node level or at a Vserver level. This depends upon the NDMP mode in which you choose to perform a tape backup and restore operation.

3. Use NDMP options to manage NDMP on your storage system.

You can use NDMP options either at a node level or at a Vserver level. This depends upon the NDMP mode in which you choose to perform a tape backup and restore operation.

4. Perform a tape backup or restore operation by using an NDMP-enabled backup application. Clustered Data ONTAP supports dump engine for tape backup and restore.

For more information about using the backup application (also called Data Management Applications or DMAs) to perform backup or restore operations, see your backup application documentation.

Related concepts

[Understanding dump engine for FlexVol volumes](#) on page 59

Related references

[Common NDMP tape backup topologies](#) on page 39

[Node-scoped NDMP mode options](#) on page 43

[Vserver aware NDMP mode options](#) on page 49

Tape backup and restore workflow for Infinite Volumes

You can perform incremental tape backup and restore operations for an Infinite Volume mounted with the NFS or CIFS protocol by using a SnapDiff-supported backup application. The workflow provides an overview of tasks required for incremental tape backup and restore operations.

Note: Infinite Volumes do not support NDMP.

You can incrementally back up the data in Infinite Volumes to tape and restore the data from tape using any backup application (also called Data Management Applications or DMAs) that meets the following requirements:

- Supports SnapDiff
- Backs up and restores data over a volume mounted with the NFS or CIFS protocol

Note: The SnapDiff API must be configured for the Infinite Volume.

The high-level tasks that are required to set up and use incremental tape backup for Infinite Volumes are as follows:

1. Set up a tape library configuration to meet the requirements specified by the backup application for a volume mounted with the NFS or CIFS protocol.
For more information about tape library requirements for the backup application, see your backup application documentation.
2. Configure the backup application to back up the mounted Infinite Volume.

3. Set up incremental tape backup of Infinite Volumes:
 - a. Verify space requirements in the Infinite Volume for SnapDiff resources.
 - b. Enable SnapDiff on the Infinite Volume.
 - c. Select a Snapshot copy of the Infinite Volume for SnapDiff to use.
 - d. Configure SnapDiff to use the Snapshot copy.
 - e. Back up all data in the Infinite Volume to create a baseline for SnapDiff to use in the future. For more information about using the backup application to perform backup operations, see your backup application documentation.
4. Back up only changed files in the Infinite Volume for each subsequent tape backup:
 - a. Use a backup application to request from SnapDiff a list of changed files in the Infinite Volume since the last tape backup.
 - b. Use a backup application to back up the changed files to tape.
5. Restore files from tape by using the backup application.
For more information about using the backup application to perform restore operations, see your backup application documentation.

Related concepts

[Understanding incremental tape backup of Infinite Volumes using SnapDiff and Snapshot copies](#)
on page 81

[Tape backup topology to use for Infinite Volumes](#) on page 81

Understanding tape drives

You must use a qualified tape drive that has been tested and found to work properly on a storage system. You can follow tape aliasing and also enable tape reservations to ensure that only one storage system accesses a tape drive at any particular time.

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. You can download the tape configuration file from 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 tasks

[Using a nonqualified tape drive](#) on page 22

Related references

[Commands for viewing tape drive information](#) on page 21

[Commands for verifying tape library connections](#) on page 25

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

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

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

Item	Size	Description
{l m h a}_description=(string)	up to 16 bytes	The string to print for the nodeshell command, <code>sysconfig -t</code> that describes characteristics of the particular density setting.
{l m h a}_density=(hex codes)		The density code to be set in the SCSI mode page block descriptor corresponding to the desired density code for l, m, h, or a.

Item	Size	Description
{l m h a}_algorithm=(hex codes)		The compression algorithm to be set in the SCSI Compression Mode Page corresponding to the density code and the desired density characteristic.

The following table describes the optional fields available in the tape configuration file:

Field	Description
autoload=(Boolean yes/no)	This field is set to <i>yes</i> 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 <i>no</i> .
cmd_timeout_0x	Individual timeout value. Use this field only if you want to specify a different timeout value from the one being used as a default by the tape driver. The sample file lists the default SCSI command timeout values used by the tape 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.

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.

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.

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.

Related concepts

Scalability limits for dump backup and restore sessions on page 65

What tape aliasing 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.

Note: When you upgrade from Data ONTAP 8.1.x to Data ONTAP 8.2.x, the tape alias feature of Data ONTAP 8.2.x modifies the existing tape alias names. In such a case you might have to update the tape alias names in the backup application.

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.

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.

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.

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[LX]

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.

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.

In the following example, the storage system accesses LUN group 0 through two initiator paths: 0b and 0d. On both of these paths, the LUN group has the same LUN number, 0 and LUN count, 5. The storage system accesses LUN group 1 through only one initiator path, 3d.

```
STSW-3070-2_cluster:> storage library config show
```

Node	LUN Group	LUN Count	Library Name	Library Target Port	Initiator
STSW-3070-2_cluster-01	0	5	IBM 3573-TL_1	510a09800000412d	0b
					0d
	1	2	IBM 3573-TL_2	50050763124b4d6f	3d

3 entries were displayed

For more information, see the man pages.

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.

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.

Managing tape drives

You can verify tape library connections and view tape drive information prior to performing a tape backup or restore operation. You can use a nonqualified tape drive by emulating this to a qualified tape drive. You can also assign and remove tape aliases in addition to viewing existing aliases.

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.

Commands for viewing tape drive information

You can view information about tape drives to help you perform a tape backup and restore operation, use tape drives supported by Data ONTAP, understand tape drive performance, and be aware of the existing tape aliases.

You can view the following tape drive information:

- Node to which the tape drive is attached
- Device ID
- NDMP path
- Tape drive description
- Tape drives supported by Data ONTAP
- Tape drive statistics
- Existing aliases of tape drives

If you want to...	Use this command...
View information about tape drives in a cluster	<code>system node hardware tape drive show</code>
View tape drives supported by Data ONTAP	<code>storage show tape supported [-v]</code> Note: You must use this command at the nodeshell. You can use the <code>-v</code> option to view more details about each tape drive.
View tape device statistics to understand tape performance and check usage pattern. You can reset the statistics reading and restart the process of displaying the statistics whenever you want.	<code>storage stats tape tape_name</code> Note: You must use this command at the nodeshell.

If you want to...	Use this command...
View existing aliases of tape drives	<pre>storage alias</pre> <p>Note: You must use this command at the nodeshell.</p>

For more information about these commands, refer to the man pages.

Using a nonqualified tape drive

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 `mt` command, go directly to Step 3. If the storage system has not accessed the tape drive through the `mt` command, go to Step 2.
2. To access the tape drive, enter the following command at the nodeshell:

```
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 at the nodeshell:

```
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)
```

To emulate a qualified tape drive, go to the [NetApp Support Site](#).

Related concepts

[What qualified tape drives are](#) on page 11

Related information

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

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 at the nodeshell:

```
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.

Note: You can view the PPN and SN information about the tape drives and tape libraries by using `system node hardware tape drive show` and `system node hardware tape library show` commands respectively.

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.

Related concepts

[What tape aliasing is](#) on page 17

Related tasks

[Removing tape aliases](#) on page 24

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 at the nodeshell:

```
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
```

Related concepts

[What tape aliasing is](#) on page 17

Related tasks

[Assigning tape aliases](#) on page 23

Enabling or disabling tape reservations

You can control how Data ONTAP manages tape device reservations by using the `tape.reservations` option. By default, tape reservation is turned off.

About this task

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 use either the SCSI Reserve/Release mechanism or SCSI Persistent Reservations or to disable tape reservations, enter the following command at the clustershell:

```
options -option-name tape.reservations -option-value {scsi | persistent
| off}
```

`scsi` selects the SCSI Reserve/Release mechanism.

`persistent` selects SCSI Persistent Reservations.

`off` disables tape reservations.

Related concepts

[What tape reservations are](#) on page 20

Commands for verifying tape library connections

You can view information about the path connected between a storage system and a tape library configuration attached to the storage system. You can use this information to verify the connection path to the tape library configuration or for troubleshooting issues related to the connection paths.

You can view the following tape library details to verify the tape library connections after adding or creating a new tape library or after restoring a failed path in a single-path or multipath access to a tape library. You can also use this information during troubleshooting of path-related errors or if the access to a tape library fails.

- Node to which the tape library is attached
- Device ID
- NDMP path
- Tape library name
- Target port and initiator port IDs
- Single-path or multipath access to a tape library for every target or FC initiator port
- Path-related data integrity details such as "Path Errors" and "Path Qual"
- LUN groups and LUN counts

If you want to...	Use this command...
Display information about a tape library in a cluster	<code>system node hardware tape library show</code>
Display information about tape library paths for every target port	<code>storage library path show</code> Note: If you want to display information related to "Path Errors" and "Path Qual", you must use the <code>-detail</code> option along with this command.

If you want to...	Use this command...
Display information about tape library paths for every initiator port	<code>storage library path show-by-initiator</code>
Display information about LUN groups and LUN counts in a tape library	<code>storage library config show</code>

For more information about these commands, refer to the man pages.

Understanding NDMP for FlexVol volumes

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.

Starting with Data ONTAP 8.2, you can perform tape backup and restore operations in either node-scoped NDMP mode or Vserver aware NDMP mode.

You must be aware of the considerations that you need to take into account while using NDMP, list of environment variables, and supported NDMP tape backup topologies. You can also enable or disable the enhanced DAR functionality. The two authentication methods supported by Data ONTAP for authenticating NDMP access to a storage system are: plaintext and challenge.

NDMP does not support backup and restore of Infinite Volumes.

About NDMP modes of operation

Starting with Data ONTAP 8.2, you can choose to perform tape backup and restore operations either at a node level as you have been doing until now or at a Vserver level. To perform these operations successfully at the Vserver level, NDMP service must be enabled on the Vserver.

In a mixed cluster where nodes are running different versions of Data ONTAP 8.2 and earlier versions in the Data ONTAP 8.x release family, NDMP follows the node-scoped behavior (at the node level). This NDMP behavior continues even after upgrading to Data ONTAP 8.2. You must explicitly disable the node-scoped NDMP mode to perform tape backup and restore operations in the Vserver aware mode.

In a newly installed cluster where all nodes are running Data ONTAP 8.2, NDMP is in the Vserver aware mode (at the Vserver level) by default. To perform tape backup and restore operations in the node-scoped NDMP mode, you must explicitly enable the node-scoped NDMP mode.

Related concepts

[Managing node-scoped NDMP mode for FlexVol volumes](#) on page 42

[Managing Vserver aware NDMP mode for FlexVol volumes](#) on page 47

Related references

Commands for managing node-scoped NDMP mode on page 42

Considerations when using NDMP

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

- NDMP services can generate file history data at the request of NDMP backup applications. File history is used by backup applications to enable optimized recovery of selected subsets of data from a backup image. File history generation and processing might be time-consuming and CPU-intensive for both the storage system and the backup application. If your data protection needs are limited to disaster recovery, where the entire backup image will be recovered, you can disable file history generation to reduce backup time. See your backup application documentation to determine if it is possible to disable NDMP file history generation.
- Firewall policy for NDMP is enabled by default on all LIF types. For information about managing firewall service and policies, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.
- In the node-scoped NDMP mode, to back up a FlexVol volume you must use the backup application to initiate a backup on a node that owns the volume. However, you cannot back up a node root volume.
- You can perform NDMP backup from any LIF as permitted by the firewall policies. If you use a data LIF, you must select one that is not configured for failover. If a data LIF fails over during an NDMP operation, the NDMP operation fails and must be reexecuted.
- In the node-scoped NDMP mode, NDMP data connection uses the same LIF as the NDMP control connection.
- NDMP backup path is of the format `/vserver_name/volume_name/path_name`, where `path_name` is the path of the directory, file, or Snapshot copy.
- When using `ndmpcopy` command for transferring data between a storage system running Data ONTAP operating in 7-Mode and a storage system running clustered Data ONTAP:
 - The `ndmpcopy` command must be initiated from a storage system running Data ONTAP operating in 7-Mode
 - In the node-scoped mode, the destination IP address is the address of a LIF on the node on which the target volume is located
 - Destination path is of the format `/vserver_name/volume_name`

Note: You should not use the `ndmpcopy` command for restoring a LUN between a storage system running Data ONTAP operating in 7-Mode and a storage system running clustered Data ONTAP because the LUN is restored as a file on the destination storage system.

For the syntax and examples of the `ndmpcopy` command, see the *Data ONTAP Data Protection Tape Backup and Recovery Guide for 7-Mode*.

- 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.

Related concepts

[Tape backup and restore between Data ONTAP 7-Mode and clustered Data ONTAP](#) on page 65

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 `/vserver1/voll/dir1`, the backup application sets the `FILESYS` environment variable to `/vserver1/voll/dir1`. 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.

Environment variables supported by Data ONTAP

Environment variables are used to communicate information about a backup or restore operation between an NDMP-enabled backup application and a storage system. Data ONTAP supports environment variables, 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 backup or restore operations might not be doing what the backup application expected them to do. But in some cases, judicious modification 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. 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 LEVEL/UPDATE based incremental backups.
DEBUG	Y or N	N	Specifies that debugging information is printed.

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. This variable is mandatory for multiple subtree backups.
DUMP_DATE	<i>return_value</i>	none	You do not change this variable directly. It is created by the backup if the BASE_DATE variable is set to a value other than -1. The DUMP_DATE variable is derived by prepending the 32-bit level value to a 32-bit time value computed by the dump software. The level is incremented from the last level value passed into the BASE_DATE variable. The resulting value is used as the BASE_DATE value on a subsequent incremental backup.

Environment variable	Valid values	Default	Description
ENHANCED_DAR_ENABLED	Y or N	N	<p>Specifies if enhanced DAR functionality is instantiated. Enhanced DAR functionality supports directory DAR, and DAR of files with NT Streams. It provides performance improvements. Enhanced DAR during restore is possible only if the following conditions are met:</p> <ul style="list-style-type: none"> • Data ONTAP 8.0 or later supports enhanced DAR • 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 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. <p>Note: Files or directories specified to be excluded for backup are not excluded if you set the NON_QUOTA_TREE to Y simultaneously.</p>
EXTRACT	Y, N, or E	N	<p>Specifies that subtrees of a backed-up data set are to be restored. The backup application specifies the names of the subtrees to be extracted. If a file 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 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.
FORCE	Y or N	N	<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>
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>

Environment variable	Valid values	Default	Description
IGNORE_CTIME	Y or N	N	Specifies that a file is not incrementally backed up if only its ctime value has changed since the previous incremental backup. Some applications, such as virus scanning software, change the ctime value of a file within the inode, even though the file or its attributes have not changed. As a result, an incremental backup might back up files which have not changed. The IGNORE_CTIME variable should be specified only if incremental backups are taking an unacceptable amount of time or space because the ctime value was modified.
IGNORE_QTREES	Y or N	N	Specifies that the restore operation does not restore qtree information from backed up qtrees.
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.
LIST_QTREES	Y or N	N	Lists the backed-up qtrees without actually restoring the data.

Environment variable	Valid values	Default	Description
MULTI_SUBTREE_NAMES	<i>string</i>	none	Specifies that the backup is a multiple subtree backup. Multiple subtrees are specified in the string which is a newline-separated, null-terminated list of subtree names. Subtrees are specified by path names relative to their common root directory, which must be specified as the last element of the list. If you use this variable, you must also use the DMP_NAME variable.
NDMP_UNICODE_FH	Y or N	N	Specifies that a Unicode name is included in addition to the NFS name of the file in the file history information. This option is not used by most backup applications and should not be set unless the backup application is designed to receive these additional file names. The HIST variable must also be set.
NDMP_VERSION	<i>return_only</i>	none	You should not modify the NDMP_VERSION variable. Created by the backup operation, the NDMP_VERSION variable returns the NDMP version. Data ONTAP sets the NDMP_VERSION variable during a backup for internal use and to pass to a backup application for informational purposes. The NDMP version of an NDMP session is not set with this variable.
NO_ACLS	Y or N	N	Specifies that ACLs not be copied when backing up data.

Environment variable	Valid values	Default	Description
NON_QUOTA_TREE	Y or N	N	<p>Specifies that files and directories in qtrees be ignored when backing up data. 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.</p> <p>Note: Files or directories specified to be excluded for backup are not excluded if you set the NON_QUOTA_TREE to Y simultaneously.</p>
NOWRITE	Y or N	N	<p>Specifies that the restore operation not write data to the disk. This variable is used for debugging.</p>
PATHNAME_SEPARATOR	<i>return_value</i>	none	<p>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.</p>

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	dump	Specifies the type of backup supported to perform tape backup and restore operations.
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.

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 or TCP/IPv6 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.

Supported NDMP authentication methods

You can specify an authentication method to allow NDMP connection requests. Data ONTAP supports two methods for authenticating NDMP access to a storage system: plaintext and challenge.

In node-scoped NDMP mode, both challenge and plaintext are enabled by default. However, you cannot disable challenge. You can enable and disable plaintext. In the plaintext authentication method, the login password is transmitted as clear text.

In the Vserver aware NDMP mode, by default the authentication method is challenge. Unlike the node-scoped NDMP mode, in this mode you can enable and disable both plaintext and challenge authentication methods.

Related concepts

[User authentication in a node-scoped NDMP mode](#) on page 46

[User authentication in a Vserver aware NDMP mode](#) on page 57

What enhanced DAR functionality is

You can use the enhanced direct access recovery (DAR) functionality for directory DAR and DAR of files and NT streams. By default, enhanced DAR functionality is enabled.

Enabling enhanced DAR functionality might impact the backup performance because an offset map has to be created and written onto tape. You can enable or disable enhanced DAR in both the node-scoped and Vserver aware NDMP modes.

Scalability limits for NDMP sessions

You must be aware of the maximum number of NDMP sessions that can be established 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	Maximum number of NDMP sessions
Less than 16 GB	8
Greater than or equal to 16 GB but less than 24 GB	20
Greater than or equal to 24 GB	36

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.

Managing node-scoped NDMP mode for FlexVol volumes

NDMP follows the node-scoped behavior when the cluster is in a mixed version state. In the node-scoped NDMP mode, you can perform tape backup and restore operations at a node level.

When a cluster is in a mixed version state where Data ONTAP 8.2 and earlier versions in the Data ONTAP 8.x release family are running on nodes, NDMP follows the node-scoped behavior. This NDMP behavior continues even after upgrading to Data ONTAP 8.2. In a newly installed cluster where all nodes are running Data ONTAP 8.2, NDMP is in the Vserver aware mode by default.

If you want to perform tape backup and restore operations in the Vserver aware mode, you must explicitly disable the node-scoped NDMP mode.

You can manage NDMP at a node level by using NDMP options and commands. You must use NDMP specific credentials to access a storage system to perform tape backup and restore operations.

Related references

[Commands for managing node-scoped NDMP mode](#) on page 42

Commands for managing node-scoped NDMP mode

You can use the `system services ndmp` commands to manage NDMP at a node level.

You can use the following NDMP commands only at the advanced privilege level:

- `system services ndmp terminate`
- `system services ndmp start`
- `system services ndmp stop`
- `system services ndmp log start`
- `system services ndmp log stop`

If you want to...	Use this command...
Enable NDMP service	<code>system services ndmp on</code>
Disable NDMP service	<code>system services ndmp off</code>
Display NDMP configuration	<code>system services ndmp show</code>
Modify NDMP configuration	<code>system services ndmp modify</code>
Display default NDMP version	<code>system services ndmp version</code>
Display all NDMP sessions	<code>system services ndmp status</code>

If you want to...	Use this command...
Display detailed information about all NDMP sessions	<code>system services ndmp probe</code>
Terminate all NDMP sessions	<code>system services ndmp kill-all</code>
Change the NDMP password	<code>system services ndmp password</code>
Enable node-scoped NDMP mode	<code>system services ndmp node-scope-mode on</code>
Disable node-scoped NDMP mode	<code>system services ndmp node-scope-mode off</code>
Display node-scoped NDMP mode status	<code>system services ndmp node-scope-mode status</code>
Forcefully terminate all NDMP sessions	<code>system services ndmp terminate</code>
Start the NDMP service daemon	<code>system services ndmp start</code>
Stop the NDMP service daemon	<code>system services ndmp stop</code>
Start logging for the specified NDMP session	<code>system services ndmp log start</code>
Stop logging for the specified NDMP session	<code>system services ndmp log stop</code>

For more information about these commands, see the man pages for the `system services ndmp` commands.

Node-scoped NDMP mode options

You can use NDMP options to manage NDMP on your storage system at a node level.

You can use the following NDMP options only at the advanced privilege level:

- `ndmpd.debug.enable`
- `ndmpd.debug.filter`
- `ndmpd.abort_on_disk_error`
- `ndmpd.dump.logical_find`
- `ndmpd.fh_dir_retry_interval`
- `ndmpd.fh_node_retry_interval`
- `ndmpd.restore.vm_cache_size`

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

Option	Description	Default value
<code>ndmpd.backup.log.enable{true false}</code>	Enables or disables event logging	true
<code>ndmpd.data_port_range {start_port-end_port all}</code>	Specifies a port range on which the NDMP server can listen for data connections Note: <ul style="list-style-type: none"> <code>start_port</code> and <code>end_port</code> can have values between 1024 and 65535; <code>start_port</code> must be less than or equal to <code>end_port</code> It is best to use <code>start_port</code> and <code>end_port</code> values between 18600 and 18699. 	all
<code>ndmpd.ignore_ctime.enabled {on off}</code>	Enables or disables incremental backup of files that have their ctime changed since the previous backup.	off
<code>ndmpd.maxversion</code>	Specifies the highest NDMP version Note: Starting with Data ONTAP 8.2, only NDMP version 4 is supported.	4
<code>ndmpd.offset_map.enable {on off}</code>	Enables or disables offset map generation during backup	on
<code>ndmpd.tcpnodelay.enable {on off}</code>	Optimizes the performance of the NDMP socket through which the storage system communicates with the backup application.	off
<code>ndmpd.tcpwinsize {tcp_window_size}</code>	Specifies the TCP window size for data connection Note: The valid TCP window size range is 8192 - 7631441.	32768
<code>ndmpd.abort_on_disk_error {on off}</code>	When a disk error occurs, dump aborts a tape backup or restore operation if the value of this option is true.	off

Option	Description	Default value
ndmpd.debug.enable {on off}	Enables or disables debug logging for NDMP	off
ndmpd.debug.filter {all none normal backend filter-expression}	Specifies the NDMP modules for which debug logging is to be enabled.	none
ndmpd.dump.logical_find {default always vol_baseline vol_incr qtree_baseline qtree_incr}	Specifies whether the dump engine should follow inode-file walk or tree walk to determine what inodes need to be backed up in the backup path.	default
ndmpd.fh_node_retry_interval{time_interval}	<p>Specifies the delay, in milliseconds, to retransfer the newly generated file history data of files by the dump engine to the NDMP server.</p> <p>Note: The valid time interval range is 0 - 4294967294.</p> <p>Dump engine sends file history of all files to the NDMP server. The NDMP server sends this file history data to the backup application. If the backup application is slow in reading this file history data, NDMP server rejects the new data sent by the dump engine until the backup application reads all the accumulated data. Hence, the dump engine delays the retransfer of newly generated file history data to the NDMP server by the time interval specified.</p>	250

Option	Description	Default value
<code>ndmpd.fh_dir_retry_interval {time_interval}</code>	<p>Specifies the delay, in milliseconds, to retransfer the newly generated file history data of directories by the dump engine to the NDMP server.</p> <p>Note: The valid time interval range is 0 - 4294967294.</p> <p>Dump engine sends file history of all directories to the NDMP server. The NDMP server sends this file history data to the backup application. If the backup application is slow in reading this file history data, NDMP server rejects the new data sent by the dump engine until the backup application reads all the accumulated data. Hence, the dump engine delays the retransfer of newly generated file history data to the NDMP server by the time interval specified.</p>	250
<code>ndmpd.restore.vm_cache_size {size}</code>	<p>Increases the cache size to allocate space for restores.</p> <p>Note: The valid cache size range is 4 - 1024.</p>	64

For more information about the `options` command, see the man pages.

User authentication in a node-scoped NDMP mode

In the node-scoped NDMP mode, you must use NDMP specific credentials to access a storage system in order to perform tape backup and restore operations.

The default user ID is “root”. Before using NDMP on a node, you must ensure that you change the default NDMP password associated with the NDMP user. You can also change the default NDMP user ID.

Related references

Commands for managing node-scoped NDMP mode on page 42

Managing Vserver aware NDMP mode for FlexVol volumes

Starting with Data ONTAP 8.2, you can perform tape backup and restore operations at a Vserver level successfully if the NDMP service is enabled on the Vserver. You can back up and restore all volumes hosted across different nodes in a cluster of a Vserver if the backup application supports the CAB extension.

You can add NDMP in the allowed or disallowed protocols list by using the `vserver modify` command. By default, NDMP is in the allowed protocols list. If NDMP is added to the disallowed protocols list, NDMP sessions cannot be established. For more information about the allowed and disallowed protocols list, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

An NDMP control connection can be established on different LIF types. In a Vserver aware NDMP mode, these LIFs belong to either the data Vserver or admin Vserver. Data LIF belongs to the data Vserver and the intercluster LIF, node-management LIF, and cluster-management LIF belong to the admin Vserver.

The NDMP control connection can be established on a LIF only if the NDMP service is enabled on the Vserver that owns this LIF. In a Vserver context, the availability of volumes and tape devices for backup and restore operations depends upon the LIF type on which the NDMP control connection is established and the status of the CAB extension. If your backup application supports the CAB extension and a volume and tape device share the same affinity, then the backup application can perform a local backup or restore operation instead of a three-way backup or restore operation.

You can also manage NDMP on a per Vserver basis by using the NDMP options and commands. In the Vserver aware NDMP mode, user authentication is integrated with the role-based access control mechanism. To perform tape backup and restore operations in the node-scoped NDMP mode, you must explicitly enable the node-scoped NDMP mode.

Related concepts

[What Cluster Aware Backup extension does](#) on page 53

Related references

[Commands for managing Vserver aware NDMP mode](#) on page 48

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

Commands for managing Vserver aware NDMP mode

You can use the `vserver services ndmp` commands to manage NDMP per Vserver basis.

If you want to...	Use this command...
Enable NDMP service	<code>vserver services ndmp on</code> Note: You must ensure that NDMP service is always enabled on all nodes in a cluster. You can enable NDMP service on a node by using the <code>system services ndmp on</code> command. By default, NDMP service is always enabled on a node.
Disable NDMP service	<code>vserver services ndmp off</code>
Display NDMP configuration	<code>vserver services ndmp show</code>
Modify NDMP configuration	<code>vserver services ndmp modify</code>
Display default NDMP version	<code>vserver services ndmp version</code>
Display all NDMP sessions	<code>vserver services ndmp status</code>
Display detailed information about all NDMP sessions	<code>vserver services ndmp probe</code>
Terminate a specified NDMP session	<code>vserver services ndmp kill</code>
Terminate all NDMP sessions	<code>vserver services ndmp kill-all</code>
Generate the NDMP password	<code>vserver services ndmp generate-password</code>
Start logging for the specified NDMP session	<code>vserver services ndmp log start</code> Note: This command is available at the advanced privilege level.
Stop logging for the specified NDMP session	<code>vserver services ndmp log stop</code> Note: This command is available at the advanced privilege level.

For more information about these commands, see the man pages for the `vserver services ndmp` commands.

Vserver aware NDMP mode options

In a Vserver aware NDMP mode, you can use certain parameters to manage NDMP on a per Vserver basis.

You can use the following NDMP options only at the advanced privilege level:

- `-debug-enable`
- `-debug-filter`
- `-abort-on-disk-error`
- `-dump-logical-find`
- `-fh-dir-retry-interval`
- `-fh-node-retry-interval`
- `-restore-vm-cache-size`
- `-secondary-debug-filter`

The following table lists the NDMP options that you can use with the `vserver services ndmp modify` command:

Option	Description	Default value
<code>-maxversion</code>	Specifies the highest NDMP version Note: Starting with Data ONTAP 8.2, only NDMP version 4 is supported.	4
<code>-ignore-ctime-enabled {true false}</code>	Enables or disables incremental backup of files that have their ctime changed since the previous backup	false
<code>-offset-map-enable {true false}</code>	Enables or disables offset map generation during backup	true
<code>-tcpnodelay {true false}</code>	Optimizes the performance of the NDMP socket through which the storage system communicates with the DMA	false

Option	Description	Default value
-tcpwinsize { <i>tcp_window_size</i> }	Specifies the TCP window size for data connection Note: The valid TCP window size range is 8192 - 7631441.	32768
-data-port-range { <i>start_port-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
-backup-log-enable {true false}	Enables or disables event logging	true
-per-qtrees-exclude-enable {true false}	Enables the users to specify the files or directories that need to be excluded when backing up data on a per qtrees basis.	false
-authtype {challenge plaintext plaintext,challenge}	Specifies whether a storage system accepts plaintext, challenge, or both to authenticate NDMP session requests	challenge
-enable {true false}	Enables NDMP per Vserver basis	false
-preferred-interface-role In the cluster context, {intercluster node-mgmt cluster-mgmt} In the Vserver context, {intercluster data}	Specifies the preferred network interface types to be used for NDMP data connections on a particular Vserver. Note: The format of this option is a comma-separated list of LIF types.	In the cluster context, {intercluster, node-mgmt, cluster-mgmt} In the Vserver context, {intercluster, data}

Option	Description	Default value
-debug-enable {true false}	Enables or disables debug logging for NDMP	false
-debug-filter {all none normal backend filter-expression}	Specifies the NDMP modules for which debug logging is to be enabled	none
-abort-on-disk-error {true false}	When a disk error occurs, dump aborts a tape backup or restore operation if the value of this option is true.	false
-dump-logical-find {default always vol_baseline vol_incr qtree_baseline qtree_incr}	Specifies whether the dump engine should follow inode-file walk or tree walk to determine what inodes need to be backed up in the backup path	default
-fh-dir-retry-interval {time_interval}	<p>Specifies the delay, in milliseconds, to retransfer the newly generated file history data of directories by the dump engine to the NDMP server.</p> <p>Note: The valid time interval range is 0 - 4294967294.</p> <p>Dump engine sends file history of all directories to the NDMP server. The NDMP server sends this file history data to the backup application. If the backup application is slow in reading this file history data, NDMP server rejects the new data sent by the dump engine until the backup application reads all the accumulated data. Hence, the dump engine delays the retransfer of newly generated file history data to the NDMP server by the time interval specified.</p>	250

Option	Description	Default value
<code>-fh-node-retry-interval {time_interval}</code>	<p>Specifies the delay, in milliseconds, to retransfer the newly generated file history data of files by the dump engine to the NDMP server.</p> <p>Note: The valid time interval range is 0 - 4294967294.</p> <p>Dump engine sends file history of all files to the NDMP server. The NDMP server sends this file history data to the backup application. If the backup application is slow in reading this file history data, NDMP server rejects the new data sent by the dump engine until the backup application reads all the accumulated data. Hence, the dump engine delays the retransfer of newly generated file history data to the NDMP server by the time interval specified.</p>	250
<code>-restore-vm-cache-size {size}</code>	<p>Increases the cache size to allocate space for restores</p> <p>Note: The valid cache size range is 4 - 1024.</p>	64
<code>-secondary-debug- filter{ipaddr=[ip-address]}</code>	<p>This option is a string of <code>tag=value</code> pairs. Currently, the supported tag is <code>ipaddr</code>. This allows debug logging to be enabled for NDMP control connections established on LIFs whose IP address is specified as the value of the tag.</p> <p>Note: You can specify more than one IP address separated by commas.</p>	none

For more information about the `vserver services ndmp modify` command, see the man pages.

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:

- CAB (Cluster Aware Backup extension)
- CAE (Connection Address Extension) for IPv6 support

What Cluster Aware Backup extension does

CAB (Cluster Aware Backup) is an NDMP v4 protocol extension. This extension enables the NDMP server to establish a data connection on a node that owns a volume. This also enables the backup application to determine if volumes and tape devices are located on the same node in a cluster.

To enable the NDMP server to identify the node that owns a volume and to establish a data connection on such a node, the backup application must support the CAB extension. CAB extension requires the backup application to inform the NDMP server about the volume to be backed up or restored prior to establishing the data connection. This allows the NDMP server to determine the node that hosts the volume and appropriately establish the data connection.

With the CAB extension supported by the backup application, the NDMP server provides affinity information about volumes and tape devices. Using this affinity information, the backup application can perform a local backup instead of a three-way backup if a volume and tape device are located on the same node in a cluster.

Availability of volumes and tape devices for backup and restore on different LIF types

You can configure a backup application to establish an NDMP control connection on any of the LIF types in a cluster. In Vserver aware NDMP mode, you can determine the availability of volumes and tape devices for backup and restore operations depending upon these LIF types and the status of the CAB extension.

The following tables show the availability of volumes and tape devices for NDMP control connection LIF types and the status of the CAB extension:

Availability of volumes and tape devices when CAB extension is not supported by the backup application

NDMP control connection LIF type	Volumes available for backup or restore	Tape devices available for backup or restore
Node-management LIF	All volumes hosted by a node	Tape devices connected to the node hosting the node-management LIF
Data LIF	Only volumes that belong to the Vserver hosted by a node that hosts the data LIF	None
Cluster-management LIF	All volumes hosted by a node that hosts the cluster-management LIF	None
Intercluster LIF	All volumes hosted by a node that hosts the intercluster LIF	Tape devices connected to the node hosting the intercluster LIF

Availability of volumes and tape devices when CAB extension is supported by the backup application

NDMP control connection LIF type	Volumes available for backup or restore	Tape devices available for backup or restore
Node-management LIF	All volumes hosted by a node	Tape devices connected to the node hosting the node-management LIF
Data LIF	All volumes that belong to the Vserver that hosts the data LIF	None
Cluster-management LIF	All volumes in the cluster	All tape devices in the cluster
Intercluster LIF	All volumes in the cluster	All tape devices in the cluster

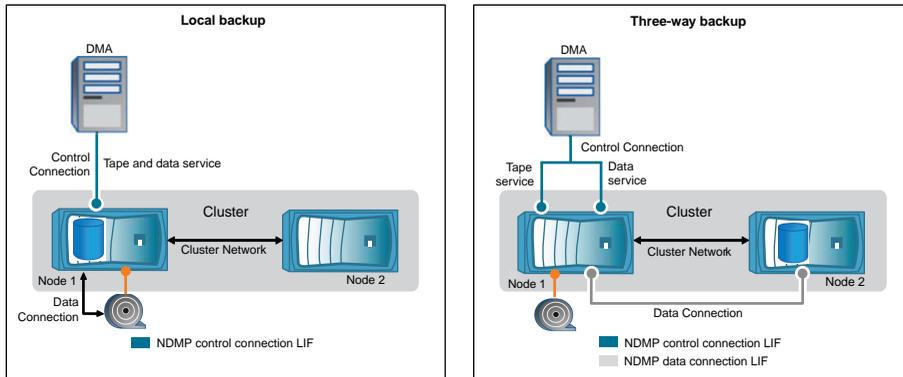
What affinity information is

With the backup application being CAB aware, the NDMP server provides unique location information about volumes and tape devices. Using this affinity information, the backup application can perform a local backup instead of a three-way backup if a volume and a tape device share the same affinity.

If the NDMP control connection is established on a node management LIF, cluster management LIF, or an intercluster LIF, the backup application can use the affinity information to determine if a

volume and tape device are located on the same node and then perform either a local or a three-way backup or restore operation. If the NDMP control connection is established on a data LIF, then the backup application always performs a three-way backup.

Local NDMP backup and Three-way NDMP backup



Using the affinity information about volumes and tape devices, the DMA (backup application) performs a local NDMP backup on the volume and tape device located on Node 1 in the cluster. If the volume moves from Node 1 to Node 2, affinity information about the volume and tape device changes. Hence, for a subsequent backup the DMA performs a three-way NDMP backup operation. This ensures continuity of the backup policy for the volume irrespective of the node to which the volume is moved to.

Related concepts

[What Cluster Aware Backup extension does](#) on page 53

What preferred-interface-role is

In a Vserver aware NDMP mode, you can control the LIF type on which an NDMP data connection is established.

You can specify the LIF types required for the NDMP data connection by using the `preferred-interface-role` option. The format of this option is a comma-separated list of LIF types. During an NDMP data connection establishment, NDMP chooses an IP address that belongs to the LIF type as specified by the `preferred-interface-role` option. The LIF type chosen is in the order as defined by this option. If the IP addresses do not belong to any of the LIF types, the NDMP data connection cannot be established.

Related references

[Vserver aware NDMP mode options](#) on page 49

NDMP data connection types

In a Vserver aware NDMP mode, the supported NDMP data connection types depend on the NDMP control connection LIF type and the status of the CAB extension. This NDMP data connection type indicates whether you can perform a local or a three-way NDMP backup or restore operation.

You can perform a three-way NDMP backup or restore operation over a TCP or TCP/IPv6 network. The following tables show the NDMP data connection types based on the NDMP control connection LIF type and the status of the CAB extension.

NDMP data connection type when CAB extension is not supported by the backup application

NDMP control connection LIF type	NDMP data connection type
Node-management LIF	LOCAL, TCP, TCP/IPv6
Data LIF	TCP, TCP/IPv6
Cluster-management LIF	TCP, TCP/IPv6
Intercluster LIF	LOCAL, TCP, TCP/IPv6

NDMP data connection type when CAB extension is supported by the backup application

NDMP control connection LIF type	NDMP data connection type
Node-management LIF	LOCAL, TCP, TCP/IPv6
Data LIF	TCP, TCP/IPv6
Cluster-management LIF	LOCAL, TCP, TCP/IPv6
Intercluster LIF	LOCAL, TCP, TCP/IPv6

Note: In Data ONTAP 8.2, you cannot configure intercluster LIFs with IPv6 addresses. For more information about IPv6 and intercluster LIF, see *Clustered Data ONTAP Network Management Guide*.

Related concepts

[What Cluster Aware Backup extension does](#) on page 53

User authentication in a Vserver aware NDMP mode

In Vserver aware NDMP mode, NDMP user authentication is integrated with the role-based access control mechanism. In a Vserver context, the NDMP user must belong to either the "vsadmin" or "vsadmin_backup" role. In a cluster context, the NDMP user must belong to either the "admin" or "backup" role.

In this mode, you must generate an NDMP password for a given user account, which is created through role-based access control mechanism. Cluster users in an admin or backup role can access a node-management LIF, cluster-management LIF, or an intercluster LIF. Users in a vsadmin_backup or vsadmin role can access only the data LIF. Hence, depending upon the role of a user, the availability of volumes and tape devices for backup and restore operations vary.

This mode also supports user authentication for NIS and LDAP users. Hence, NIS and LDAP users can access multiple Vservers with a common userid and password. However, NDMP authentication does not support Active Directory users.

In this mode, a user account must be associated with the SSH application and the authentication method "User password".

For more information about role-based access control, see *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Related references

[Commands for managing Vserver aware NDMP mode](#) on page 48

Generating an NDMP-specific password for NDMP users

In the Vserver aware NDMP mode, you must generate a password for a specific user ID. The generated password is based on the actual login password for the NDMP user. If the actual login password changes, you must generate the NDMP-specific password again.

Steps

1. Use the `vserver services ndmp generate-password` command to generate an NDMP-specific password.

You can use this password in any current or future NDMP operation that requires password input.

Note: From a Vserver context, you can generate NDMP passwords for users belonging only to that Vserver.

Example

The following example shows how to generate an NDMP-specific password for a user ID user1:

```
cluster1::vserver services ndmp> generate-password -vserver vs1 -user
user1

Vserver: vs1
User: user1
Password: jWZiNt57huPOoD8d
```

2. If you change the password to your regular storage system account, repeat this procedure to obtain your new NDMP-specific password.

Understanding dump engine for FlexVol volumes

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.

Dump does not support backup and restore of Infinite Volumes.

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

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.

Note: After reverting to Data ONTAP 8.1.x, to perform an incremental backup operation you must first perform a baseline backup operation.

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.

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
- Junction paths
Junction paths are backed up as symbolic links.
- 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 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

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.

You can use a blocking factor from 4 to 256. 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`.

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.

Starting with Data ONTAP 8.2, you can restore data to an online mapped LUN. However, host applications cannot access this LUN until the restore operation is complete. After the restore operation is complete, the host cache of the LUN data should be flushed to guarantee coherency with the restored data.

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.

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.
 - A 7-Mode LUN is restored as a regular file on a clustered Data ONTAP volume.
- A 7-Mode volume can be restored to a clustered Data ONTAP volume.
- 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.

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.

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.

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.

Tape backup and restore between Data ONTAP 7-Mode and clustered Data ONTAP

You can restore data backed up from a 7-Mode or a clustered Data ONTAP storage system to either a 7-Mode or a clustered Data ONTAP storage system.

The following tape backup and restore operations are supported between Data ONTAP operating in 7-Mode and clustered Data ONTAP:

- Backing up a 7-Mode volume to a tape drive connected to a clustered Data ONTAP system
- Backing up a clustered Data ONTAP volume to a tape drive connected to a 7-Mode system
- Restoring backed up data of a 7-Mode volume from a tape drive connected to a clustered Data ONTAP system
- Restoring backed up data of a clustered Data ONTAP volume from a tape drive connected to a 7-Mode system
- Restoring a 7-Mode volume to a clustered Data ONTAP volume

Note: A 7-Mode LUN is restored as a regular file on a clustered Data ONTAP volume.

- Restoring a clustered Data ONTAP volume to a 7-Mode volume

Note: A clustered Data ONTAP LUN is restored as a regular file on a 7-Mode volume.

How dump backs up data from a vault destination volume

You can perform tape backup operations on data that is mirrored on the vault destination volume. You can back up only the data that is mirrored on the vault destination volume to tape, but not the vault relationship metadata.

When you break the data protection mirror relationship (`snapmirror break`) or when a SnapMirror resynchronization occurs, you must always perform a baseline backup.

For more information about vault destination volumes, see *Clustered Data ONTAP Data Protection Guide*

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

How dump works with SFO and ARL

Storage Failover (SFO) and Aggregate Relocate (ARL) allow an aggregate to be relocated from one node to another in a cluster. The `-override-vetoes` option determines the behavior of dump engine during an SFO or ARL operation.

When a dump backup or restore operation is running and the `-override-vetoes` option is set to false, a user-initiated SFO or ARL operation is stopped. However, if the `-override-vetoes` option is set to true then the SFO or ARL operation is continued and the dump backup or restore operation is aborted. When an SFO or ARL operation is automatically initiated by the storage system, an active dump backup or restore operation is always aborted. Dump backup and restore operations are not restartable after SFO or ARL operations are complete.

After an SFO or ARL operation is complete, in the Vserver aware NDMP mode, you can continue performing incremental tape backup and restore operations without reconfiguring backup policies. In the node-scoped NDMP mode, you must reconfigure the backup policy to enable backup and restore operations from the destination aggregate. To reconfigure the backup policy, see your backup application documentation.

For more information about the SFO operations, see the *Clustered Data ONTAP High-Availability Configuration Guide* and for more information about ARL operations, see the *Clustered Data ONTAP Physical Storage Management Guide*.

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

How dump works with volume movement operation

Starting with Data ONTAP 8.2, you can perform tape backup and restore operations using the dump engine and volume movement operations simultaneously. You can continue performing incremental backups both on read/write and read-only volumes after the volume movement operation.

Note: When a volume movement occurs, if the volume belonging to a different Vserver on the destination node has the same name as that of the moved volume, then you cannot perform incremental backup operations of the moved volume.

In the Vserver aware NDMP mode, you can move volumes without reconfiguring the backup policy. In the node-scoped NDMP mode, you must reconfigure the backup policy to enable backup and restore operations from the destination node after the volume movement operation is complete. To reconfigure the backup policy, see your backup application documentation.

Monitoring tape backup and restore operations for FlexVol volumes

You can view the event log files to monitor the tape backup and restore operations. Data ONTAP automatically logs significant dump and restore events and the times at which they occur in a log file named `backup` in the controller's `/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).

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 at the nodeshell. You can view these event log files to monitor tape backup and restore operations.

Steps

1. To access the nodeshell, enter the following command:

```
node run -node node_name
```

node_name is the name of the node.

2. To access the event log files for tape backup and restore operations, enter 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 *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

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	NDMP dump is started	Dump level and the type of dump

Event	Description	Event information
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
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 internal metafile has been located	The level and time of the base dump (for incremental dumps only)

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	NDMP restore is started	Restore level and the type of restore
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

Event	Description	Event information
Error	Restore encounters an unexpected event	Error message

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 at the clustershell:

```
options -option-name backup.log.enable -option-value {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 of FlexVol volumes

You might encounter an error message when performing a dump 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.

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.

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:
- 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**
- 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, *size* 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.

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

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

Cause The blocking factor specified for the local restore is incorrect.

Corrective action Retry the restore job with `block_size1` 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 `ndmpd invalid version number: version_number`

Cause The NDMP version specified is not supported by the storage system.

Corrective action Specify NDMP version 4.

ndmpd session `session_ID` not active

Message `ndmpd session session_ID not active`

Cause The NDMP session might not exist.

Corrective action Use the `ndmpd status` command to view the active NDMP sessions.

Could not obtain vol ref for Volume `volume_name`

Message `Could not obtain vol ref for Volume vol_name`

Cause The volume reference could not be obtained because the volume might be in use by other operations.

Corrective action Retry the operation later.

Data connection type ["NDMP4_ADDR_TCP"|"NDMP4_ADDR_TCP_IPv6"] not supported for ["IPv6"|"IPv4"] control connections

Message `Data connection type ["NDMP4_ADDR_TCP"|"NDMP4_ADDR_TCP_IPv6"] not supported for ["IPv6"|"IPv4"] control connections`

Cause	In node-scoped NDMP mode, the NDMP data connection established must be of the same network address type (IPv4 or IPv6) as the NDMP control connection.
Corrective action	Contact your backup application vendor.

DATA LISTEN: CAB_DATA_CONN_PREPARE not sent or addr type mismatch

Message	DATA LISTEN: CAB_DATA_CONN_PREPARE not sent or addr type mismatch
Cause	The backup application has not negotiated the CAB extension, yet has sent the CAB message. Or, the backup application has negotiated the CAB extension, but sends the incorrect CAB message.
Corrective action	Disable the CAB extension using the backup application by contacting your backup application vendor.

DATA CONNECT: CAB_DATA_CONN_PREPARE not sent or addr type mismatch

Message	DATA CONNECT: CAB_DATA_CONN_PREPARE not sent or addr type mismatch
Cause	The backup application has not negotiated the CAB extension, yet has sent the CAB message. Or, the backup application has negotiated the CAB extension, but sends the incorrect CAB message.
Corrective action	Disable the CAB extension using the backup application by contacting your backup application vendor.

Error:show failed: Cannot get password for user '<username>'

Message	Error: show failed: Cannot get password for user '<username>'
Cause	Incomplete user account configuration for NDMP
Corrective action	Ensure that the user account is associated with the SSH access method and the authentication method is user password.

Dump error messages

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

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.

Dumps temporarily disabled on volume, try again

Message	<code>Dumps temporarily disabled on volume, try again</code>
Cause	Volume movement is in progress.
Corrective action	Wait for volume movement operation to complete and then perform the backup operation again.

No files were created

Message	<code>No files were created</code>
Cause	A directory DAR was attempted without enabling the enhanced DAR functionality.
Corrective action	Enable the enhanced DAR functionality and retry the DAR.

Unable to lock a snapshot needed by dump

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

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.

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.

Tape backup and restore of Infinite Volumes

You can incrementally back up data in Infinite Volumes to tape and restore data from tape using a SnapDiff-supported backup application that supports backup and restore over a volume mounted with the NFS or CIFS protocols. NDMP is not supported for Infinite Volumes.

The cluster administrator can perform all of the tape backup tasks. The Vserver administrator can perform only the following tape backup tasks:

- Enabling SnapDiff on an Infinite Volume
- Selecting a Snapshot copy of an Infinite Volume for SnapDiff to use
- Creating a baseline for incremental tape backup
- Backing up only changed files to tape
- Restoring files to disk
- Creating SnapDiff resources after expanding a destination Infinite Volume onto new nodes
- Disabling SnapDiff on an Infinite Volume

Protocols for tape backup and restore of Infinite Volumes

You can use some NFS and CIFS protocols to back up files in Infinite Volumes to tape and restore files from tape. However, some protocols cannot back up the streams or access control lists (ACLs) associated with the files.

You should consider the following implications when you choose a protocol to back up and restore files for Infinite Volumes:

- If you are using NFSv3 to back up and restore files that use ACLs, NFSv3 backs up and restores only the files, not the ACLs.
You should use another supported protocol to back up and restore both the files and the ACLs.
- If you want to back up and restore the files and the streams, you should use CIFS.
When a file uses streams, only CIFS can back up and restore the files and the streams associated with the files.

For more information about the protocols supported for Infinite Volumes, see the *Clustered Data ONTAP File Access and Protocols Management Guide*.

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

Tape backup topology to use for Infinite Volumes

For Infinite Volumes, you must use the tape backup topology specified by the backup application for volumes mounted with the NFS or CIFS protocol. You cannot use an NDMP-supported tape backup topology for Infinite Volumes.

Tape backup topologies for mounted volumes do not have any components installed on NetApp controllers and use the network to communicate with Data ONTAP. For example, when you back up data for Infinite Volumes, the data flows as follows:

1. The NetApp controller sends data over the network to the backup application client.
2. The backup application client communicates over the network with the backup application.
3. The backup application communicates over a direct connection with the tape device.

Each backup application specifies how to set up tape libraries or tape drives to back up data over a mounted volume. For more information about tape library and tape drive requirements for backup applications, see your backup application documentation.

How storage classes affect tape backup of Infinite Volumes

Setup and management of incremental tape backup is the same for Infinite Volumes with or without storage classes. You must configure tape backup for the entire Infinite Volume. You cannot configure tape backup for individual storage classes in Infinite Volumes.

Understanding incremental tape backup of Infinite Volumes using SnapDiff and Snapshot copies

The storage capacity potential of an Infinite Volume is larger than what a traditional file-scanning backup application can back up in a reasonable time. An incremental backup of Infinite Volumes to tape by using SnapDiff and Snapshot copies is the only viable solution for large Infinite Volumes.

What SnapDiff is

SnapDiff is an internal Data ONTAP engine that quickly identifies the file and directory differences between two Snapshot copies.

By finding the differences between two Snapshot copies, SnapDiff eliminates the file scanning requirements of a traditional backup application during an incremental backup, which reduces the backup processing to only the time it takes to write the changed or added data.

When incrementally backing up an Infinite Volume to tape using SnapDiff, the backup application uses the SnapDiff application programming interfaces (APIs) to communicate with the SnapDiff engine to identify new, changed, and deleted files between two Snapshot copies of the active file system in an Infinite Volume. The differencing process uses the namespace constituent and

namespace mirror constituents in an Infinite Volume to determine names for the list of new, changed, and deleted files. Once these changes are identified, the backup application backs up the identified data from the list produced during the differencing process.

SnapDiff support for backup applications

You must enable SnapDiff on an Infinite Volume for a SnapDiff-supported backup application to incrementally back up the data to tape. Only supported third-party backup applications can use SnapDiff.

Because the backup application requires integration with the SnapDiff APIs to use SnapDiff for incremental tape backup of Infinite Volumes, only supported third-party backup applications can use SnapDiff for incremental tape backup of Infinite Volumes.

Guidelines for using LIFs with SnapDiff

You should connect backup applications to the data LIF for the Vserver with Infinite Volume to allow SnapDiff to communicate large amounts of information. You should not connect backup applications to the cluster-management LIF or to the node-management LIF.

SnapDiff differs from other application programming interfaces (APIs) in the amount of information handled. Where APIs typically retrieve a small amount of control and management information, SnapDiff retrieves information about each and every file in an Infinite Volume. As a result, the amount of information reported by SnapDiff to the backup application can be large, depending on the number of files in the Infinite Volume. Because of the potentially large amount of information, you should connect the backup application to the data LIF for the Vserver with Infinite Volume. For this network configuration, you must change the firewall settings on the data LIF to allow SnapDiff to communicate.

You should not connect the backup application to the cluster-management LIF for the Vserver with Infinite Volume. However, if you do connect to the cluster-management LIF, backup applications must include the name of the Vserver and the name of the Infinite Volume in the SnapDiff APIs to identify which Vserver to use.

You should not connect the backup application to the node-management LIF because node-management LIFs are aware of nodes, not clusters, and an Infinite Volume spans multiple nodes.

What Snapshot copies you should back up to tape

You can back up any of the Snapshot copies that display a valid state for an Infinite Volume when you use the `volume snapshot show` command.

Note: When using CIFS to back up Snapshot copies on an Infinite Volume, you should create a share at `/NS/.snapshot` or below.

For more information about Snapshot copies, see the *Clustered Data ONTAP Data Protection Guide*.

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

What a namespace mirror constituent is

A namespace mirror constituent is an intracluster data protection mirror copy of the namespace constituent in an Infinite Volume. The namespace mirror constituent performs two roles: It provides data protection of the namespace constituent, and it supports SnapDiff for incremental tape backup of Infinite Volumes.

When namespace mirror constituents are created

When you create an Infinite Volume, one namespace mirror constituent is automatically created to provide data protection for the namespace constituent. When you enable SnapDiff for an Infinite Volume, additional namespace mirror constituents are automatically created for SnapDiff for incremental tape backup of Infinite Volumes.

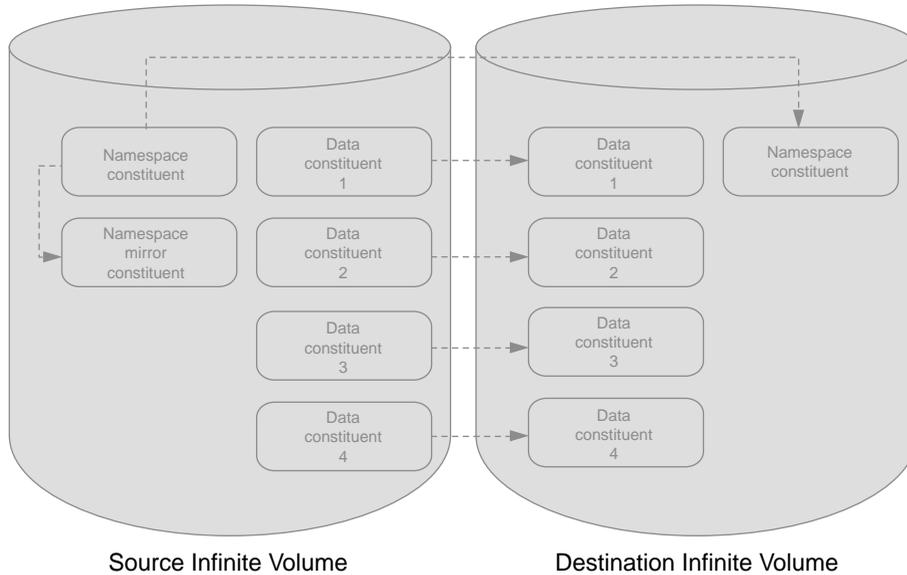
When you create a read/write Infinite Volume that spans two or more nodes in a cluster, one namespace mirror constituent is automatically created, and a data protection mirror relationship is automatically created between the namespace constituent and the namespace mirror constituent. The data protection mirror relationship is updated every five minutes. The data protection mirror relationship is an automatic process for an Infinite Volume.

When you enable SnapDiff on an Infinite Volume that spans three or more nodes, additional namespace mirror constituents are automatically created for SnapDiff to use for incremental tape backup of Infinite Volumes. A namespace mirror constituent is created on each node with a data constituent, except the node with the namespace constituent and the node with the namespace mirror constituent that was created to provide data protection for the namespace constituent. Namespace mirror constituents created to support SnapDiff are updated daily or as configured for SnapDiff. A SnapMirror license is not required to enable SnapDiff.

When you create a destination Infinite Volume for a data protection mirror relationship, a namespace mirror constituent is not created on the destination Infinite Volume. However, if you enable SnapDiff on a destination Infinite Volume, namespace mirror constituents are automatically created for use by SnapDiff. You must initialize the data protection mirror relationship between the source and destination Infinite Volumes before you can enable SnapDiff.

The following illustration shows an Infinite Volume in a data protection mirror relationship. SnapDiff is disabled on the source and the destination Infinite Volumes. With SnapDiff disabled, a namespace mirror constituent is created on the source Infinite Volume to provide data protection for the namespace constituent, and no other namespace mirror constituents are created.

Data Protection Mirror Relationship



How modifying an Infinite Volume affects namespace mirror constituents

Namespace mirror constituents are automatically created or deleted when you modify Infinite Volumes, depending on the SnapDiff setting for the volume. The behavior differs slightly between a read/write Infinite Volume and a read-only destination Infinite Volume in a data protection mirror relationship.

When you modify a read/write Infinite Volume to increase its size and create data constituents on new nodes, and SnapDiff is already enabled on the Infinite Volume, namespace mirror constituents are automatically created on new nodes as required. When you modify a read/write or read-only Infinite Volume, and SnapDiff is already disabled, all namespace mirror constituents—except one namespace mirror constituent—are automatically deleted. One namespace mirror constituent is retained to provide data protection for the namespace constituent.

The following table summarizes when namespace mirrors are automatically created and deleted when you modify a read/write or read-only Infinite Volume:

Infinite Volume	SnapDiff setting	Task	Result
Read/write Infinite Volume	Enabled	Add new nodes to the cluster and increase volume size by using <code>volume modify</code> Note: You must increase the size enough to create new data constituents on the new nodes.	Namespace mirror constituents are automatically created on new nodes as required.
Read-only destination Infinite Volume in a data protection mirror relationship			Namespace mirror constituents are not automatically created on new nodes, although SnapDiff is enabled on the destination Infinite Volume. You must expand the destination Infinite Volume onto new nodes and update the data protection mirror relationship before you reenable SnapDiff to automatically create namespace mirror constituents on new nodes.
Read/write Infinite Volume and read-only destination Infinite Volume in a data protection mirror relationship	Disabled	Modify the volume by using <code>volume modify</code>	All namespace mirror constituents—except one namespace mirror constituent—are automatically deleted.

How namespace mirror constituents affect the state of an Infinite Volume

Infinite Volumes with SnapDiff enabled temporarily display a mixed state while namespace mirror constituents are created. After the namespace mirror constituents are created, Infinite Volumes automatically display an online state. You cannot start new operations until the namespace mirror constituents are created, and the Infinite Volume displays an online state.

The amount of time required to create all of the namespace mirror constituents depends on the size of the namespace constituent and the number of nodes an Infinite Volumes spans.

If you enable SnapDiff when you create an Infinite Volume, the namespace constituent is empty, and little time is required to create namespace mirror constituents on all of the nodes that an Infinite Volume spans. As a result, the Infinite Volume displays a mixed state for a small amount of time.

If you enable SnapDiff after an Infinite Volume has existed for some time, the namespace constituent contains data, and more time is required to create namespace mirror constituents. The Infinite Volume displays a mixed state until all of the namespace mirror constituents are created and initialized on all of the nodes with data constituents. The number of nodes that the Infinite Volume spans also affects the amount of time it takes to create and initialize namespace mirror constituents.

When you resize an Infinite Volume with SnapDiff enabled, namespace mirror constituents are created and initialized on all the new nodes with data constituents. Depending on the size of the namespace constituent and the number of new nodes, it can take some time to create and initialize a namespace mirror constituent on each new node. The Infinite Volume displays a mixed state while the namespace mirror constituents are being initialized.

Before you start a tape backup operation, you should wait until all of the namespace mirror constituents are created and initialized, and the Infinite Volume displays an online state.

How aggregates are selected for namespace mirror constituents

One process is used to automatically select an aggregate for the namespace mirror constituent created to provide data protection of the namespace constituent in an Infinite Volume. Another process is used to automatically select aggregates for the namespace mirror constituents created for SnapDiff and incremental tape backup of Infinite Volumes.

Depending on the number of nodes, Data ONTAP executes the following process to select an aggregate for the namespace mirror constituent used for data protection of the namespace constituent:

Number of nodes	Aggregate selection process
For a Vserver with Infinite Volume that uses aggregates from two nodes	Data ONTAP considers the node without the namespace constituent, and chooses an aggregate with the largest available space.
For a Vserver with Infinite Volume that uses aggregates from multiple nodes	Data ONTAP considers all nodes that contain a data constituent, and chooses an aggregate with the largest available space. The aggregate selection process excludes the node that contains the namespace constituent. The aggregate selection process prefers to also exclude the HA pair of the node that contains the namespace constituent, but uses the node when no other nodes are available.

Data ONTAP executes the following process to select aggregates for namespace mirror constituents used for SnapDiff and incremental tape backup:

1. Identifies all the nodes in the cluster that an Infinite Volume spans
2. Identifies the nodes that have a data constituent for the Infinite Volume
3. Creates a namespace mirror constituent on each node that has a data constituent for the Infinite Volume, except for the node that has the namespace constituent and the node in which the first namespace mirror constituent is created

When you expand an Infinite Volume with SnapDiff enabled onto new nodes, additional namespace mirror constituents are automatically created, and aggregates are automatically selected for the new namespace mirror constituents.

Space and node requirements for namespace mirror constituents

Up to 25% of the total size of an Infinite Volume can be used for the namespace constituent and namespace mirror constituents. The number of namespace mirror constituents created depends on whether SnapDiff is enabled for the Infinite Volume and the number of nodes the Infinite Volume spans.

An Infinite Volume must span two or more nodes before a namespace mirror constituent is created to provide data protection for the namespace constituent. An Infinite Volume must span three or more nodes before additional namespace mirror constituents are created for SnapDiff to use.

The following table summarizes the impact of nodes and SnapDiff on the number of namespace mirror constituents created for an Infinite Volume:

When an Infinite Volume spans...	The following constituents are created...
Two or more nodes in a cluster and SnapDiff is disabled for the Infinite Volume	<ul style="list-style-type: none"> • One namespace constituent • One namespace mirror constituent for data protection
Two nodes in a cluster and SnapDiff is enabled for the Infinite Volume	<ul style="list-style-type: none"> • One namespace constituent • One namespace mirror constituent for data protection SnapDiff can also use this namespace mirror constituent. • No additional namespace mirror constituents SnapDiff has no additional namespace mirror constituents to use.
Three or more nodes in a cluster and SnapDiff is enabled for the Infinite Volume	<ul style="list-style-type: none"> • One namespace constituent • One namespace mirror constituent for data protection • One namespace mirror constituent on each node that contains a data constituent, except for the node that contains the namespace constituent and the node that contains the namespace mirror constituent for data protection

Each namespace mirror constituent is the same size as the namespace constituent. The maximum size for the namespace constituent is 10 TB. The space available in an Infinite Volume is divided evenly among the namespace constituent and the namespace mirror constituents. The sum of the space used

by the namespace constituent and each namespace mirror constituent can consume up to 25% of the total size of the Infinite Volume or 100 TB of space. For example, consider an Infinite Volume that is 80 TB in size, spans two nodes, and has SnapDiff disabled. The Infinite Volume contains a 10 TB namespace constituent and a 10 TB namespace mirror constituent. The namespace constituent and the namespace mirror constituent consume 25% of the total size of the Infinite Volume.

Note: The total size of the namespace constituent and namespace mirror constituents should not exceed 100 TB of space because an Infinite Volume should span a maximum of 10 nodes, and each namespace constituent or namespace mirror constituent can be a maximum of 10 TB. As a result, 10 x 10 TB = 100 TB of space.

How the file system size setting for an Infinite Volume affects SnapDiff

You must be aware of how the `filesystem-size-fixed` parameter for the `volume modify` command affects whether you can enable or disable SnapDiff for an Infinite Volume. SnapDiff is required for incremental tape backup of Infinite Volumes.

When you modify an Infinite Volume by using the `volume modify` command, you can set the `filesystem-size-fixed` parameter. The setting for the `filesystem-size-fixed` parameter can affect whether you can enable or disable SnapDiff for an Infinite Volume. The parameter affects SnapDiff as follows:

- For a read/write Infinite Volume, the `filesystem-size-fixed` parameter must be set to `false` to enable SnapDiff.
- For an Infinite Volume in a data protection mirror relationship, the `filesystem-size-fixed` parameter must be set to `false` for the source Infinite Volume, but the parameter can be set to `true` or `false` for the destination Infinite Volume.

Setting up incremental tape backup of Infinite Volumes

Setting up incremental tape backup involves verifying that the Infinite Volume has sufficient space for SnapDiff resources, enabling SnapDiff on the Infinite Volume, and creating a baseline for incremental tape backup.

Steps

1. *Verifying space requirements in an Infinite Volume for SnapDiff resources (cluster administrators only)* on page 89
2. *Enabling SnapDiff on an Infinite Volume* on page 91
3. *Selecting a Snapshot copy of an Infinite Volume for SnapDiff to use* on page 93
4. *Creating a baseline for incremental tape backup* on page 93

Verifying space requirements in an Infinite Volume for SnapDiff resources (cluster administrators only)

Before you enable SnapDiff on an Infinite Volume, you must ensure that the Infinite Volume has enough space for all of the namespace mirror constituents that will be created for SnapDiff to use for tape backup.

Before you begin

- You must have cluster administrator privileges.
- Determine whether to use this procedure.
 - When an Infinite Volume spans two nodes in a cluster, no new namespace mirror constituents are created. Instead, SnapDiff uses the namespace constituent and the namespace mirror constituent created for data protection of the namespace constituent. In this case, you do not need to verify space requirements.
 - When an Infinite Volume spans three or more nodes, additional namespace mirror constituents are created when you enable SnapDiff, and you need to verify that the Infinite Volume has enough space.
 - When an Infinite Volume is in a data protection mirror relationship, you can enable SnapDiff on the destination Infinite Volume. Additional namespace mirror constituents are created for destination Infinite Volumes that span two or more nodes when you enable SnapDiff. When a destination Infinite Volume spans two or more nodes, you must verify that the destination Infinite Volume has enough space before enabling SnapDiff.

Steps

1. View the size of the namespace constituent on the Infinite Volume by using the `volume show` command with the `-is-constituent true` parameter.

Each namespace mirror constituent is the same size as the namespace constituent.

Example

In the following example, the namespace constituent is named `repo_vol_ns`, and it is 256 MB in size. The namespace mirror constituent created for data protection of the namespace constituent is named `repo_vol_ns_mirror0001`, and it is 256 MB in size.

```
cluster1::> volume show -is-constituent true
Vserver Volume                Aggregate State  Type  Size  Available Used%
-----
vs0    repo_vol_1024_data0001 aggr1    online RW   384MB 364.7MB 5%
vs0    repo_vol_1024_data0002 aggr2    online RW   384MB 364.7MB 5%
vs0    repo_vol_1024_data0003 aggr3    online RW   384MB 364.7MB 5%
vs0    repo_vol_1024_data0004 vs_aggr  online RW   384MB 364.7MB 5%
vs0    repo_vol_ns            aggr1    online RW   256MB 243.0MB 5%
vs0    repo_vol_ns_mirror0001 aggr2    online DP   256MB 243.1MB 5%
6 entries were displayed.
```

2. Make a list of the aggregates that have a data constituent.

In the above example, the data constituents are named `repo_vol_1024_data0001`, `repo_vol_1024_data0002`, `repo_vol_1024_data0003`, and `repo_vol_1024_data0004`, and the data constituents use the following aggregates: `aggr1`, `aggr2`, `aggr3`, and `vs_aggr`.

3. Identify which of the aggregates that contain a data constituent also have a namespace constituent or a namespace mirror constituent.

In the above example, the following aggregates with data constituents also contain either a namespace constituent or a namespace mirror constituent:

- `aggr1` contains a data constituent named `repo_vol_1024_data0001` and a namespace constituent named `repo_vol_ns`.
- `aggr2` contains a data constituent named `repo_vol_1024_data0002` and a namespace mirror constituent named `repo_vol_ns_mirror0001`.

The aggregates named `aggr3` and `vs_aggr` contain data constituents, but no namespace constituent or namespace mirror constituent.

4. View the aggregates to identify the nodes that contain them by using the `storage aggregate show` command.

Example

In this example, aggregates `aggr3` and `vs_aggr` have data constituents, but do not have a namespace constituent or a namespace mirror constituent, and `aggr3` is on `node3` and `vs_aggr` is on `node4`. As a result, `node3` and `node4` require namespace mirror constituents for SnapDiff to use.

```
cluster1::> storage aggregate show
Aggregate      Size    Available Used%  State  #Vols  Nodes  RAID Status
-----
aggr0          900MB  42.66MB  95%   online  1     node2  raid_dp,normal
aggr0_node_0  900MB  42.55MB  95%   online  1     node1  raid_dp,normal
aggr0_node3_0 900MB  42.66MB  95%   online  1     node3  raid_dp,normal
aggr0_node4_0 900MB  42.66MB  95%   online  1     node4  raid_dp,normal
aggr1          1.76GB 1.13GB   36%   online  2     node1  raid4,normal
aggr2          1.76GB 1.13GB   36%   online  2     node2  raid4,normal
aggr3          1.76GB 1.38GB   22%   online  1     node3  raid4,normal
vs_aggr        1.76GB 1.36GB   23%   online  2     node4  raid4,normal
8 entries were displayed.
```

You have a list of the nodes with data constituents that require namespace mirror constituents for SnapDiff to use for tape backup.

5. For each node with a data constituent that requires a namespace mirror constituent for SnapDiff to use, ensure that at least one aggregate contained by the node has enough available space for a namespace mirror constituent by using the `storage aggregate show` command.

The namespace mirror constituent will be the same size as the namespace constituent. In this example, `node3` and `node4` require a namespace mirror constituent, and the namespace mirror constituent will be 256 MB. As a result, `node3` and `node4` must contain aggregates with enough space for a namespace mirror constituent that is 256 MB in size.

Example

In the following example, aggr3 is 1.76 GB in size, and only 387.5 MB of its space is used; therefore, aggr3 has enough space for a namespace mirror constituent.

```
cluster1::> storage aggregate show -aggregate aggr3

Aggregate: aggr3
...
Size: 1.76GB
State: online
Used Size: 387.5MB
Number Of Volumes: 1
Volume Style: flex
```

6. Increase the aggregate space if required.

For information about how to increase aggregate space, see the *Clustered Data ONTAP Physical Storage Management Guide*.

Related concepts

[Space and node requirements for namespace mirror constituents](#) on page 87

Related information

[Documentation on the NetApp Support Site: support.netapp.com](#)

Enabling SnapDiff on an Infinite Volume

You must enable SnapDiff on an Infinite Volume to set up incremental tape backup.

Before you begin

- The Infinite Volume must span two or more nodes in the cluster.
- All nodes in the cluster for the Infinite Volume must be running Data ONTAP 8.2 or later.
- The `filesystem-size-fixed` parameter must be set to `false` for read/write Infinite Volumes to enable SnapDiff.
- You must have written permission from NetApp to enable SnapDiff.

About this task

You can have cluster administrator or Vserver administrator privileges to perform this task. However, if you receive a message about insufficient aggregate space, you require cluster administrator privileges to add and modify aggregates. With cluster administrator and advanced privileges, you can identify the list of aggregates to use for the namespace mirror constituents by using the `volume modify` command with the `-ns-mirror-aggr-list` parameter.

It is best to enable SnapDiff when you create an Infinite Volume. However, you can enable SnapDiff when you create or modify an Infinite Volume. When you modify an Infinite Volume to enable or disable SnapDiff, you cannot modify any other parameters at the same time. For example, you cannot

increase the size of the Infinite Volume at the same time as you enable or disable SnapDiff. After you enable or disable SnapDiff, you can modify other parameters on the Infinite Volume.

Step

1. Enable SnapDiff on the Infinite Volume by using the `volume modify` command with the `-enable-snapdiff` parameter.

Example

In the following example, SnapDiff is enabled for an Infinite Volume named `repo_vol` that spans four nodes:

```
cluster1::> volume modify repo_vol -enable-snapdiff true
[Job 39] Job succeeded: Modified Infinite Volume successfully.
```

Before enabling SnapDiff, the Infinite Volume contains one namespace mirror constituent named `repo_vol_ns_mirror0001`:

```
cluster1::> volume show -is-constituent true
Vserver Volume                Aggregate State   Type  Size  Available  Used%
-----
vs0    repo_vol_1024_data0001 aggr1    online  RW   384MB  364.7MB  5%
vs0    repo_vol_1024_data0002 aggr2    online  RW   384MB  364.7MB  5%
vs0    repo_vol_1024_data0003 aggr3    online  RW   384MB  364.7MB  5%
vs0    repo_vol_1024_data0004 vs_aggr  online  RW   384MB  364.7MB  5%
vs0    repo_vol_ns            aggr1    online  RW   256MB  243.0MB  5%
vs0    repo_vol_ns_mirror0001 aggr2    online  DP   256MB  243.1MB  5%
6 entries were displayed.
```

After enabling SnapDiff, the Infinite Volume contains three namespace mirror constituents named `repo_vol_ns_mirror0001`, `repo_vol_ns_mirror0002`, and `repo_vol_ns_mirror0003`:

```
cluster1::> volume show -is-constituent true
Vserver Volume                Aggregate State   Type  Size  Available  Used%
-----
vs0    repo_vol_1024_data0001 aggr1    online  RW   384MB  364.7MB  5%
vs0    repo_vol_1024_data0002 aggr2    online  RW   384MB  364.7MB  5%
vs0    repo_vol_1024_data0003 aggr3    online  RW   384MB  364.7MB  5%
vs0    repo_vol_1024_data0004 vs_aggr  online  RW   384MB  364.7MB  5%
vs0    repo_vol_ns            aggr1    online  RW   256MB  243.0MB  5%
vs0    repo_vol_ns_mirror0001 aggr2    online  DP   256MB  243.1MB  5%
vs0    repo_vol_ns_mirror0002 aggr3    online  DP   256MB  243.1MB  5%
vs0    repo_vol_ns_mirror0003 vs_aggr  online  DP   256MB  243.1MB  5%
8 entries were displayed.
```

Namespace mirror constituents are created for SnapDiff to use.

Related concepts

[How the file system size setting for an Infinite Volume affects SnapDiff](#) on page 88

[How namespace mirror constituents affect the state of an Infinite Volume](#) on page 85

Selecting a Snapshot copy of an Infinite Volume for SnapDiff to use

You must select a Snapshot copy of an Infinite Volume for SnapDiff to use for incremental tape backup. SnapDiff uses the Snapshot copy to help identify new, changed, or deleted files in an Infinite Volume.

About this task

You can use a Snapshot copy created on a schedule by Data ONTAP, or you can use a Snapshot copy created by a backup application. The Snapshot copy of the Infinite Volume must be in an online state for SnapDiff to use it.

Steps

1. Perform one of the following actions:

To...	Do this...
Use a Snapshot copy created by Data ONTAP on a regular schedule for an Infinite Volume	Select a Snapshot copy in an online state from the available Snapshot copies by running the <code>volume snapshot show</code> command.
Create a regularly scheduled Snapshot copy for SnapDiff to use	Create a Snapshot copy of the Infinite Volume by using the <code>volume snapshot create</code> command. Note: The Infinite Volume must be in an online state to create a Snapshot copy.
Use the backup application to initiate the creation of a Snapshot copy of the Infinite Volume	See the documentation provided with the backup application.

You have selected a Snapshot copy of an Infinite Volume to use with SnapDiff.

2. Configure SnapDiff to use the Snapshot copy.

For more information about SnapDiff configuration, see the documentation provided with the backup application.

Creating a baseline for incremental tape backup

With SnapDiff enabled, you can use the NFS or CIFS protocol to mount an Infinite Volume, and you can use a backup application to back up all the files from the mounted Infinite Volume to tape. The tape backup creates a baseline for SnapDiff to use.

Before you begin

- A backup application must be installed and running, and it must support tape backup of files from a volume that is mounted with a supported version of the NFS or CIFS protocol.
- The server running the backup application must be on the same network as the Vserver with Infinite Volume that contains the mounted Infinite Volume.

- SnapDiff API must be configured for the Infinite Volume.
- SnapDiff must be enabled on the Infinite Volume.
- A Snapshot copy for the SnapDiff API to use must exist on the Infinite Volume, and the SnapDiff API must be configured to use the Snapshot copy.

For more information about SnapDiff configuration, see the documentation provided with the backup application.

About this task

The backup application backs up data from the entire Infinite Volume to tape. You cannot configure the backup application to back up data from specific, individual data constituents to tape.

Steps

1. Use the NFS or CIFS protocol to mount the Infinite Volume.
2. Ensure that the server running the backup application can read and write data to the Infinite Volume.

If you are backing up a destination Infinite Volume, ensure that the server running the backup application can read data from the destination Infinite Volumes. The server running the backup application cannot write data because the active file system on a destination Infinite Volume is read-only.

3. Configure the backup application to back up the mounted Infinite Volume to tape.

For more information about backing up files to tape, see the documentation provided with the backup application.

Result

All of the files in the Infinite Volume are backed up to tape, and a baseline is created for SnapDiff to use.

After you finish

You can now use subsequent backups to back up only the new and changed files in the Infinite Volume.

Related tasks

[Incrementally backing up changed files from an Infinite Volume to tape](#) on page 97

Setting up incremental tape backup of destination Infinite Volumes

When an Infinite Volume is in a data protection mirror relationship, you can use a backup application and SnapDiff to back up to tape a Snapshot copy on the destination Infinite Volume rather than on the source Infinite Volume. This setup moves the processing workload to the destination volume.

Steps

1. [Transferring a Snapshot copy to the destination Infinite Volume](#) on page 95
2. [Creating a baseline for incremental tape backup](#) on page 96

Transferring a Snapshot copy to the destination Infinite Volume

You can create a Snapshot copy on the source Infinite Volume, and use the `snapmirror update` command to transfer the Snapshot copy to the destination Infinite Volume. Backup applications can use the Snapshot copy on the destination volume for incremental tape backup.

Before you begin

- The source and destination Infinite Volumes must be in a data protection mirror relationship, and the data protection mirror relationship must be initialized.
For more information about data protection mirror relationships for Infinite Volumes, see the *Clustered Data ONTAP Data Protection Guide*.
- The destination Infinite Volume must span two or more nodes.
- The destination Infinite Volume must have enough space for SnapDiff resources.
- SnapDiff must be enabled on the destination Infinite Volume.

Steps

1. On the source Infinite Volume, create a Snapshot copy to use for tape backup by using the `volume snapshot create` command.

Example

In the following example, a Snapshot copy named `tape_backup` is created for a source Infinite Volume named `repo_vol`:

```
cluster1::> volume snapshot create -volume repo_vol -snapshot tape_backup
```

2. On the destination cluster, manually update the destination Infinite Volume by using the `snapmirror update` command.

Alternately, you can wait for the regularly scheduled data protection mirror update to occur to transfer the Snapshot copy to the destination Infinite Volume.

Example

In the following example, the data protection mirror relationship is updated for a destination Infinite Volume named `repo_vol_dest` on a destination Vserver with Infinite Volume named `vs0_dest`:

```
cluster2::> snapmirror update -destination-path vs0_dest:repo_vol_dest
```

The latest data protection mirror copy of the source Infinite Volume is transferred to the destination Infinite Volume, and the transfer includes the latest Snapshot copy that you created.

Related tasks

[Verifying space requirements in an Infinite Volume for SnapDiff resources \(cluster administrators only\)](#) on page 89

[Enabling SnapDiff on an Infinite Volume](#) on page 91

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

Creating a baseline for incremental tape backup

With SnapDiff enabled, you can use the NFS or CIFS protocol to mount an Infinite Volume, and you can use a backup application to back up all the files from the mounted Infinite Volume to tape. The tape backup creates a baseline for SnapDiff to use.

Before you begin

- A backup application must be installed and running, and it must support tape backup of files from a volume that is mounted with a supported version of the NFS or CIFS protocol.
- The server running the backup application must be on the same network as the Vserver with Infinite Volume that contains the mounted Infinite Volume.
- SnapDiff API must be configured for the Infinite Volume.
- SnapDiff must be enabled on the Infinite Volume.
- A Snapshot copy for the SnapDiff API to use must exist on the Infinite Volume, and the SnapDiff API must be configured to use the Snapshot copy.
For more information about SnapDiff configuration, see the documentation provided with the backup application.

About this task

The backup application backs up data from the entire Infinite Volume to tape. You cannot configure the backup application to back up data from specific, individual data constituents to tape.

Steps

1. Use the NFS or CIFS protocol to mount the Infinite Volume.

2. Ensure that the server running the backup application can read and write data to the Infinite Volume.

If you are backing up a destination Infinite Volume, ensure that the server running the backup application can read data from the destination Infinite Volumes. The server running the backup application cannot write data because the active file system on a destination Infinite Volume is read-only.

3. Configure the backup application to back up the mounted Infinite Volume to tape.

For more information about backing up files to tape, see the documentation provided with the backup application.

Result

All of the files in the Infinite Volume are backed up to tape, and a baseline is created for SnapDiff to use.

After you finish

You can now use subsequent backups to back up only the new and changed files in the Infinite Volume.

Incrementally backing up changed files from an Infinite Volume to tape

Backup applications can use SnapDiff to identify new, modified, and deleted files in an Infinite Volume since the last tape backup, and then back up to tape only the changed files instead of all of the files in the Infinite Volume.

Before you begin

- You must have set up incremental tape backup to create a baseline against which SnapDiff can identify changed files in the Infinite Volume since the last tape backup.

Steps

1. Use a backup application to request from SnapDiff a list of changed files in the Infinite Volume since the last tape backup.
2. Use a backup application to back up the changed files to tape.

Related tasks

[Setting up incremental tape backup of Infinite Volumes](#) on page 88

[Setting up incremental tape backup of destination Infinite Volumes](#) on page 95

Restoring files from tape for an Infinite Volume

You can use the NFS or CIFS protocol to mount the Infinite Volume and restore files from a tape backup to the mounted Infinite Volume.

Before you begin

- A backup application must be installed and running, and it must support the retrieval of files from a tape backup to a volume that is mounted with a supported version of the NFS or CIFS protocol.
- The server running the backup application must be on the same network as the Vserver with Infinite Volume that contains the mounted Infinite Volume.
- Some but not all versions of NFS and CIFS are supported for Infinite Volumes.

For more information about the protocols supported for Infinite Volumes, see the *Clustered Data ONTAP File Access and Protocols Management Guide*.

About this task

When an Infinite Volume is in data protection mirror relationship, you can only restore files to the source Infinite Volume because the destination Infinite Volume is read-only. If you want to restore files to the destination Infinite Volume, you must first break the data protection mirror relationship to change the destination Infinite Volume from read-only to read/write. For more information about data protection mirror relationships, see the *Clustered Data ONTAP Data Protection Guide*.

Steps

1. Use the NFS or CIFS protocol to mount the Infinite Volume.
2. Ensure that the server running the backup application can read and write data to the Infinite Volume.
3. Restore files from the tape backup to the mounted Infinite Volume.

For more information about restoring files, see the documentation provided with the backup application.

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

Reenabling SnapDiff after expanding a destination Infinite Volume onto new nodes

If you are using SnapDiff for incremental tape backup of a Snapshot copy on a destination Infinite Volume, you must manually reenable SnapDiff after you expand the destination volume onto a new

node to automatically create namespace mirror constituents on the new node. Otherwise, tape backup operations might fail.

Before you begin

- The Infinite Volume must not contain storage classes.
If you want to expand an Infinite Volume with storage classes, you cannot use the command-line interface; you must use OnCommand Workflow Automation instead. When an Infinite Volume with storage classes is in a data protection mirror relationship, the expansion workflow provided by OnCommand Workflow Automation reenables SnapDiff if necessary.
- SnapDiff must be enabled on the destination Infinite Volume.
- You must have increased the size of the destination Infinite Volume to expand the volume onto one or more new nodes.
- You must have updated the data protection mirror relationship to transfer new data constituents from the source Infinite Volume to the new nodes for the destination Infinite Volume.

About this task

When Infinite Volumes are in a data protection mirror relationship, you can use SnapDiff for incremental tape backup of a Snapshot copy on the destination Infinite Volume. When you increase the size of the source and destination Infinite Volumes in a data protection mirror relationship to expand onto new nodes, you must also manually reenables SnapDiff on the destination Infinite Volume to create namespace mirror constituents on the new nodes for SnapDiff to use.

Step

1. On the destination Infinite Volume, reenables SnapDiff by using the `volume modify` command with the `-enable-snapdiff` parameter.

Example

In the following example, SnapDiff is reenables for the Infinite Volume `repo_vol`:

```
cluster2::> volume modify repo_vol_dest -enable-snapdiff true
```

Namespace mirror constituents are automatically created on the new nodes for SnapDiff to use for incremental tape backup of Infinite Volumes.

Disabling SnapDiff for an Infinite Volume

You can disable SnapDiff on an Infinite Volume to disable incremental tape backup. With SnapDiff disabled, the namespace mirror constituents that are created for SnapDiff are deleted, and backup applications no longer receive a list of files that were changed since the last tape backup.

About this task

When you disable SnapDiff, all namespace mirror constituents, except one namespace mirror constituent, are automatically deleted from the nodes in the cluster. One namespace mirror constituent is retained to provide data protection for the namespace constituent.

Step

1. Disable SnapDiff on the Infinite Volume by using the `volume modify` command with the `-enable-snapdiff` parameter.

Example

In the following example, SnapDiff is disabled for the Infinite Volume named `repo_vol`:

```
cluster1::> volume modify repo_vol -enable-snapdiff false
[Job 50] Job succeeded: Modified Infinite Volume successfully.
```

Before disabling SnapDiff, the Infinite Volume spans four nodes and contains three namespace mirror constituents named `repo_vol_ns_mirror0001`, `repo_vol_ns_mirror0002`, and `repo_vol_ns_mirror0003`:

```
cluster1::> volume show -is-constituent true
Vserver Volume                Aggregate State  Type Size  Available Used%
-----
vs0  repo_vol_1024_data0001 aggr1  online RW   384MB 364.7MB 5%
vs0  repo_vol_1024_data0002 aggr2  online RW   384MB 364.7MB 5%
vs0  repo_vol_1024_data0003 aggr3  online RW   384MB 364.7MB 5%
vs0  repo_vol_1024_data0004 vs_aggr online RW   384MB 364.7MB 5%
vs0  repo_vol_ns            aggr1  online RW   256MB 243.0MB 5%
vs0  repo_vol_ns_mirror0001 aggr2  online DP   256MB 243.1MB 5%
vs0  repo_vol_ns_mirror0002 aggr3  online DP   256MB 243.1MB 5%
vs0  repo_vol_ns_mirror0003 vs_aggr online DP   256MB 243.1MB 5%
8 entries were displayed.
```

After disabling SnapDiff, the Infinite Volume spans four nodes and contains one namespace mirror constituent named `repo_vol_ns_mirror0001`, which is used for data protection of the namespace constituent named `repo_vol_ns`:

```
cluster1::> volume show -is-constituent true
Vserver Volume                Aggregate State  Type Size  Available Used%
-----
vs0  repo_vol_1024_data0001 aggr1  online RW   384MB 364.7MB 5%
```

```

vs0      repo_vol_1024_data0002 aggr2      online  RW   384MB 364.7MB 5%
vs0      repo_vol_1024_data0003 aggr3      online  RW   384MB 364.7MB 5%
vs0      repo_vol_1024_data0004 vs_aggr    online  RW   384MB 364.7MB 5%
vs0      repo_vol_ns              aggr1      online  RW   256MB 243.0MB 5%
vs0      repo_vol_ns_mirror0001 aggr2      online  DP   256MB 243.1MB 5%
6 entries were displayed.

```

All namespace mirror constituents are deleted, except the namespace mirror constituent used for data protection of the namespace constituent.

Related concepts

How the file system size setting for an Infinite Volume affects SnapDiff on page 88

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**
- affinity information
 - about [54](#)
 - aggregates
 - selection, for namespace mirror constituents [86](#)
- B**
- backup
 - interoperability between Data ONTAP 7-Mode and clustered Data ONTAP [65](#)
 - See also* tape backup
 - backup applications
 - guidelines for using network interfaces for SnapDiff [82](#)
 - selecting a Snapshot copy for SnapDiff [93](#)
 - SnapDiff support for [82](#)
 - supported tape backup topology for mounted volumes [81](#)
 - blocking factor
 - about [62](#)
- C**
- CAB
 - about [53](#)
 - NDMP v4 protocol extension [53](#)
 - challenge
 - supported NDMP authentication method [40](#)
 - CIFS
 - considerations for tape backup and restore of Infinite Volumes [80](#)
 - using to back up Infinite Volumes to tape [80](#)
 - using to restore files to Infinite Volumes [98](#)
 - Cluster Aware Backup extension
 - See* CAB
 - collocation of volumes and tapes
 - detecting [54](#)
- D**
- data protection mirror copies
 - namespace constituent in Infinite Volumes, of [83](#)
 - data restore
 - using dump engine for [63](#)
 - destination Infinite Volumes
 - transferring Snapshot copies to [95](#)
 - different LIF types
 - volumes and tape devices available on [53](#)
 - DMAs
 - incremental tape backups of Infinite Volumes, using Snapshot copies and SnapDiff [81](#)
 - dump
 - about [59](#)
 - backing up directories using [59](#)
 - backing up files using [59](#)
 - dump and restore events
 - viewing log messages for [69](#)
 - dump and volume movement operations
 - interoperability of [67](#)
 - dump backup
 - how it works [59](#)
 - dump backup and restore sessions
 - scalability limits for [65](#)
 - dump backup from vault destination volume
 - about [66](#)
 - dump engine
 - See* dump
 - dump engine and SFO/ARL operations
 - interoperability of [66](#)
 - dump error messages
 - destination qtree is read-only [78](#)
 - destination volume is read-only [78](#)
 - dumps temporarily disabled on volume, try again [78](#)
 - no files were created [78](#)
 - unable to locate bitmap files [79](#)
 - unable to lock a snapshot needed by dump [78](#)
 - volume is temporarily in a transitional state [79](#)
 - dump events
 - about [69](#)
 - dump restores
 - about [62](#)
- E**
- enhanced DAR functionality
 - about [40](#)
 - environment variables
 - about [29](#)
 - uses [29](#)

event logging
enabling or disabling [71](#)

F

files and data backup
using dump engine for [60](#)
format of dump and restore event log messages
about [69](#)

I

increment chains
understanding [61](#)
incremental tape backup
disabling for Infinite Volumes [100](#)
enabling SnapDiff [91](#)
for Infinite Volumes [97](#)
setup, destination Infinite Volumes [95](#)
setup, Infinite Volumes [88](#)
SnapDiff support
of backup application [82](#)
transferring Snapshot copies to destination Infinite
Volumes [95](#)

incremental tape backups
of Infinite Volumes [81](#)
tape backup topology for [81](#)

Infinite Volumes
backing up files incrementally to tape [97](#)
backing up over NAS protocols to tape [93](#), [96](#)
backup applications supported by SnapDiff [82](#)
CIFS and NFS considerations, tape backup of [80](#)
deleting namespace mirror constituents used by
SnapDiff [100](#)
disabling incremental tape backups [100](#)
disabling SnapDiff [100](#)
effect of namespace mirror constituents on the state
of [85](#)
enabling SnapDiff [91](#)
how namespace mirror constituents provide data
protection [83](#)
how namespace mirror constituents support tape
backup [83](#)
how SnapDiff supports [81](#)
how the file system size affects SnapDiff [88](#)
namespace mirror constituents, definition of [83](#)
reenabling SnapDiff [98](#)
restoring files to [98](#)
selecting a Snapshot copy for SnapDiff [93](#)
setting up incremental tape backup [88](#)

setting up incremental tape backup for destination
Infinite Volumes [95](#)
Snapshot copies you can back up to tape [82](#)
space required for namespace mirror constituents [87](#)
support for incremental tape backup [81](#)
support for tape backup [8](#)
supported tape backup topology [81](#)
transferring Snapshot copies to a destination Infinite
Volumes [95](#)
verifying if they have sufficient space for SnapDiff
[89](#)
workflow for tape backup and restore [9](#)

Infinite Volumes with storage classes
tape backup, impact of storage classes on [81](#)

L

levels of incremental backup
specifying [61](#)
LIF types for NDMP data connection
controlling [55](#)

LIFs
guidelines for using with SnapDiff [82](#)

M

managing tape backup and restore operations
using environment variables for [29](#)
multipath tape access
about [19](#)
considerations for [19](#)
understanding [19](#)

N

namespace mirror constituents
definition of [83](#)
deleting by disabling SnapDiff [100](#)
effect of modifying the Infinite Volume [84](#)
effect on Infinite Volume state [85](#)
how aggregates are selected for [86](#)
space requirements for [87](#)
verifying if Infinite Volumes have enough space for
[89](#)

NAS protocols
using to back up Infinite Volumes to tape [93](#), [96](#)

NDMP
about [27](#)
considerations [28](#)
data connection types

- firewall policy [28](#)
 - Infinite Volumes, unsupported [8](#)
 - `ndmptcopy` command [28](#)
 - tape backup topologies
 - Storage system-to-data server-to-tape [39](#)
 - Storage system-to-local-tape [39](#)
 - Storage system-to-network attached tape library [39](#)
 - Storage system-to-tape attached to another storage system [39](#)
 - NDMP authentication methods
 - specifying [40](#)
 - NDMP control connections
 - about [53](#)
 - NDMP data connection type
 - determining [56](#)
 - NDMP error messages
 - Data connection type
 - ["NDMP4_ADDR_TCP"|"NDMP4_ADDR_TCP_IPv6"] plaintext not supported for ["IPv6"|"IPv4"] control connections [76](#)
 - could not obtain vol ref for Volume *volume_name* [76](#)
 - DATA CONNECT: CAB_DATA_CONN_PREPARE not sent or addr type mismatch [77](#)
 - DATA LISTEN: CAB_DATA_CONN_PREPARE not sent or addr type mismatch [77](#)
 - Error:show failed: Cannot get password for user '<username>' [77](#)
 - message from Read Socket: error_string [75](#)
 - message from Write Dirnet: error_string [75](#)
 - ndmpd invalid version number: *version_number* [76](#)
 - ndmpd session *session_ID* not active. [76](#)
 - network communication error [75](#)
 - read Socket received EOF [76](#)
 - NDMP modes of operation
 - understanding [27](#)
 - NDMP sessions
 - scalability limits for [41](#)
 - NDMP user in node-scoped mode
 - authenticating [46](#)
 - NDMP-specific password
 - generating [57](#)
 - network interfaces
 - guidelines for using with SnapDiff [82](#)
 - NFS
 - protocols for tape backup and restore of Infinite Volumes [80](#)
 - using to restore files to Infinite Volumes [98](#)
 - using to tape back up Infinite Volumes [80](#)
 - node-scoped NDMP mode
 - about [27](#), [42](#)
 - commands for managing [42](#)
 - managing [42](#)
 - using NDMP options in [43](#)
 - node-scoped NDMP mode options
 - managing NDMP by [43](#)
 - nonqualified tape drives
 - using [22](#)
- ## O
- options
 - backup.log.enable (turns event logging on or off) [71](#)
- ## P
- physical path names (PPNs)
 - about [17](#)
 - preferred-interface-role option
 - about [55](#)
- ## Q
- qualified tape drives
 - about [11](#)
- ## R
- restore
 - interoperability between Data ONTAP 7-Mode and clustered Data ONTAP [65](#)
 - restore command
 - disk space required for [64](#)
 - information required for using [64](#)
 - restore events
 - about [70](#)
 - restoring
 - files to an Infinite Volume [98](#)
- ## S
- serial numbers
 - about [18](#)
 - SFO and ARL operations
 - performing tape operations during [66](#)
 - simultaneous backup or restore sessions
 - supported number of [16](#)
 - SnapDiff

- aggregate selection for namespace mirror constituents [86](#)
- backup applications
 - requirements to use [82](#)
- defined [81](#)
- disabling for Infinite Volumes [100](#)
- enabling for Infinite Volumes [91](#)
- guidelines for using network interfaces [82](#)
- how it works with namespace mirror constituents [81](#)
- impact of the volume modify command on [88](#)
- incremental tape backups of Infinite Volumes, using [81](#)
- namespace mirror constituents creation for tape backup of Infinite Volumes [83](#)
- namespace mirror constituents for [83](#)
- nodes
 - number required for namespace mirror constituents [87](#)
- reenabling it after expanding destination Infinite Volume [98](#)
- required for incremental tape backup of Infinite Volumes [80](#)
- selecting a Snapshot copy for [93](#)
- space requirements for namespace mirror constituents [87](#)
- verifying space requirements in Infinite Volumes [89](#)
- workflow for tape backup and restore of Infinite Volumes [9](#)
- Snapshot copies
 - available for tape backup [82](#)
 - incremental tape backups of Infinite Volumes, using [81](#)
 - selecting for SnapDiff to use [93](#)
 - used by SnapDiff to identify file and directory differences [81](#)
- storage (aliasing) commands
 - storage unalias (removes tape alias) [24](#)
- storage classes
 - tape backup of Infinite Volumes, impact on [81](#)
- storage systems
 - adding Fiber Channel-attached drives dynamically to [19](#)
 - dynamically adding tape drives and libraries to [19](#)
- supported NDMP extensions
 - about [53](#)

T

- tape aliases
 - assigning [23](#)

- definition [17](#)
- using serial numbers for [18](#)
- tape backup
 - (incremental), setting up for destination Infinite Volumes [95](#)
 - (incremental), transferring a Snapshot copy for [95](#)
 - creating baseline for incremental [93](#), [96](#)
 - how namespace mirror constituents support SnapDiff for [83](#)
 - Infinite Volumes with storage classes, impact on [81](#)
 - of Infinite Volume, unsupported features [80](#)
 - of Infinite Volume, using CIFS or NFS [80](#)
 - selecting a Snapshot copy of Infinite Volumes for [93](#)
 - setting up incremental [88](#)
 - Snapshot copies available for [82](#)
 - using NDMP [27](#)
- tape backup and recovery
 - Infinite Volumes, support for [8](#)
 - using NDMP [27](#)
 - using NDMP for [8](#)
 - using the dump engine [59](#)
- tape backup and restore
 - Infinite Volumes, workflow for [9](#)
- tape backup and restore error messages
 - already at the end of tape [74](#)
 - could not initialize media [72](#)
 - media error on tape read [74](#)
 - media error on tape write [73](#)
 - resource limitation: no available thread [72](#)
 - tape read error [74](#)
 - tape record size is too small [74](#)
 - tape record size must be in the range between 4KB and 256KB [75](#)
 - tape record size should be block_size1 and not block_size2 [75](#)
 - tape reservation preempted [72](#)
 - tape write error [74](#)
 - tape write failed [73](#)
 - tape write failed - new tape encountered media error [73](#)
 - tape write failed - new tape is already at the end of media [73](#)
 - tape write failed - new tape is broken or write protected [73](#)
 - too many active dumps/restores currently in progress [73](#)
- tape backup and restore event log files
 - accessing [68](#)
- tape backup and restore operations
 - accessing the event log files for [68](#)

- monitoring [68](#)
- performing per Vserver basis [47](#)
- workflow for [8](#)
- tape configuration files
 - accessing [11](#)
 - format of [12](#)
- tape device name
 - format [15](#)
- tape devices
 - about [14](#)
- tape drive connections
 - supported number of [16](#)
- tape drive information
 - viewing [21](#)
- tape drives
 - managing [21](#)
 - supported topology for Infinite Volumes [81](#)
 - understanding [11](#)
- tape drives and libraries to storage systems
 - dynamically adding [19](#)
- tape drives dynamically
 - qualifying [14](#)
- tape drives to storage systems
 - dynamically adding [19](#)
- tape libraries
 - supported topology for Infinite Volumes [81](#)
- tape libraries to storage systems
 - dynamically adding [19](#)
- tape library connections

- verifying [25](#)
- tape library details
 - viewing [25](#)
- tape reservations
 - what are [20](#)

V

- vault destination volume
 - backing up data from [66](#)
- volumes and tape devices for backup and restore operations
 - determining availability of [53](#)
- Vserver aware NDMP mode
 - about [27, 47](#)
 - authenticating user in [57](#)
 - commands for managing [48](#)
 - generating passwords [57](#)
 - using NDMP options in [49](#)
- Vserver aware NDMP mode options
 - managing NDMP by using [49](#)
- Vserver level
 - performing backup and restore operations at [47](#)
- Vservers with Infinite Volume
 - guidelines for using network interfaces for SnapDiff [82](#)