

NetBackup™ NAS Administrator's Guide

Release 11.2.0.1

NetBackup™ NAS Administrator's Guide

Last updated: 2026-09-07

Legal Notice

Copyright © 2026 Cohesity, Inc. All rights reserved.

© 2026 Cohesity, Inc. All Rights Reserved. Cohesity, the Cohesity Logo and other Cohesity Marks are trademarks of Cohesity, Inc. in the US and/or internationally. The information supplied herein is the confidential and proprietary information of Cohesity and may only be used (a) by the intended recipients and (b) in conjunction with validly licensed Cohesity software and services. Find the terms of Cohesity licenses at www.cohesity.com/agreements.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. COHESITY SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to [Cohesity Support](#), to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the [Cohesity Support Portal](#) to create a new case.
- Click the (?) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the [Cohesity Knowledge Base](#).
- Log in to the [Cohesity Support Portal](#) to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Contents

Section 1	About NAS backups	12
Chapter 1	Introduction	13
	About NAS backups	13
	Backups using the NAS-Data-Protection policy	13
	Backups using NDMP policy	14
	Terminology	14
Section 2	Using NAS-Data-Protection (D-NAS)	
		20
Chapter 2	D-NAS overview	21
	Dynamic data streaming for D-NAS Policy	21
	Understanding the features of D-NAS	22
	Dynamic streaming parameter	23
	Dynamic backup host pool	24
	About the All media server pool option	24
	Limitations and considerations	25
Chapter 3	D-NAS Planning and Tuning	27
	Sizing guidelines for D-NAS	27
	Tuning parameters for Server Message Block (SMB)	29
	Tuning parameters for Network File System (NFS)	30
	NetBackup tuning parameters for the backup host	30
Chapter 4	Pre-requisites for D-NAS configuration	32
	Prerequisites for D-NAS configuration	32
	Required firewall ports	33
	Domain user requirement for SMB share backups	33
	Minimum permissions for SMB share backup and recovery	34
	Minimum supported backup host versions for different features	34
	Configuring a backup host pool	35

Chapter 5	Configuring Storage Lifecycle Policies for D-NAS	37
		37
	About storage lifecycle policies	37
	Snapshot operation in SLP	39
	Primary snapshot storage unit	40
	Primary + Replication source snapshot storage unit	41
	Primary + Replication source + Replication target storage unit	42
	Replication target snapshot storage unit	42
	Creating a storage lifecycle policy for snapshots and snapshot replication	43
	Replication operation in the SLP	44
	Index from snapshot operation in an SLP	45
	Determining where and when the Index from snapshot operation occurs	47
	Backup from snapshot operation in an SLP	47
	About retention period for backup from snapshot images	49
	Duplication operation in an SLP	49
	Retention types for SLP operations	50
	Expire after copy retention type for SLP operations	51
	Fixed retention type for SLP operations	52
	Maximum snapshot limit retention type for SLP operations	52
Chapter 6	Multi-host backup for volumes	54
	About multi-host backup	54
	Stream distribution across multiple backup hosts	54
	Considerations for multi-host backups	55
	Monitoring and troubleshooting backup jobs	55
Chapter 7	Configure D-NAS policy for NAS backups	57
	About policies for NAS backups	58
	Planning for policies	58
	Prerequisites for D-NAS policies	59
	Configure D-NAS policy for NAS volumes	60
	Policy attributes	61
	Creating schedule attributes for policies	65
	Configuring the Start window	68
	Adding, changing, or deleting a time window in a policy schedule	68
	Example of schedule duration	69
	Configuring the exclude dates	70

	Configuring clients	71
	Configuring backup selections	72
	Configuring exclude lists	73
	Ordering of backup from snapshot jobs	73
	About mixed mode volumes	74
	Configuring include and exclude lists	74
	Auto-resume backup for incomplete backup jobs	77
Chapter 8	Using Accelerator	78
	Accelerator for D-NAS	78
	About the track logs for Accelerator	79
	Track log sizing considerations	80
	Notes on accelerator for D-NAS	80
	Accelerator forced rescan option	80
Chapter 9	Using Vendor Change Tracking	82
	About Vendor Change Tracking	82
	About NetApp SnapDiff support	83
	Using VCT with accelerator for D-NAS	83
	Using VCT for indexing	84
	Changing the number of backup streams when VCT and accelerator are enabled	85
	Index from snapshot for D-NAS	85
	Using VCT with NetBackup client exclude list	87
Chapter 10	Using true image restore	88
	About true image restore	88
	Configuring TIR information retention time	90
	Considerations for using TIR	90
Chapter 11	Replication using D-NAS policy	91
	Replication using D-NAS policy	91
Chapter 12	Restoring from D-NAS backups	93
	Considerations for restoring from D-NAS backups	93
	About the Overwrite existing file option during restore	94
	Multi-stream restores from D-NAS backups	95
	RBAC role for D-NAS restores	95
	Scanning for malware	95
	Restore everything to a different location	95

Restore individual files and folders to different locations	97
Original location restores for D-NAS Policy	98
Restore Azure Files backups to the original SMB volume	98
Point-in-time rollback	100
Chapter 13 Troubleshooting	102
Troubleshooting	103
Setting the log level	103
Logging directories for Linux platforms	103
Logging folders for Windows platforms	106
Logging folders for multi-stream restore	109
Exclude list is not working during backup	110
Restore from a snapshot fails with status 133	110
Backup from snapshot jobs do not start after the snapshot job completes successfully	110
Backup from snapshot fails with error 50	111
Backup from snapshot parent job fails with error 4213: Snapshot import failed	111
Backup host pool creation fails with the error "Failed to fetch host list"	112
Snapshot job fails and the snapshot command does not recognize the volume name	112
Accelerator enabled incremental backup of NetApp NAS volume	113
Snapshot method: Auto	113
Backup from snapshot jobs for NAS-Data-Protection policy fail with error 4213	114
A full VCT-enabled indexing job runs, when followed by a non-VCT indexing job with a backup host prior to version to 10.3	114
Backup from snapshot jobs for NAS data protection policy fail with error 927	115
Error code: 930: No supported media server is available in the All_Media_Server_Pool to use to backup the NAS shares.	115
Restore from NAS array volume fails with the status: 174 Media manager – system error occurred.	116
NAS job fails with the error: Crawler process timed out after 600 seconds waiting for streams to attach with shared memory.	116
D-NAS backup fails with the error: The file system crawler process timed-out waiting for streams to attach with shared memory. (3003)	117
Isilon backup from snapshot failed with the Snapshot cannot be mounted error.	117

	Discovery and snapshot operations fail with the errors 156 and 1542	118
Section 3	Using NDMP	119
Chapter 14	Introduction to NetBackup for NDMP	120
	About NetBackup for NDMP	121
	NetBackup for NDMP features	121
	NetBackup for NDMP terminology	123
	About Network Data Management Protocol (NDMP)	126
	Types of NDMP backup	126
	NDMP local backup	127
	NDMP three-way backup	127
	Backup to Media Manager storage units (remote NDMP)	128
	About NDMP policies in NetBackup	129
	About NetBackup storage units	130
	About assigning tape drives to different hosts	130
	About robotics control	131
	About the NDMP backup process	132
	About the NDMP restore process	134
	About Direct Access Recovery (DAR)	136
	Snapshot Client assistance	137
	About NDMP multiplexing	137
	About NDMP support for Replication Director	138
	Limitations of Replication Director with NDMP	138
	About NDMP support for NetApp clustered Data ONTAP (cDOT)	139
Chapter 15	Installation Notes for NetBackup for NDMP	142
	NetBackup for NDMP installation prerequisites	142
	Adding the NetBackup for NDMP license	143
	About existing NetApp cDOT configurations before you upgrade	144
Chapter 16	Configuring NDMP backup to NDMP-attached devices	149
	About configuring NDMP-attached devices	149
	Authorizing NetBackup access to a NAS (NDMP) host	150
	About access for three-way backups and remote NDMP	151
	About Media and Device Management configuration	152
	Adding a robot or tape drive directly attached to an NDMP host	152
	Checking the device configuration	153

	Configure an NDMP filer	153
	About verifying NDMP password and robot connection	154
	About adding volumes	154
	Adding NDMP storage units	155
	About creating an NDMP policy	156
	About appropriate host selection for NetApp cDOT backup policies	156
	Attributes tab options for an NDMP policy	157
	Schedules tab options for an NDMP policy with Accelerator for NDMP enabled	158
	About backup types in a schedule for an NDMP policy	158
	Clients tab options for an NDMP policy	159
	Backup selection options for an NDMP policy	159
	About enabling or disabling DAR	167
	Disabling DAR for file and directory restores	168
	Disabling DAR for directory restores only	168
	Configuring NetBackup for NDMP in a clustered environment	169
Chapter 17	Configuring NDMP backup to NetBackup media servers (remote NDMP)	171
	About remote NDMP	171
	Configuring NDMP backup to Media Manager storage units	172
Chapter 18	Configuring NDMP DirectCopy	174
	About NDMP DirectCopy	174
	Prerequisites for using NDMP DirectCopy	175
	NDMP DirectCopy with VTL	175
	NDMP DirectCopy without VTL	177
	Configuring NDMP DirectCopy	178
	Using NDMP DirectCopy to duplicate a backup image	179
	Requirements to use NDMP DirectCopy for image duplication	179
	Initiating NDMP DirectCopy with the NetBackup web UI	180
Chapter 19	Accelerator for NDMP	181
	About NetBackup Accelerator for NDMP	181
	About the track log for Accelerator for NDMP	184
	How to redirect track logs for Accelerator for NDMP	185
	Accelerator messages in the NDMP backup job details log	187
	NetBackup logs for Accelerator for NDMP	190

Chapter 20	Remote NDMP and disk devices	192
	About remote NDMP and disk devices	192
	Configuring remote NDMP	193
Chapter 21	Using the Shared Storage Option (SSO) with NetBackup for NDMP	195
	About the Shared Storage Option (SSO) with NetBackup for NDMP	195
	Setting up SSO with NetBackup for NDMP	196
Chapter 22	NAS appliance information for NDMP	198
	About NAS appliances support	198
	Non-vendor-specific information	198
	Vendor-specific information	200
	Dell EMC Isilon	200
	Dell EMC VNX	201
	Dell EMC Unity	204
	EMC Celerra	206
	Hitachi HDI/VFP	209
	Hitachi NAS (HNAS)	209
	HP X9000 NAS	211
	Huawei OceanStor V3	212
	IBM System Storage Nxxxx	213
	NEC Storage NV series	215
	NetApp	217
	Nexenta	223
	Nexsan	224
	Oracle Axiom Series	225
	Oracle Solaris Server	226
	Stratus V Series	228
Chapter 23	Backup and restore procedures	229
	Performing a manual backup with an NDMP policy	229
	Perform an NDMP restore	230
Chapter 24	Troubleshooting	232
	About NetBackup for NDMP logs	232
	Viewing NetBackup for NDMP logs	232
	NDMP backup levels	234
	General NetBackup for NDMP operating notes and restrictions	235

NetBackup for NDMP troubleshooting suggestions	237
Troubleshooting NDMP media and devices on Windows	237
Troubleshooting NDMP media and devices on UNIX	238
Troubleshooting NDMP DirectCopy	238
Troubleshooting Direct Access Recovery (DAR) with NetBackup for NDMP	239
About robot tests	240
TLD robot test example for UNIX	240

Chapter 25	Using NetBackup for NDMP scripts	242
	About the NetBackup for NDMP scripts	242
	ndmp_start_notify script (UNIX)	243
	ndmp_start_notify.cmd script (Microsoft Windows)	245
	ndmp_end_notify script (UNIX)	247
	ndmp_end_notify.cmd script (Microsoft Windows)	249
	ndmp_start_path_notify script (UNIX)	251
	ndmp_start_path_notify.cmd script (Microsoft Windows)	254
	ndmp_end_path_notify script (UNIX)	256
	ndmp_end_path_notify.cmd script (Microsoft Windows)	258
	ndmp_moving_path_notify script (UNIX)	260
	ndmp_moving_path_notify.cmd script (Microsoft Windows)	262

About NAS backups

- [Chapter 1. Introduction](#)

Introduction

This chapter includes the following topics:

- [About NAS backups](#)
- [Backups using the NAS-Data-Protection policy](#)
- [Backups using NDMP policy](#)
- [Terminology](#)

About NAS backups

NetBackup Snapshot Manager and NDMP V4 snapshot extension can make snapshots of client data on a NAS host. A NAS snapshot is a point-in-time disk image. You can retain the Snapshots on the disk for any duration. Using the Instant Recovery feature in NetBackup, you can efficiently restore the data from the disk. In NetBackup, you can use snapshot-based data protection to protect your NAS data, using two policies: NAS-Data-Protection and NDMP.

Backups using the NAS-Data-Protection policy

NAS-Data-Protection policy is a robust approach to back up the data residing on NAS storage. It is also known as dynamic NAS or D-NAS policy. NetBackup Snapshot Manager and the storage array plug-ins can make snapshots of NAS volumes and shares. The dynamic data streams can access the snapshots on the backup hosts and read them to create point-in-time backup copies. For more details about D-NAS policy, see *Section 2* of this guide.

Backups using NDMP policy

NetBackup can make snapshots of client data on a NAS (NDMP) host using the NDMP V4 extension. The snapshot data is read over NDMP, and backup copies are created per configured target. For more details about NDMP policy, see *Section 3: Using NDMP* of this guide.

Terminology

The following table describes the concepts and terms in D-NAS data protection.

Table 1-1 D-NAS terminology

Term	Definition
Backup	The process of creating a copy of user data and creating backup images of the data. Can be any of the two: ■ The process of creating a new backup image of the client's data that is tar-formatted. ■ The process of creating a snapshot of the client's data.
Backup host	The backup host acts as a proxy client where the snapshot of the NAS share is staged for reading purpose. All the backup and restore operations are run through the backup host. You can configure NetBackup media servers, clients, or a primary server as a backup host. The backup host is also used as a destination client during restores.
Backup job	A backup job in D-NAS is a compound job. ■ The backup job runs a discovery job for getting information of the data to be backed up. ■ Child jobs are created for each backup host that performs the actual data transfer. ■ After the backup is complete, any temporary files or transient information is cleaned up and then job is marked complete.

Table 1-1 D-NAS terminology (*continued*)

Term	Definition
Child job	For backup, a separate child job is created for each backup host to transfer data to the storage media.
Copy	An instance of a NetBackup image which can be standalone; it can be read or deleted without affecting any other copy.
Data mover	The mechanism that is used to copy data from storage on the production client to backup storage. Or, to duplicate, the data mover copies data from backup storage to different backup storage. Traditionally, NetBackup functions as the data mover and data travels through clients and media servers. Storage devices can provide more efficient mechanisms to move the data, such as NDMP, built-in replication, or OST (as in Optimized Duplication).
Discovery of NAS shares	When a storage array plug-in is created, a discovery task starts on the Snapshot Manager for Data Center host. The discovery job communicates with the arrays and gathers information of storage array clusters, arrays, volumes, and shares. The discovery runs periodically every 4 hours in a day to refresh its asset information. NetBackup presents this asset information for user selection.
Disk array	A disk array which exposes storage or network shares to a host server over SAN, NAS, NFS, CIFS, or iSCSI protocols.
Dynamic streaming	NetBackup Dynamic streaming is a framework that engages multiple backup and restore streams to read data in a distributed manner and send them for backup storage or the restore location.

Table 1-1 D-NAS terminology (*continued*)

Term	Definition
Media server	Media servers provide additional storage by allowing NetBackup to use the storage devices that are attached to them. Media servers can also increase performance by distributing the network load. Media servers can also be referred to by using the following terms: ■ Device hosts, when tape devices are present. ■ Storage servers, when I/O is directly to disk. ■ Data movers, when data is sent to independent, external disk devices like OpenStorage appliances.
MSDP	Media Server Deduplication Storage Pool is a NetBackup deduplication technology engine to optimize backup storage.
NetBackup Accelerator	A backup technology that speeds up the backup process by reducing the amount of data sent to the media server. It can be used for full and incremental backups.
NetBackup certificate	A security certificate that is issued from the NetBackup CA.
NetBackup Replication	The process of copying and transferring backups created in one NetBackup domain to the storage of another NetBackup domain. This process creates a duplicate set of backups at a different location. Replication is typically used for disaster recovery purposes. This function is primarily known as Auto Image Replication (AIR) within NetBackup.

Table 1-1 D-NAS terminology (*continued*)

Term	Definition
NetBackup Snapshot Manager for Data Center	Undertakes on-premises storage array snapshot management and replication tasks. NetBackup Snapshot Manager for Data Center has plug-ins which integrate with REST APIs and SDK of storage array vendors for interaction with storage arrays. NetBackup also enables NetBackup Snapshot Manager for Data Center for snapshot management of Cloud offerings of Storage arrays, viz., NetApp CVO, and Azure Files.
Primary copy	The Primary copy or Copy 1 refers to the snapshot copy of the D-NAS backup job. The backup copies created from the primary snapshot copy are called Copy 2 or secondary copies.
Primary server	The primary server manages backups, archives, and restores. The primary server is responsible for media and device selection for NetBackup. Typically, the primary server contains the NetBackup catalog. The catalog contains the internal databases that contain information about NetBackup backups and configuration.
Primary volume	A unit of storage space that a disk array exposes to a host in the form of a network share (NFS or CIFS) or LUN block device. Primary volumes store an application's active data.
RBAC	Role-based access control. The role administrator can delegate or limit access to the NetBackup UI through the roles that are configured in RBAC.

Table 1-1 D-NAS terminology (*continued*)

Term	Definition
Role	For RBAC, defines the operations that a user can perform and the NAS shares that they can access. For example, you can configure a role to manage recovery of specific NAS shares and the credentials that are needed for backups and restores. 'Default NAS Administrator' is an RBAC role tailored for NAS administrators.
Replication job	A replication operation is specified in an SLP that was added to a D-NAS policy. Generates a replication parent-child job in the Activity monitor.
Snapshot	Refers to a point-in-time copy of the NAS volume or share on the storage arrays. An image copy that is a snapshot is also considered a replica. A snapshot copy consists of one or more snapshot fragments.
Snapshot job	A NetBackup job that creates a hardware snapshot for the NAS volume or share specified in the policy. NetBackup creates a parent-child job hierarchy, where each child job represents a NAS volume or share in the backup selection.
Storage lifecycle policy (SLP)	NetBackup uses SLPs to manage the lifecycle of a backup or snapshot image. An SLP controls image migration, duplication, and replication within a single NetBackup primary server domain.
Storage server	A storage device that is configured in NetBackup. A storage server is a NetBackup component that represents a disk array.

Table 1-1 D-NAS terminology (*continued*)

Term	Definition
Storage unit	A storage unit is configured for one of two types of data: ■ Backup storage units contain backup images. A backup storage unit cannot contain snapshots. ■ Snapshot storage units contain snapshots. A snapshot storage unit cannot contain backups. The replication process uses snapshot storage units in snapshot replication configurations.
Vendor Change Tracking (VCT)	Several NAS storage array vendors feature difference engines that identify the list of changed files and directories between two snapshot copies of the same volume. When VCT is enabled for a D-NAS policy, NetBackup does not perform any file system tracking for backup or index of NAS volumes. Instead, NetBackup relies solely on the change-list from the difference engine of the storage array to perform backup of files and directories. This process optimizes the backup process.

Using NAS-Data-Protection (D-NAS)

- [Chapter 2. D-NAS overview](#)
- [Chapter 3. D-NAS Planning and Tuning](#)
- [Chapter 4. Pre-requisites for D-NAS configuration](#)
- [Chapter 5. Configuring Storage Lifecycle Policies for D-NAS](#)
- [Chapter 6. Multi-host backup for volumes](#)
- [Chapter 7. Configure D-NAS policy for NAS backups](#)
- [Chapter 8. Using Accelerator](#)
- [Chapter 9. Using Vendor Change Tracking](#)
- [Chapter 10. Using true image restore](#)
- [Chapter 11. Replication using D-NAS policy](#)
- [Chapter 12. Restoring from D-NAS backups](#)
- [Chapter 13. Troubleshooting](#)

D-NAS overview

This chapter includes the following topics:

- [Dynamic data streaming for D-NAS Policy](#)
- [Understanding the features of D-NAS](#)
- [Dynamic streaming parameter](#)
- [Dynamic backup host pool](#)
- [Limitations and considerations](#)


Dynamic data streaming for D-NAS Policy

NetBackup can make snapshots of NAS volumes and shares using the Snapshot management server and the storage array plug-ins. The snapshots are accessed on backup hosts and read by dynamic streams to create point-in-time backup copies.

You can perform a snapshot-enabled, off-host backup of NAS volumes, where a volume is backed up using dynamic backup streams.

Each NAS volume or share is read over NFS or SMB, and backed up using a configured number of backup streams. Files from these NAS volumes or shares are optimally distributed in real-time across streams to realize the full potential of backup streams. You cannot mix NAS volumes or shares of different storage array vendors in a single policy. In other words, using a single policy, you can only protect assets of a single vendor, on a single NAS protocol.

Dynamic streaming is built on the NetBackup client framework and uses NAS-Data-Protection policy type for snapshot and backup orchestration of NAS data. This policy supports SLP only for the data lifecycle.



Understanding the features of D-NAS

This table explains the salient features of data protection using D-NAS.

Table 2-1

Feature	Description
Integration with NetBackup Role-based Access Control (RBAC)	NetBackup web UI provides the Default NAS Administrator RBAC role to control which NetBackup users can perform backup and restore of NAS volumes using NAS-Data-Protection policy. The user need not be a NetBackup administrator to perform these operations on NAS volumes using the NAS-Data-Protection policy.
Convenience of backup host pool	The backup host pool is a group of NetBackup backup hosts where the snapshot of the volume is staged for the backup process to read. These hosts can be NetBackup clients, media servers, or primary servers.

Table 2-1 (continued)

Feature	Description
Vendor change tracking	Vendor Change Tracking (VCT) is a mechanism to get the difference in the content of the volume or share between two points-in-time snapshots. See “About Vendor Change Tracking” on page 82.
Exclude volumes	You can exclude the volumes from the backup selection list that you do not want to back up. For example, if <code>/prodVol*</code> is the backup selection, there may be a volume <code>/prodVol-Scratch</code> , which you do not want to back up.
NetBackup accelerator	You can use NetBackup's robust accelerator feature and dynamic streaming for optimized and fast backups.
Checkpoint restart	You can leverage NetBackup's checkpoint restart feature along with dynamic streaming. By taking checkpoints periodically during the backup, NetBackup can retry a failed backup from the beginning of the last checkpoint without restarting the entire job.

Dynamic streaming parameter

Dynamic streaming is a group of backup streams running in parallel that dynamically distributes the files for backups among them. This feature optimizes and speeds up the backup of dense NAS volumes or shares.

Maximum number of streams per volume: The value determines the number of backup streams that are deployed for backing up each volume. For example, if a policy contains 10 volumes, and the value of this parameter is set to 4, then you see a group of four backup streams for each volume. So, a total of 40 child backup streams and 10 parent backup streams run during the backup.

You can set this parameter in the web UI. The default maximum value is 40. You can select a value between 1 to 40.

You can configure the maximum value this parameter to be available in the web UI, using the `MAX_NUMBER_OF_DYNAMIC_STREAMS` parameter in the `bp.conf` file. You can specify any numeric value greater than 0.

For example, if you assign the value 100 to this parameter in the `bp.conf` file, you can set values from 1 to 100 streams per volume in a D-NAS policy in the web UI.

Note the following about the dynamic streaming parameter:

- Setting the dynamic streaming parameter to a default value of 10 or lower, NetBackup forces the D-NAS policy to use only a single host. NetBackup does not consider multiple hosts, even if they are present in the backup host pool.
- For optimum utilization of resources, set the dynamic streaming parameter to more than 10, when combined with the multi-host feature. This way, the streams are distributed across multiple hosts, thus reducing the load on only one host.

Dynamic backup host pool

NetBackup version 10.4 or higher, can create a dynamic backup host pool. No need to manually add or remove backup hosts from the backup host pool. All media servers configured with the primary server constitute this backup host pool. Any new media server configured with the primary server automatically becomes part of this backup host pool.

About the All media server pool option

You can use the **All media server pool** option, while selecting a backup host pool for the D-NAS data protection policies. This option selects all the supported media servers that are configured on the primary server and available at run time.

Note: The All media server pool option does not include the primary servers that are also configured as a media server.

This option is present for both NFS and SMB protocols. You can use it for D-NAS policy jobs like snapshot and backup from snapshot. These jobs need a host from the backup host pool.

When you select this option in the policy, NetBackup creates a backup host pool named `All_Media_Server_Pool`. So, if you already have an existing backup host pool by that name, delete or rename the existing pool.

During run time, if NetBackup cannot find any suitable media server(s) to run the job, you can see an error with code 930. See [“Error code: 930: No supported media server is available in the All_Media_Server_Pool to use to backup the NAS shares.”](#) on page 115.

Advantages of the All media server pool option

- You do not need to create and frequently update the member hosts of the backup host pool after you add a new media server.
- No need to create separate backup host pools for NFS and SMB volumes. However, NetBackup still uses Linux media servers to back up NFS volumes. It uses Windows media servers to back up SMB volumes.

Limitations and considerations

Before setting up a NAS-Data-Protection policy for your workloads, consider the following.

Note: If you use cloud as a storage unit, you must configure the appropriate buffer size. Refer to the *NetBackup Cloud Administrator's Guide*.

Note the following important points about the NAS-Data-Protection policy.

- The NAS-Data-Protection is not supported in the DNAT and CloudScale environments.
- This policy does not support copy-based retention for Snapshot images. Ensure that you carefully plan your policy scheduling and snapshot retention in SLP.
- Client-side deduplication is not supported for the NAS-Data-Protection policy.
- A Vendor Change Tracking (VCT) enabled backup with an incremental schedule requires an initial base snapshot to identify the variances between the current and base snapshots. A differential incremental schedule refers to creating a base snapshot copy from a previous differential incremental, cumulative incremental, or full schedule. The cumulative incremental schedule refers to creating a base snapshot copy from a full schedule. During VCT-enabled backups with incremental schedule, if the base snapshot copy is not available, then the backup operation might fail with an error in the Activity monitor.
- NAS-Data-Protection policy is a snapshot-enabled data protection policy. You can configure only the Storage Lifecycle Policy (SLP) for the policy's storage destination. Additionally, the SLP should always have Snapshot as the primary job and backup from snapshot as the secondary job.
- If the NAS-Data-Protection policy is used in a backup host that is running any antivirus software, the parent backup from snapshot job might hang. The interaction between the NetBackup processes may be hindered by the antivirus software, resulting in process hang-ups. In this particular scenario, the nbcs process on the backup host might hang resulting in the

backup-from-snapshot job to hang. Create an antivirus exclusion for the nbcs process on the backup host.

To cancel the hung job:

- Note down the process ID of the nbcs process that is running on the backup host. This can be obtained from the job details section.
- Log on to the backup host and manually kill the nbcs process.
- Refer to the Technote for more details regarding how to exclude the NetBackup processes from virus scanning:
<https://support.cohesity.com/s/article/article-100004864>
- If the above steps cannot resolve the issue (and the nbcs hang persists), uninstall the network component from the antivirus. On Symantec Endpoint Protection, this is called the "Network and Host Exploit Mitigation" component.
- For the NAS-Data-Protection policy, multiple images are created for a single volume that is backed up. The number of images is equal to the value configured for the **Maximum number of streams per volume** in the policy. Since a single image cannot be referred from a single volume, NetBackup groups the images associated with a volume. When an operation is performed on one of the images in a volume, the same operation is also performed on the other grouped images in the volume. For example, if the **Maximum number of streams per volume** parameter is set as four, and you select one image for a volume to expire, the other three images also expire. The image grouping is applicable for the following operations:
 - Browse and restore
 - Image expiration
 - Image import
 - Image duplication
 - Image verification
 - Set primary copy

Note: Image grouping is not applicable for importing images as part of the Image Sharing operation.

- To enable checkpoint restart for NAS-Data-Protection policies created before upgrading to version 9.0, you must select the **Take checkpoints every** check box and enter a value in minutes.

D-NAS Planning and Tuning

This chapter includes the following topics:

- [Sizing guidelines for D-NAS](#)
- [Tuning parameters for Server Message Block \(SMB\)](#)
- [Tuning parameters for Network File System \(NFS\)](#)
- [NetBackup tuning parameters for the backup host](#)

Sizing guidelines for D-NAS

The sizing for D-NAS environments is based on your business requirements. Sizing depends on the storage array. It also depends on the characteristics of the NAS data that you protect. You can configure the D-NAS policies to use the NetBackup media server as a backup host. As a result, D-NAS backup processes may organically scale in terms of both performance and throughput.

CPU considerations

When a D-NAS policy runs, NetBackup uses the nbcs (crawler) and the bpbkar processes on the backup host. It uses the bpbrm and bptm on the media server. Each running process consumes CPU cycles. The nbcs process uses CPU the most. The nbcs is a multi-threaded crawler. It traverses the NAS share during backup and index operations. Multiple nbcs processes handle concurrent jobs for backup and indexing operations of NAS volumes. One nbcs process corresponds to one NAS volume backup. If you use multiple backup hosts to back up a single NAS volume, each backup host uses a separate nbcs process.

The crawler process is multi-threaded. Multiple threads traverse one NAS share during D-NAS policy execution. This can lead to spikes in CPU use for the nbcs process. You can ease CPU use by decreasing the number of threads used by the nbcs process. You can set the `MULTI_THREADED_CRAWLER_THREADS` parameter in the `bp.conf` file. This changes the thread count used by each nbcs process. The default value is 20, and you can specify a value in the range of 1 to 200. You must set this parameter on the backup hosts used for NAS backups. This is applicable for the hosts on NetBackup version 10.4 and above.

Memory considerations

The amount of memory used by a backup job of a single NAS share depends on the number of streams configured in the D-NAS policy. Each backup stream on a host uses one bpbrm, bptm, and bpbkar process. For example, if the policy is set to use 10 streams, then a single NAS share backup runs 10 instances of the bpbrm, bptm, and bpbkar processes. It runs only one nbcs process.

The amount of memory used by the bpbrm, bptm, and bpbkar processes is static and does not fluctuate much. The memory used by the nbcs process depends on certain data traits. These include file system hierarchy, number of files, and folders in a NAS share. If a NAS share has a very dense directory structure, then the nbcs process uses 200 MB of memory at its peak. If a NAS share has a flat hierarchy with millions of files in its directories, then nbcs uses 20% – 30% additional memory at its peak. Note that memory usage is not always at peak and decreases as the backup progresses. You may observe spikes in memory consumption depending on the data characteristics.

Memory usage is the same for all backup types. These include index tasks, first full, incremental, and accelerated full backups.

Table 3-1 Memory consumption for a single NAS backup with 5, 10, and 20 backup streams

NAS share memory consumption	Memory for 5 backup streams (MB)	Memory for 10 backup streams (MB)	Memory for 20 backup streams (MB)
On backup host	115 + 200 (crawler)	230 + 200 (crawler)	460 + 200 (crawler)
On media server *	780	1560	3120
Total memory consumption (approximate)	1095	1990	3780

* For index from snapshot operations, only the media server memory consumption is applicable.

Backup host and media server sizing

You can estimate the approximate memory needed for D-NAS backups using the table: **Memory consumption for a single NAS backup with 5, 10, and 20 backup streams**. Using this table, you can also estimate the number of backup hosts and media server hosts that you need to provision for D-NAS backups. This also helps to schedule the NAS backup jobs. For example:

- During backup, a policy with 10 streams and 10 NAS shares uses about 20 GB of system memory.
- During backup, a policy with 20 streams and 10 NAS shares uses about 38 GB of system memory.

Consider the following:

- When a single host serves as both the media server and the backup, all memory usage occurs on that host.
- When you use the multi-host feature for NAS share backups, the memory utilization of the crawler is distributed equally among all backup hosts.
- Use multiple media servers if the overall number of backup streams for D-NAS backups exceeds 200.

Tuning parameters for Server Message Block (SMB)

You must verify that SMB 3 is enabled to gain SMB Multichannel capabilities. SMB Multichannel enables the file servers to use multiple network connections simultaneously.

Configure the parameters as specified in the **SMB tuning parameters** table, on the backup host to gain better throughput. Note that the mentioned values for the parameters are from the test environment. You can set these values based on the configuration of the backup hosts.

NetBackup administrators can also refer to the SMB vendor's documentation to configure the parameters in the **SMB tuning parameters** table.

Table 3-2 SMB tuning parameters

Parameters	Values
Set-SmbClientConfiguration -ConnectionCountPerRssNetworkInterface	8
Set-SmbClientConfiguration -DirectoryCacheEntriesMax	4096

Table 3-2 SMB tuning parameters (*continued*)

Parameters	Values
Set-SmbClientConfiguration -DirectoryCacheLifetime	60
Set-SmbClientConfiguration -EnableBandwidthThrottling	0
Set-SmbClientConfiguration -FileInfoCacheEntriesMax	32768
Set-SmbClientConfiguration -FileInfoCacheLifetime	60
Set-SmbClientConfiguration -FileNotFoundCacheEntriesMax	32768
Set-SmbClientConfiguration -FileNotFoundCacheLifetime	60
Set-SmbClientConfiguration -MaxCmds	32768

Tuning parameters for Network File System (NFS)

Using the *nconnect* mount option, you can specify the number of connections (network flows) that must be established between the NFS client and the NFS endpoint. You can specify up to 16 connections. By default, NFS clients use a single connection between themselves and the endpoint. You can set the *nconnect* value for the mount command in the `/etc/nfsmount.conf` file.

[NFSMount_Global_Options]: set default options for each mount command run. The recommended setting for *nconnect* is 2, for the NFS volume backups to gain backup performance.

NetBackup tuning parameters for the backup host

Following are the NetBackup tuning parameters for the backup host:

- **DNAS_LOOKAHEAD_CACHE_SIZE_PER_VOLUME_MB**: The crawler uses the `DNAS_LOOKAHEAD_CACHE_SIZE_PER_VOLUME_MB` memory for cache purposes during each volume's backup job.
The default value is 100 MB per volume. During a backup, the crawler looks ahead to the snapshot file system and uses this additional cache space to store information for the backup job.
You can configure the parameters in the `bp.conf` file for the backup hosts of NetBackup 10.3 onwards. You must set this configuration on all the backup hosts that are associated with the backup host pool.
For example, if you configure the value to 512, then it means that 512 MB memory per volume is used as cache.

- **IGNORE_FILE_ACLS:** Set this parameter in the `bp.conf` file to 1, to ignore the backup of file-level ACLs and user group information. Note that only the directory ACLs are backed up, the file-level ACLs are not backed up.
During the restore operation, the *Directory* permissions are inherited. You can configure the parameters in the `bp.conf` file for backup hosts of NetBackup 10.3 onwards. Administrators must set this configuration on all the backup hosts that are associated with the backup host pool.

Pre-requisites for D-NAS configuration

This chapter includes the following topics:

- [Prerequisites for D-NAS configuration](#)
- [Domain user requirement for SMB share backups](#)
- [Minimum permissions for SMB share backup and recovery](#)
- [Minimum supported backup host versions for different features](#)
- [Configuring a backup host pool](#)

Prerequisites for D-NAS configuration

You need to meet the following prerequisites.

- Ensure that you have installed the NetBackup Snapshot Manager component. For more details, see the *NetBackup Snapshot Manager Install and Upgrade Guide*.
- Prepare the plug-in that you want to use for the NetBackup D-NAS configuration. For more details, refer to the *NetBackup™ Snapshot Manager for Data Center Administrator's Guide*.
- Identify the backup host(s) that you want to use for the configuration.
- If the NAS data protection policy uses a TAPE storage unit in SLP for protecting NAS volumes, then the number of tape drives must be greater than or equal to the maximum number of streams per volume, otherwise, backups fail. The other parameters of TAPE, like Media multiplexing and maximum concurrent write drives, do not have any effect on NetBackup D-NAS backups.

- For SMB backups using the NAS-Data-Protection policy the primary, media and backup host version should be 9.1 onwards.

Required firewall ports

If your NAS server is behind a firewall, open the following ports for NetBackup access.

- For bi-directional NFS access:
 - TCP port: 2049 (NFS v4)
 - TCP port: 111 (NFS v2/v3)
- For bi-directional SMB access:
 - TCP port: 445

Domain user requirement for SMB share backups

This step is required for Windows backup hosts for SMB share backups only. You must log on to the NetBackup client service and the NetBackup legacy network service as a domain user to perform the tasks described in the following sections.

Note: The Windows domain user must be a part of the local administrator's group.

To log on to the NetBackup services as a domain user:

- 1 Make sure that the NetBackup client service and the NetBackup legacy network service are running.
- 2 In Windows Services, double-click the NetBackup service.
- 3 Check the **Log on** tab: if any of these services is not logged on as the domain user, change the logon to the domain account and restart the service. If both services are not logged on as the domain user, you must do it in the following sequence:
 - Log on to the first service as a domain user and restart the service.
 - Log on to the second service as a domain user and restart the service.
- 4 Make sure that all NetBackup services are running.
- 5 Relaunch the NetBackup UI.

Minimum permissions for SMB share backup and recovery

To back up or recover SMB shares in NetBackup, the local or domain user must have these privileges:

- **SeBackupPrivilege**: Lets the SMB user back up SMB data without changing the ACL on the SMB shares.
- **SeRestorePrivilege**: Lets the user restore files and directories, override ACLs, and set a valid user or group SID as the file owner.
- **SeChangeNotifyPrivilege**: Lets the user bypass traverse checking without requiring traverse permissions for folders, symlinks, or junctions.

Minimum supported backup host versions for different features

Different features of NAS data protection policy require a backup host with a NetBackup version greater than or equal to the minimum supported backup host version. The following table specifies which feature is supported from which NetBackup version.

Table 4-1 NAS data protection policy features

Supported features	Minimum supported backup host version
Only NFS backup	8.3
NFS and Vendor change tracking	8.3
NFS and Checkpoint restart enabled backups	9.0
NFS and Accelerator enabled backups	9.0.1
SMB backups (including CPR, accelerator, Vendor change tracking)	9.0.1
NFS and SMB backups with Vendor Change Tracking (VCT) and accelerator	10.2
Multi-stream Restore	10.2
Replication	10.0
VCT support for Indexing jobs	10.3

Table 4-1 NAS data protection policy features (*continued*)

Supported features	Minimum supported backup host version
Forever incremental	10.3
Multi-host support	10.4
True Image Restore (TIR)	10.5

Configuring a backup host pool

Backup hosts and backup host pools are used for the NAS-Data-Protection policy based on dynamic multistreams.

You can use a NetBackup primary server, media server, or a standalone client as a backup host. For the hosts that you add to the backup host pool, their volumes are distributed for backup purposes on the backup hosts. This configuration results in a better backup performance.

Note: A NetBackup primary or media server running on Cohesity Flex Appliance is not supported as a backup host for a NAS-Data-Protection policy.

You can create a backup host pool with different versions of NetBackup hosts. You can create Windows backup host pools only with version 9.0.1 or later. Windows hosts with a version earlier than 9.0.1 are not displayed.

Note the following important points:

- In a backup host pool you can either have Linux hosts or Windows hosts only. A pool does not support hosts with both platforms.
- If you want to backup SMB shares along with the SMB ACLs, use Windows hosts in the backup host pool.
- All the hosts in the backup host pool must use the same Linux OS version. This way each host has the same version of NFS for consistent backups.
- For backup hosts with a multi-NIC setup, add the hostname that is already used on the NetBackup primary server. Do not add an alias name or any other host names in the backup host pool.

To configure a backup host pool

- 1 In the web UI, click **Host > Host properties**.
- 2 Select and connect to the primary server that you want to configure, and click **Edit primary server**.

- 3** Click **Backup Host Pools**.
- 4** Click **Add**.
- 5** Enter the backup host pool name.
- 6** (Conditional) This step is applicable only to the clients that you want to add to the list. In the **Enter hostname to add to the list** field, add the client name and click **Add to list**.
- 7** Select the **OS Type**.
- 8** Select the backup hosts that you want to add to the list.
- 9** Click **Save**.

Note: You cannot delete a backup host pool, if it is configured with an existing NAS-Data-Protection policy.

Configuring Storage Lifecycle Policies for D-NAS

This chapter includes the following topics:

- [About storage lifecycle policies](#)
- [Snapshot operation in SLP](#)
- [Creating a storage lifecycle policy for snapshots and snapshot replication](#)
- [Replication operation in the SLP](#)
- [Index from snapshot operation in an SLP](#)
- [Backup from snapshot operation in an SLP](#)
- [About retention period for backup from snapshot images](#)
- [Duplication operation in an SLP](#)
- [Retention types for SLP operations](#)

About storage lifecycle policies

A storage lifecycle policy (SLP) contains instructions in the form of storage operations to store data. Operations are added to the SLP that determine how the data is stored and copied or replicated. For example, the NetBackup administrator creates an operation that determines where the data exists as a snapshot, as a replication, or as a duplication. The administrator also determines the retention of the data at each storage unit or storage unit group.

An SLP that is configured for snapshots or snapshot replication must contain a specific, hierarchical combination of operations.

See [Figure 5-1](#) on page 38.. It represents an SLP for a replication scenario. In the example, the following operations are used:

- A **Snapshot** operation creates a snapshot.
- A **Replication** operation replicates the snapshot to another volume.
- A **Backup from Snapshot** operation creates a tar-formatted backup from the snapshot.
- A **Duplication** operation copies the backup to tape.

[Table 5-1](#) describes the four types of operations that are required in this example replication scenario.

Figure 5-1 Four types of operations in this example replication scenario




Table 5-1 Example of a storage lifecycle policy configured for snapshots and snapshot replication

Operation order in SLP	Operation	Description
1	Snapshot	Operation 1 creates a snapshot in the primary storage. The snapshot serves as the source for the other operations in the SLP. ■ The operation must be a Snapshot operation. ■ The storage unit can be any of the following types: Snapshot, AdvancedDisk, or MSDP storage unit
2 (Child to Operation 1)	Replication	Operation 2 replicates the snapshot that the first operation created. ■ The operation must be a Replication operation. ■ The storage must be any one of the following: Auto or <Vendor>_<ReplicationType>. Note: <Vendor>_<ReplicationType> is the replication type supported by the storage array vendor.

Table 5-1 Example of a storage lifecycle policy configured for snapshots and snapshot replication (*continued*)

Operation order in SLP	Operation	Description
3 (Child to Operation 2)	Backup from Snapshot	Operation 3 creates a tar-formatted backup copy of the snapshot. ■ The operation must be a Backup form Snapshot operation. This operation creates a backup image from the snapshot. ■ The storage must be a backup storage unit.
4 (Child to Operation 3)	Duplication	Operation 4 makes a duplicate copy from the tar-formatted backup copy. In this example, the copy is duplicated to tape media. ■ The operation must be a Duplication operation. This operation creates a backup copy of the tar-formatted image. ■ The storage must be a backup storage unit.

After the SLP is configured for different operations, the NetBackup administrator configures a backup policy that points to the snapshot SLP.

The **SLP Parameters** host properties in the **NetBackup Web UI** allow administrators to customize how SLPs are maintained and how SLP jobs run.

Snapshot operation in SLP

The **Operation** selections are the instructions in the storage lifecycle policy. A snapshot operation creates a point-in-time, read-only, disk-based copy of data. NetBackup provides several types of snapshots, depending on the device where the snapshot occurs.

Use a snapshot operation as the first operation in a storage lifecycle policy for a Snapshot Manager for Data Center snapshot and replication operations.


Table 5-2 Snapshot operation characteristics

Characteristic	Description
Storage unit selection	The snapshot operation does not write data to a NetBackup storage unit. However, you need to select a storage unit to specify the media server to use to launch the snapshot job. Use these storage units for snapshot operation: ■ Snapshot label ■ Media Server Deduplication Pool storage unit ■ AdvancedDisk storage unit Considerations for the 'Snapshot' label as the storage unit: ■ If the SLP contains only a snapshot operation, then NetBackup uses any available media server later than NetBackup version 10.0.1 to launch the snapshot job. ■ If a subsequent replication operation uses the snapshot, then the same media server performs the snapshot and replication operation. ■ If a subsequent Backup from snapshot operation uses the snapshot, then snapshot operation uses the storage unit that is selected for the Backup from snapshot operation. Note: To use the Snapshot label all the NetBackup hosts must be version 10.1 or later.
Child of	A snapshot operation cannot serve as the child of any other operation. Therefore, do not click on any other operation in the SLP when adding a snapshot operation.
Source for	A snapshot operation can be the source for the following operations: ■ Backup form Snapshot ■ Index form Snapshot ■ Replication
Hierarchy notes	If a snapshot operation appears in an SLP, it must be first in the operations list.
Job type	A snapshot operation generates a snapshot job in the Activity Monitor.
Window	Snapshot operations do not offer the option to create an SLP window.

Primary snapshot storage unit

A snapshot operation can use a Primary snapshot storage unit. That is, the storage unit represents an AdvancedDisk storage unit, Media Server Deduplication Pool storage unit or a Snapshot label.

The following figure shows an SLP that contains one primary-only snapshot operation, one Backup from snapshot operation, and one Duplication operation. The Backup from snapshot operation is used to create a backup from the snapshot on the primary-only snapshot operation. After the backup is created, it is duplicated to a Duplication operation.




Primary + Replication source snapshot storage unit

An SLP operation can use a primary storage unit for snapshots and an Auto or Vendor-supported replication type for the replication destination. If a subsequent replication operation uses the snapshot, then the same media server is used for each snapshot and replication operation.

The following figure shows an SLP that contains a snapshot as the storage unit for snapshot operation, and one Replication target snapshot storage unit as another operation.


Figure 5-2 SLP that contains a snapshot operation and a replication operation



Primary + Replication source + Replication target storage unit

An SLP operation can use a primary storage unit for snapshots and an Auto or Vendor-supported replication type for the replication destination.


Figure 5-3 SLP that contains a snapshot operation and two replication operations



Replication target snapshot storage unit

An SLP operation can use a primary storage unit for snapshots and an Auto or Vendor-supported replication type for the replication destination. It can also have backup from snapshot operation with a backup storage unit.

Figure 5-4 SLP that contains a snapshot operation, a replication operation, a Backup from snapshot operation, and a Duplication operation



Creating a storage lifecycle policy for snapshots and snapshot replication

Use the following procedure to configure a storage lifecycle policy that creates snapshots and snapshot replications with Snapshot Manager for Data Center. Only those options that are necessary to configure an SLP for Snapshot Manager for Data Center Replication are listed. You can use the NetBackup web UI to configure a storage lifecycle policy to create snapshots and snapshot replication.

To configure a storage lifecycle policy to create snapshots and snapshot replication

- 1 On the left, click **Storage Lifecycle Policies** under **Storage**.
- 2 Click **Add** on the right pane.
- 3 Enter a **Storage lifecycle policy name**.
- 4 Click **Add** to add operations to the SLP. The operations are the instructions for the SLP to follow and apply to the data that is specified in the backup policy.
- 5 In the **Properties** tab of the **New operation** page, select **Snapshot** from the **Operation** drop-down menu.

This **Snapshot** operation creates a snapshot of the primary data and serves as the source for other operations in the SLP. For example:

- A **Replication** operation.
- A **Backup form Snapshot** operation.
- An **Index form Snapshot** operation.

- 6 In the **Destination Storage** drop-down menu, select a storage unit. NetBackup displays only those storage units that are configured to contain primary snapshots.
- 7 Select the **Retention type** and the **Retention period** for the data in this storage unit. The **Retention period** option does not appear for all **Retention type** selections. Click **Create**.
- 8 To replicate the primary snapshot, create a **Replication** operation that is based on the snapshot. Select the check box in the row of the snapshot, and click **Add child**.
- 9 In the **Operation** drop-down menu, select **Replication**.
- 10 Under **Destination storage attributes**, select a **Replication target** that is configured to contain replicated snapshots. NetBackup displays only those targets that can act as target destinations.
- 11 Select the **Retention type** and the **Retention period** for the data in this storage unit.
- 12 The **Window** tab displays the following operation types: **Backup form Snapshot**, **Duplication**, **Import**, **Index form Snapshot**, and **Replication**.
Create a window during which secondary operations can run.
- 13 Click **OK** to create the SLP.

Continue to create operations, depending on the needs of your environment.

To cascade storage operations in the SLP, make sure to select the correct parent operation as the source for the child operation. If the correct operation is not selected, you unintentionally operate on an incorrect source.

Replication operation in the SLP

Use the **Replication** operation for the following types of replication:

- NetBackup Snapshot Manager for Data Center replication to replicate a snapshot.
- NetBackup Auto Image Replication to replicate a backup to a different domain or a different NetBackup primary server.

Table 5-3 Replication operation characteristics

Characteristic	Description
Storage unit selection	Under Destination storage attributes: For Snapshot Manager for Data Center replication, the following destinations are supported: ■ Auto ■ <Vendor>_<ReplicationType>
Child of	Click the appropriate operation when adding a replication operation. Using Snapshot Manager for Data Center Replication, a replication operation can be the child of a snapshot operation or the child of another replication operation.
Source for	A replication operation can be the source for the following operations: ■ Replication ■ Backup from snapshot ■ Index from snapshot
Job type	A Replication operation generates a Snapshot Replication job in the Activity Monitor.
Window	An SLP window can be created for a replication operation.

Index from snapshot operation in an SLP

The Index from snapshot operation indexes the contents of existing snapshots. When NetBackup indexes a snapshot, it creates an image catalog file in the NetBackup catalog for each snapshot. The presence of an image catalog file assists the user when a file needs to be restored from the snapshot, as described in the table.

Table 5-4 Restore operation

Type of restore	Where performed?	Description	Requirements
Live browse restore	NetBackup Backup, Archive, and Restore interface	You can navigate the directory structure to locate and select the files for restoration.	During a live browse restore, NetBackup automatically mounts the snapshot so that you can see what files it contains. Mounting and unmounting the snapshot can be time-consuming.

The Backup from snapshot operation also creates an image catalog file. An Index from snapshot may not be required if a Backup from snapshot occurs frequently enough for the restore needs in your environment. For example, if the Backup from snapshot runs once per week but file restores are required daily, consider using the Index from snapshot feature.

Snapshot restore requires that the snapshot is mounted, regardless of whether an index from snapshot is performed or not.

Table 5-5 Index from snapshot operation characteristics

Characteristic	Description
Storage unit selection	The Index from snapshot operation does not write data to a storage unit. However, a storage unit selection is needed to select the media server that is to be used to access the snapshot. As a best practice, use the storage unit from the snapshot or replication operation that is the source for this operation.
Child of	When an Index from snapshot operation appears in an SLP, it must be the child of a snapshot or replication operation. Therefore, select either a snapshot or a replication operation in the SLP when adding an Index from snapshot operation.
Source for	While an Index from snapshot operation cannot be the source for any operation, a replication operation can follow it.
Hierarchy notes	The index from snapshot operation can consume system resources and requires that each snapshot is mounted to create the image catalog file. See “Determining where and when the Index from snapshot operation occurs” on page 47.
Job type	An Index from snapshot operation generates an Index form Snapshot job in the Activity Monitor .
Window	An SLP window can be created for an Index from snapshot operation.

Consider the following items before using the Index from snapshot operation:

- Standard, NAS-Data-Protection, and VMware backup policy types support the use of storage lifecycle policies that contain the Index from snapshot operation.

Determining where and when the Index from snapshot operation occurs

The index from snapshot operation may be time-consuming and slow down the system resources. To populate the catalog, it is required that the snapshot is mounted or NetBackup gathers content details from the file system.

To help mitigate the extra resource and time that the operation may take, administrators can control when and where the index from snapshot operation runs:

- Use the Storage lifecycle policy option in the **Window** tab to schedule when the **Index form Snapshot** operation can run. Schedule the operation to run when it is least likely to interfere with other jobs.
- Use the following points to determine where to position the Index from snapshot operation in the SLP operations list:
 - Each NetBackup environment needs to determine where the operation works best in a specific SLP. To place the Index from snapshot operation too early (toward the top of the operations list), may consume time when the restore capabilities are not needed. To place the operation toward the end of the operations list may cause a delay to restore operations until earlier snapshots or replications are complete.
 - Use the Index from snapshot operation in an SLP only once. You can perform a restore from any snapshot after one `image.f` file is created.
 - Any operations list that includes a backup from snapshot operation, does not need an Index from snapshot operation. The backup from snapshot operation creates an `image.f` file. The only exception may be if the index is needed for restores before the backup from snapshot operation occurs.


Backup from snapshot operation in an SLP

Use the Backup from snapshot operation to create a tar-formatted copy of the snapshot. The new copy is a backup copy. The process is sometimes referred to as a *snapdupe* job.

Table 5-6 Backup from snapshot operation characteristics

Characteristic	Description
Storage unit selection	The selection must be a backup storage unit or a backup storage unit group. The selection cannot be a snapshot storage unit or a snapshot storage unit group.
Child of	A Backup from snapshot operation must use a snapshot or replication operation as its source. Therefore, select the snapshot operation in the SLP when adding a Backup from snapshot operation.
Source for	A Backup from snapshot operation can be the source for a Duplication operation.
Hierarchy notes	An SLP may contain more than one backup from snapshot operation. If the first backup from snapshot operation fails with an unrecoverable error, NetBackup does not attempt the second one. For a NAS-Data-Protection policy, the SLP supports only one backup from snapshot operation.
Job type	A Backup from snapshot operation generates a Backup job in the Activity Monitor. The Backup job that results from the Backup from snapshot operation is under the control of the SLP Manager. If an SLP window is configured, the Backup job runs during the configured SLP window. If no SLP window is configured, the Backup job can run at any time; possibly outside of the backup window as configured in the backup policy. You may experience a slight degradation in performance on the client or the client storage device while NetBackup accesses the snapshot.
Window	An SLP window can be created for a Backup from snapshot operation.

Figure 5-5 SLP that contains a Backup from snapshot operation



About retention period for backup from snapshot images

In a NAS-data-protection policy, expiration of a backup from snapshot image depends on the SLP retention period.

In a NAS-data-protection policy, a backup from snapshot job follows a hierarchy with a parent job and multiple child jobs. The number of child jobs depends on the number of streams that you define in the policy.

After all child backup jobs complete, the parent backup from snapshot job calculates the expiration time as follows:

$$I_{exp}=T_{child_max}+R_{slp}$$

Where:

- I_{exp} =Image expiration time of the NAS-data-protection policy backup job.
- T_{child_max} = Completion time of the longest-running child backup job.
- R_{slp} =Retention period defined in the SLP for the backup from snapshot jobs.

For example, if the last child backup job completes on Friday at 9 PM and the SLP retention period is 6 days, the expiration time for the backup from snapshot image is on the following Thursday at 9 PM.

Duplication operation in an SLP

Use the Duplication operation to create a copy of a Backup, a Backup from snapshot, or another Duplication operation. A media server performs the operation and writes the copy.

Note: Use the replication operation to create a copy of a snapshot operation.


Table 5-7 Duplication operation characteristics

Characteristic	Description
Storage unit selection	The selection must be a backup storage unit or a backup storage unit group.
	The selection cannot be a snapshot storage unit or a snapshot storage unit group.

Table 5-7 Duplication operation characteristics (*continued*)

Characteristic	Description
Child of	A Duplication operation can be the child of the following operations: ■ Backup operation ■ Backup from snapshot operation ■ A Duplication operation Therefore, select one of these operations in the SLP when adding a Duplication operation.
Source for	A Duplication operation can be the source of a Duplication operation.
Hierarchy notes	When a Duplication operation appears in an SLP, it cannot be the first operation.
Job type	A Duplication operation generates a Duplication job in the Activity Monitor .
Window	An SLP window can be created for a Duplication operation.

Figure 5-6 SLP that contains one Backup operation and two Duplication operations



Retention types for SLP operations

The **Retention type** for an operation in a storage lifecycle policy determines how long the data is kept on that storage media.

Note: You can set the Retention types for storage lifecycle policy operations from the NetBackup web UI.

Table 5-8 Operation and retention type configurations

Retention type	Backup operation	Snapshot operation	Replication operation	Backup from snapshot operation	Duplication operation
Fixed	Valid	Valid	Valid	Valid	Valid
Expire after copy	Valid	Valid	Valid	Valid	Valid
Maximum Snapshot limit	Invalid	Valid; SLP honors the policy setting.	Invalid	Invalid	Invalid
Target retention	Invalid	Invalid	Valid if the first operation in the SLP is an Import and if the storage is of the backup type.	Invalid	Valid if the first operation in the SLP is an Import.

Note: Retention is not associated with the **Index form Snapshot** operation because the operation does not create any copy.

Expire after copy retention type for SLP operations

The Expire after copy retention indicates that after all direct (child) copies of an image are successfully duplicated to other storage, the data on this storage is expired. The last operation in the SLP cannot use the **Expire after copy** retention type because no subsequent copy is configured. Therefore, an operation with this retention type must have a child.

It is not recommended that you enable Expire after copy retention for any storage units that are to be used with SLPs with either of the following: Accelerator or synthetic backups. The Expire after copy retention can cause images to expire while the backup runs. To synthesize a new full backup, the SLP backup needs the previous backup image. If the previous image expires during the backup, the backup fails.

For VCT-enabled incremental backups, the previous snapshot is required to generate the file change list with respect to the current snapshot. The Expire after copy retention expires the previous snapshot after the backup associated with that snapshot is complete.

Note: Although synthetic backups do support the use of storage lifecycle policies, SLPs cannot be used for the multiple-copy synthetic backups method.

If a policy is configured to use an SLP for the backup, the retention that is indicated in the SLP is the value that is used. The Retention attribute in the schedule is not used.

Expire after copy retention type cannot be used for the snapshot operation when a sync replication stage is added in the SLP.

An image copy with an Expire after copy retention expires as soon as all of its direct child copies have been successfully created. Any mirrored children must also be eligible for expiration.

Fixed retention type for SLP operations

The Fixed retention indicates that the data on the storage is retained for the specified length of time, after which the backups or snapshots are expired.

An image copy with a Fixed retention is eligible for expiration when all of the following criteria are met:

- The Fixed retention period for the copy has expired.
- All child copies have been created.
- All child copies that are mirror copies are eligible for expiration.

The **Fixed** retention period is always marked from the original backup time of the image. For example, if a tape device is down, causing a 2-day delay in creating a duplicate tape copy, the expiration time of the duplicate copy is not different due to the 2-day delay. The expiration time of the duplicate copy is still x days from the time that the original backup was completed. It does not matter when the copy was created.

If the replica copy has any dependency on its source copy, and the Fixed retention type is selected for the replica copy, then whichever copy (snapshot or any replica copy) in the SLP has the highest retention level, that copy's retention level is set as the retention level for the snapshot copy and all the replica copies in the SLP.

Maximum snapshot limit retention type for SLP operations

The **Maximum snapshot limit** determines the maximum number of snapshots that can be stored for a particular policy and client pair.

When the maximum is reached, the next snapshot causes the oldest job-complete snapshot to be deleted. A snapshot job is considered to be complete once all of its configured dependent copies are complete. (Dependent copies are created as a

result of Backup from snapshot, Index from snapshot, or Replication operations.) The practice is referred to as *rotation*. This retention type applies only to snapshots, and not to backups.

For example, Policy P1 contains two clients: C1 and C2. After the policy runs four times, it creates four snapshot images for C1 and four images for C2. If the **Maximum snapshot limit** is set to four, when the policy runs for the fifth time, NetBackup deletes the first snapshot that was created for both C1 and C2 to accommodate the fifth snapshot.

The **Maximum snapshots** parameter in the **Perform snapshot backups options** dialog determines the maximum number of snapshots. To access the dialog box, click **Options** under the Snapshot Client section in the backup policy.

Multi-host backup for volumes

This chapter includes the following topics:

- [About multi-host backup](#)
- [Stream distribution across multiple backup hosts](#)
- [Considerations for multi-host backups](#)
- [Monitoring and troubleshooting backup jobs](#)

About multi-host backup

Starting with NetBackup 10.4, you can back up a single NAS volume using multiple hosts from the backup host pool specified in the D-NAS policy. This multi-host backup capability, makes a single NAS volume accessible from multiple backup hosts. NetBackup can use a higher number of backup streams, across multiple hosts for backing up large NAS volumes. This significantly improves the performance of the NAS backups.

Stream distribution across multiple backup hosts

No additional options are available to enable multi-host backup of a NAS share. NetBackup automatically determines when to use multiple hosts based on the number of streams specified in the D-NAS policy. Note the following:

- If the policy specifies fewer than 20 streams, NetBackup uses a single host to back up the NAS share.

- If the policy specifies 20 or more streams, NetBackup distributes the backup streams evenly across an optimal number of hosts. Any host ready to start backup streams is included automatically.
- To optimize host resource usage, NetBackup starts at least ten backup streams on each host.

Considerations for multi-host backups

- NetBackup recommends using up to four backup hosts for a single NAS volume.
- If enough resources are not available on the backup hosts for stream distribution, then the backup job is queued until the hosts gain enough resources to start backup streams.
- If a backup stream running on any backup host is canceled or suspended, then the backup streams running on all the backup hosts for that NAS volume are canceled or suspended.
- If a backup stream running on any backup host fails due to an error, then the backup streams running on all the backup hosts for that NAS volume are terminated and the backup job goes into an incomplete state.

The rolling upgrade scenario for NetBackup 10.4

In a rolling upgrade scenario, you can upgrade the primary server to version 10.4 while the backup host pool still includes media servers and clients running older versions. Consider the following during this upgrade:

- D-NAS backups run as they did in versions prior to NetBackup 10.4, using a single host.
- If you upgrade one media server to NetBackup 10.4, D-NAS backups use only the upgraded server. Multi-host backups are not supported.
- If you upgrade multiple media servers and clients to NetBackup 10.4, D-NAS backups use all upgraded hosts in the backup host pool.

Monitoring and troubleshooting backup jobs

Backup streams in a multi-host environment can be identified by looking at the job details section of the parent backup-from-snapshot job. The example snippet mentions child jobs (streams) started on multiple backup hosts.

\r

```
Jan 23, 2024 11:36:12 AM - Info nbjm (pid=6493) Started child jobs 563, 564,
```

```
Jan 23, 2024 11:36:12 AM - Info nbjm (pid=6493) Started child jobs 573, 574,  
"
```

If there is a failure in resource allocation to any stream, the entire backup job for that volume fails and the parent job fails with error code 927. If you have selected the backup host pool as All media server, then the error 930 is shown.

You can resolve the majority of the problems for multi-host backups by checking the logs. For more details about troubleshooting, see:

- See [“Troubleshooting”](#) on page 103.
- See [“Logging directories for Linux platforms ”](#) on page 103.
- See [“Logging folders for Windows platforms ”](#) on page 106.

Configure D-NAS policy for NAS backups

This chapter includes the following topics:

- [About policies for NAS backups](#)
- [Planning for policies](#)
- [Prerequisites for D-NAS policies](#)
- [Configure D-NAS policy for NAS volumes](#)
- [Policy attributes](#)
- [Creating schedule attributes for policies](#)
- [Configuring the Start window](#)
- [Configuring the exclude dates](#)
- [Configuring clients](#)
- [Configuring backup selections](#)
- [Configuring exclude lists](#)
- [Ordering of backup from snapshot jobs](#)
- [About mixed mode volumes](#)
- [Configuring include and exclude lists](#)
- [Auto-resume backup for incomplete backup jobs](#)

About policies for NAS backups

Backup policies provide the instructions that NetBackup follows to back up any NAS share or volume. You can create a single policy to protect multiple NAS backups. You can select the shares or volumes that you want to protect using a policy. The NAS arrays are automatically discovered in the NetBackup environment and backed up. You need different policies to apply different backup logic to the shares and volumes.

You can configure the following using a policy:

- The Storage Lifecycle Policy (SLP) and media to use.
- Backup schedules: Full, Differential incremental, and Cumulative incremental.
- Backup selection: You can add an entire NAS share or volume to a policy, or select what to backup using backup selection and the exclude share features.

Planning for policies

Policy configuration is flexible enough to meet the various needs of all NAS backups in a NetBackup environment. To take advantage of this flexibility, take time to plan before starting to configure the policies in the **Policies** utility.

The following table outlines the steps to take to ensure that you get optimal results from your policy configurations.

Table 7-1 Steps for planning policies

Step	Action	Description
Step 1	Gather information about the NAS backups	Gather the following information about each NAS share or volume: ■ The NFS share or volume name. ■ The approximate size of each share or volume. ■ Total size of the data that you want to backup. One share may contain a large amount of data while the other shares are smaller. To avoid long backup times, include the large share in one policy and the smaller ones in another policy. It may be beneficial to create more than one policy for the large share.
Step 2	Consider the storage requirements	For NAS policies you must use a Storage lifecycle policy as the storage. The storage unit and volume pool settings apply to everything that are backed up by a policy. If the files that you are backing up have special storage requirements, create separate policies for the files, even if other factors are the same, such as schedules.

Table 7-1 Steps for planning policies (*continued*)

Step	Action	Description
Step 3	Consider the backup schedule	Create additional backup policies if the schedules in one policy do not accommodate all NAS backups. Consider the following factors when deciding to create additional policies: ■ Best times for backups to occur? To back up different NAS shares or volumes on different schedules may require additional policies with different time schedules. For example, create different policies for night and day time backups. ■ How frequently the shares/volumes change? If some share or volume changes more frequently than others, the difference may be enough to warrant creating another policy with a different backup frequency. This way you can backup the frequently occurring changes in the protected share or volume. ■ How long backups need to be retained? Each schedule includes a retention setting that determines how long NetBackup keeps the shares or volumes that are backed up by the schedule. Because the schedule backs up all the shares/volumes in the backup selection list, all shares/volumes should have similar retention requirements. Do not include the shares/volumes whose full backups must be retained forever, together in a policy where full backups are retained for only four weeks.
Step 4	Evaluate backup times	Evaluate total backup times for each schedule and further subdivide policies to reduce backup times to an acceptable level. For example, if the backup of <code>D:\User</code>, <code>D:\h001</code>, and <code>E:\h002\Projects</code> on NAS share1 takes too much time, create a new policy for <code>E:\h002\Projects</code>.
Step 5	Select exactly what to back up.	You do not need to back up entire shares or volumes, unless required. Create include and exclude lists to select and back up only the required file(s). See "Configuring include and exclude lists" on page 74.

Prerequisites for D-NAS policies

Before you begin creating a policy for NAS backups, consider the following prerequisites.

- Array credentials and plug-ins already added in NetBackup Snapshot Manager for Data Center.
- Keep handy information about the share(s)/volume(s) and the criteria that you want to use for selecting backups from them.

- Evaluate the requirement for NetBackup accelerators in your environment. If you want to use accelerators, you need to specify when creating the policy. See [“Accelerator for D-NAS ”](#) on page 78.
- Ensure that the necessary ports are open in the backup host and the configurations are in place. This is crucial for enabling server communication with the arrays through REST API calls.
- Evaluate the requirement for NetBackup multi-streaming in your environment. To use multi-stream, set the maximum number of data streams as required.
- Evaluate if you want to use multiple backup hosts and use the policy to perform load balance for the backup hosts during run-time, for multiple streams.

Configure D-NAS policy for NAS volumes

Using the NetBackup Snapshot Manager for Data Center you can perform hardware snapshots of shares and volumes. The snapshots are accessed on backup hosts and read by dynamic streams to create point-in-time backup copies. The following procedure describes how to configure a D-NAS policy to use hardware snapshots of NAS volumes.

Table 7-2 Configuration steps

Step	Description	Reference topic
1	Configure the NetBackup Snapshot Manager server in NetBackup	For more details, refer to the <i>Configure NetBackup snapshot manager for Data Center</i> chapter of the <i>NetBackup™ Snapshot Manager for Data Center Administrator's Guide</i> .
2	Configure the NAS storage array plug-in.	For more details, refer to the <i>Configure NetBackup snapshot manager for Data Center</i> chapter in the <i>NetBackup™ Snapshot Manager for Data Center Administrator's Guide</i> .
3	Add the backup hosts to a backup host pool. The backup hosts are responsible for data streaming.	See “Configuring a backup host pool ” on page 35.
4	Configure the SLP to use snapshot	See “About storage lifecycle policies” on page 37.

Table 7-2 Configuration steps (continued)

Step	Description	Reference topic
5	Configure a NAS-Data-Protection policy to perform the operations that are specified in the SLP.	- See “Policy attributes” on page 61. - See “Creating schedule attributes for policies ” on page 65. - See “Configuring the Start window” on page 68. - See “Configuring the exclude dates” on page 70. - See “Configuring clients” on page 71. - See “Configuring backup selections” on page 72. - See “Configuring exclude lists” on page 73.

Note: For all the supported NAS storage arrays, refer to the *NetBackup Snapshot Manager* section, under *Snapshot Solutions* in the *NetBackup Hardware and Cloud Storage Compatibility List (HCL)*.

Policy attributes

The following procedure describes how to select the attributes for the backup policy.

Select the policy attributes

- On the left, click **Protection > Policies**.
- Enter a name for the policy in the **Policy name** field.
- Select the **NAS-Data-Protection** option from the **Policy type** dropdown.
- In the **Destination** section, configure the following data storage parameters:
 - The **Data classification** attribute specifies the classification of the storage lifecycle policy that stores the backup. For example, a backup with a gold classification must go to a storage unit with a gold data classification. By default, NetBackup provides four data classifications: platinum, gold, silver, and bronze.

This attribute is optional and applies only when the backup is to be written to a storage lifecycle policy. If the list displays **No data classification**, the policy uses the storage selection that is displayed in the **Policy storage** list. If a data classification is selected, all the images that the policy creates are tagged with the classification ID.

- The **Policy storage** attribute specifies the storage destination for the policy's data. For NAS backups you must select a Storage lifecycle policy as storage. You can override this selection from the **Schedule** tab.

- 5 Take checkpoints every**-Specify the frequency for taking checkpoints during a backup. By taking checkpoints during a backup, you can save time if the backup fails. By taking checkpoints periodically during the backup, NetBackup can retry a failed backup from the beginning of the last checkpoint. A retry is often quicker than restarting the entire job.

The checkpoint frequency indicates how often NetBackup takes a checkpoint during a backup. The default is 15 minutes. The administrator determines checkpoint frequency on a policy-by-policy basis. When you select the checkpoint frequency, balance the loss of performance due to frequent checkpoints with the possible time lost when failed backups restart. If the frequency of checkpoints affects performance, increase the time between checkpoints.

Checkpoints are saved at object boundaries and point to the next object in the list to be backed up. Checkpoints cannot occur in the middle of an object backup. After the object is backed up, the checkpoint is saved.

- 6 The Limit jobs per policy** attribute limits the number of jobs that NetBackup performs concurrently when the policy is run. By default the box is cleared and NetBackup performs an unlimited number of backup jobs concurrently. Other resource settings can limit the number of jobs.

A configuration can contain enough devices so that the number of concurrent backups affects performance. To specify a lower limit, select **Limit jobs per policy** and specify a value from 1 to 999.

- 7** In the **Job priority** field, enter a value from 0 to 999999. This number specifies the priority that a policy has as it competes with other policies for resources. The higher the number, the greater the priority of the job. NetBackup assigns the first available resource to the policy with the highest priority.

- 8** The **Media owner** field is available when the **Policy storage** attribute is set to **Any Available**. The **Media owner** attribute specifies which media server or server group should own the media that backup images for this policy are written to.

- **Any** (default)-Allows NetBackup to select the media owner. NetBackup selects a media server or a server group (if one is configured).
- **None**-Specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.

9 Under **Snapshot Client and Replication Director** section, **Perform snapshot backups** and **Retain snapshot for Instant Recovery or SLP management** parameters are selected by default and are read only.

- Select **Enable vendor change tracking for incremental backups** to enable vendor change tracking. See [“About Vendor Change Tracking”](#) on page 82.
- **Perform snapshot backups**
Ensures that the policy creates snapshots of the disk array.
- **Retain snapshots for Instant Recovery or SLP management**
Ensures that the policy retains the snapshot after the backup completes.
- Click **Snapshot options** to configure the snapshot type that you want to capture, and the NetBackup Snapshot Manager that you want to use.
Here are the snapshot types:
 - **Auto** (default): The OpenStorage partner uses the best snapshot technology available to that partner to create the snapshot.
 - **Mirror**: The OpenStorage partner creates a copy dependent on the existence of the source. (The source can be the original snapshot or another replica.) Therefore, the retention of the replica depends on the retention of the source. If the source is deleted, the mirror is automatically deleted.
 - **Clone**: The OpenStorage partner creates an independent copy of the volume. The copy process can take some time as the entire copy must be complete. The snapshot that is created is independent of the source.
 - **Cow**: Copy-on-write snapshots. The OpenStorage partner makes a copy of modified data in a new location. Read requests for data that are unchanged are directed to the original volume. Read requests for changed data are sent to the copied blocks of the snapshot. Data blocks that have changed since the snapshot was created are described in metadata.
- The **Perform off-host backup** options are not available for selection.

From the **Snapshot Manager** list, select the snapshot manager that you want to use. The snapshot manager that you select must be configured to support the workload that you want to protect with the policy.

10 To activate the policy, select the option **Go into effect at**, and set the date and time of activation. The policy must be active for NetBackup to use it. Make sure that the date and time are set to the time that you want to resume backups.

To deactivate a policy, clear the option. Inactive policies are available in the **Policies** list.

- 11** The **Collect true image restore information** attribute specifies whether the policy collects the information necessary to perform a true image restore. A true image restore (TIR) restores the contents of a directory to reflect the contents of the directory at the time of an incremental or a full backup. Files that were deleted before the backup are not restored.

With the attribute enabled, a restore based on an incremental backup includes all files that were backed up since the last full backup. The restore also includes those files that were deleted at any time during that period.

NetBackup starts to collect the true image restore information with the next full or incremental backup for the policy. The true image restore information is collected for each client regardless of whether any files were changed.

NetBackup does not provide true image restores based on the time of a user backup or archive. However, NetBackup uses a user backup for a true image restore if the backup is more recent than the latest automatic full or incremental backup.

See [“About true image restore”](#) on page 88.

- 12** The **Allow multiple data streams** option is selected by default and is read only. This option allows NetBackup to divide automatic backups for each query into multiple jobs. Because the jobs are in separate data streams, they can occur concurrently.

Multi-stream jobs consist of a parent job to perform stream discovery and child jobs for each stream. Each child job displays its job ID in the Job ID column in the **Activity monitor**. The job ID of the parent job appears in the Parent Job ID column, which is not displayed by default. Parent jobs display a dash (-) in the Schedule column.

Optionally, for Flex Scale and MSDP Volume Group (MVG) environments, select the **Use multiple MSDP nodes** option.

This option improves load balancing for multi-stream backups, by enabling distribution of backup streams across the multiple MSDP nodes, optimizing throughput, and reducing potential congestion on individual nodes.

- 13** Select the **Use Accelerator** option to enable accelerator for the policy.

NetBackup accelerator increases the speed of backups. The increase in speed is made possible by the change detection techniques on the array. The backup host uses the change detection techniques to identify the changes that occurred between the last backup and the current state of the array. The array sends the changed data to the media server in a more efficient backup stream. The media server combines the changed data with the rest of the array's data that is stored in previous backups.

If a share/volume or a portion of the share/volume is already in storage and has not been changed, the media server uses the copy in storage rather than reading it from the array. The result is a full NetBackup backup.

See [“Accelerator for D-NAS”](#) on page 78.

- 14** Under the **Client-side deduplication** options, **Disable for all clients** is selected by default, and it is read only.
- 15** The **Keyword phrase** attribute is a phrase that NetBackup associates with all backups or archives based on the policy. Only the Windows and UNIX client interfaces support keyword phrases.

Clients can use the same keyword phrase for more than one policy. The same phrase for multiple policies makes it possible to link backups from related policies. For example, use the keyword phrase “legal department documents” for backups of multiple clients that require separate policies, but contain similar types of data.

The phrase can be a maximum of 128 characters in length. All printable characters are permitted, including spaces and periods. By default, the keyword phrase is blank.

- 16** The **Dynamic data streaming attributes** is selected by default and is read only. If required change the value for **Maximum number of streams per volume**. The maximum number of streams per volume determines the number of backup streams that are deployed for backing up each volume. For example, if a policy contains 5 volumes and the value of this parameter is set to 4, then a group of 4 backup streams for each volume is seen, thereby a total of 20 child backup streams and 5 parent backup streams as part of backup execution of the policy.

Creating schedule attributes for policies

This topic describes how to configure certain schedule properties for NAS backups. The schedule properties vary according to your specific backup strategy and system

configuration. Additional information about other schedule properties is available in the *NetBackup Administrator's Guide, Volume I*.

To create a schedule:

- 1 On the left, click **Policies**, under **Protection**. Click the **Schedules** tab. Under **Backup schedules**, click **Add**. Click the **Attributes** tab.
- 2 In the **Attributes** tab, enter a name for the schedule in the **Name** field.
- 3 Select the **Type of backup**:
 - **Full Backup**-A complete backup of the objects that contain all of the data objects and the log(s).
 - **Differential Incremental Backup**-Backup of the changed blocks since the last backup. If you configure a differential incremental backup, you must also configure a full backup.
 - **Cumulative Incremental Backup**-Backs up all the changed objects since the last full backup. All objects are backed up if no previous backup was done.
- 4 Select the **Accelerator forced rescan** option to activate NetBackup accelerator for this policy. This option creates a checksum of the content of each object during backup. It uses checksums for change detection. It provides a safety net by establishing a new baseline for the next accelerator backup. See [“Accelerator forced rescan option”](#) on page 80.
- 5 The **Override policy storage selection** attribute works as follows:
 - **Disabled**-Instructs the schedule to use the **Policy storage** as specified on the policy **Attributes** tab.
 - **Enabled**-Instructs the schedule to override the **Policy storage** as specified on the policy **Attributes** tab.
Select the storage from the list of previously configured storage units and storage lifecycle policies. If the list is empty, no storage is configured.
- 6 The **Override policy volume pool** attribute works as follows:
 - **Disabled**-Instructs the schedule to override the volume pool that is specified as the Policy volume pool on the policy Attribute tab. If no policy volume pool is specified, NetBackup uses NetBackup as the default.
 - **Enabled**-Instructs the schedule to override the volume pool that is specified as the **Policy volume pool** on the policy **Attribute** tab. Select the volume pool from the list of previously configured volume pools.
- 7 The **Override media owner** selection attribute works as follows:

- **Disabled**-Instructs the schedule to use the media owner that is specified as the **Media owner** in the policy **Attribute** tab.
- **Enabled**-Instructs the schedule to override the media owner that is specified as the **Media owner** in the policy **Attribute** tab.
 Select the new media owner from the list:
 - **Any.**
 NetBackup selects the media owner, either a media server or server group.
 - **None.**
 Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.

8 Under **Schedule type, select **Calendar** or **Frequency**.**

- **Calendar**-Calendar-based schedules let you create a job schedule based on a calendar view. Select **Calendar** to display the **Include dates** tab.
 Enable **Retries allowed after run day** to have NetBackup attempt to complete the schedule until the backup is successful. With this attribute enabled, the schedule attempts to run, even after a specified run day has passed.
- **Frequency**-Use the **Frequency** attribute to specify how much time must elapse between the successful completion of a scheduled task and the next attempt.
 For example, assume that a schedule is set up for a full backup with a frequency of one week. If NetBackup successfully completes a full backup for all clients on Monday, it does not attempt another backup for this schedule until the following Monday.
 To set the frequency, select a frequency value from the list. The frequency can be seconds, minutes, hours, days, or weeks.

9 Specify a **Retention period for the backups. This attribute specifies how long NetBackup retains the backups. To set the retention period, select a period (or level) from the list. When the retention period expires, NetBackup deletes information about the expired backup. After the backup expires, the objects in the backup are unavailable for restores. For example, if the retention is 2 weeks, data can be restored from a backup that this schedule performs for only 2 weeks after the backup.**

- 10 The **Media multiplexing** attribute specifies the maximum number of jobs from the schedule that NetBackup can multiplex to any drive. Multiplexing sends concurrent backup jobs from one or several clients to a single drive and multiplexes the backups onto the media.

Specify a number from 1 through 32, where 1 specifies no multiplexing. Any changes take effect the next time a schedule runs.

- 11 Click **Add** to add the attributes, or click **Add and add another** to add a different set of attributes for another schedule.

Configuring the Start window

The **Start window** tab provides controls for setting periods during which NetBackup can start jobs when using a schedule. Periods are referred to as windows. Configure the windows to satisfy the requirements necessary to complete a job.

For example, create different windows:

- One for the backups that open each day for a specific amount of time.
- Another for the backups that keep the window open all week.

Adding, changing, or deleting a time window in a policy schedule

Use one of the following procedures to add, change, or delete a time window.

To configure the Start window:

- 1 On the left, click **Policies**, under **Protection**. Click the **Schedules** tab. Under **Backup schedules**, click **Add**. Click the **Start window** tab.
- 2 To indicate the opening of the time window, do the following:

Drag your cursor in the time table.

Click the day and time when you want the window to start and drag it to the day and time when you want the window to close.

Use the settings in the dialog box.

- In the **Start day** field, select the first day that the window opens.
- In the **Start time** field, select the time that the window opens.

- 3 To indicate the closing of the time window, do one of the following:

Drag your cursor in the time table.

Click the day and time when you want like the window to start and drag it to the day and time when you want the window to close.

Enter the duration of the time window.

Enter a length of time in the **Duration (days, hours, minutes)** field.

Indicate the end of the time window.

- Select a day in the **End day** list.
- Select a time in the **End time** field.

Time windows show as bars in the schedule display.

Specify enough time to allow all clients in the policy to complete a backup.

Consider allowing extra time in the schedule in case the schedule starts late due to factors outside of NetBackup. (Delays due to unavailable devices, for example.) Otherwise, all backups may not have a chance to start.

4 As necessary, do any of the following:

Click **Delete**.

Deletes the selected time window.

Click **Clear**.

Deletes all the time windows from the schedule display.

Click **Duplicate**.

Replicates the time window for the entire week.

Click **Undo**.

Erases the last action.

5 Do one of the following:

Click **Add**.


To save the time window and leave the dialog box open.

Click **Add and Add another**.

To save the time window and add another.

Example of schedule duration

This example illustrates the effect of schedule duration on two full backup schedules. The start time for Schedule B begins shortly after the end time for the previous Schedule A. Both schedules have three arrays with backups due.



The diagram illustrates the following points:

Point 1

Array A3 starts within the Schedule A time window but doesn't complete until after the Schedule B start time. However, Array A3 runs to completion even if the window closes while the backup is running. Array B1, on Schedule B, begins as soon as Array A3 completes.

Point 2

Schedule A does not leave enough time for all the Arrays on Schedule B to be backed up. Consequently, Array B3 is unable to start because the time window has closed. Array B3 must wait until the next time NetBackup runs Schedule B.

Configuring the exclude dates

Use the **Exclude dates** tab to exclude specific days from a schedule for a backup policy. If a day is excluded from a schedule, jobs do not run on that day. The tab displays a calendar of three consecutive months. Use the lists at the top of the calendar to change the first month or year displayed.

To exclude a day from a schedule:

- 1 On the left, click **Policies**, under **Protection**. Click the **Schedules** tab. Under **Backup schedules**, click **Add**. Click the **Exclude dates** tab.
- 2 Use one or more methods to indicate the days to exclude:
 - Select the day(s) on the 3-month calendar that you want to exclude. Use the drop-down lists at the top of the calendar to change the months or years.
 - To indicate **Recurring week days**:
 - Click **Set all** to select all of the days in every month for every year.
 - Click **Clear all** to remove all existing selections.
 - Check a box in the matrix to select a specific day to exclude for every month.
 - Click the column head of a day of the week to exclude that day every month.
 - Click the **1st**, **2nd**, **3rd**, **4th**, or **Last** row label to exclude that week every month.
 - To indicate **Recurring days of the month**:
 - Click **Set all** to select all of the days in every month.
 - Click **Clear all** to remove all existing selections.
 - Check a box in the matrix to select that day to exclude each month.
 - Click **Last** to exclude the last day of every month.
 - To indicate **Specific dates**:
 - Click **New**. Enter the month, day, and year in the dialog box. The date appears in the **Specific dates** list.
 - To delete a date, select the date in the list. Click **Delete**.
- 3 Click **Add** to save the changes.

Configuring clients

The **Clients** tab contains a list of clients to be backed up (or acted upon) by the selected policy. A client must be included in the list of at least one backup policy to be backed up.

Placing a client in more than one backup policy can be useful. For example, place the client name in two policies to back up different sets of files on the client according to different policy rules.

A client must be included in the list of at least one active backup policy to be backed up.

Adding clients in a policy:

- 1 On the left, click **Policies**, under **Protection**. Click the **Schedules** tab. Under **Backup schedules**, click **Add**. Click the **Clients** tab.
- 2 To change an existing client, select a **NAS supported vendor** from the list. Click **Add**.
- 3 In the **NAS array and array head selection** dialog, click the array on the left to see the list of array heads on the right. select one or more required array heads from the list. Click **Save**.
- 4 In the **Clients** tab, select a client from the list, and click **Save**.

To delete a client, select a client in the **Clients** tab, and click **Delete**. Hold down Shift to select multiple clients.

Configuring backup selections

The **Backup selections** tab contains a list of what to back up on each client, host, or instance when NetBackup runs an automatic schedule (for example, a full backup). The list does not apply to user backups or archives, where users select the shares/volumes to back up before they start the operation.

To configure backup selections:

- 1 On the left, click **Policies**, under **Protection**. Click the **Schedules** tab. Under **Backup schedules**, click **Add**. Click the **Backup selections** tab.
- 2 Select NFS or SMB as required.
- 3 (Optional) Select **Include mixed volumes** to select the volumes protected by both NFS and SMB protocols, under Backup selections.
- 4 From the **Backup host pool** dropdown, select any of the following:
 - From the Backup host pool dropdown, select any of the following:
 - Select an already created backup host.
 - Select **All media server pool**. See [“About the All media server pool option”](#) on page 24.
 - Select **Create new backup host pool** to create a new backup host pool. See [“Configuring a backup host pool”](#) on page 35.
 - If you have not created a backup host pool already, the **Backup selections** dialog box appears. Click **Yes** to configure. See [“Configuring a backup host pool”](#) on page 35.

- 5 From the **Backup host pool** list, select the preferred pool, and select one or more required volumes from the table.
- 6 To add a new volume, click **Add**.
 - In the **Add backup selection** dialog, do one of the following:
 - From the **Pathname or directive** list, select the preferred and click **Add to list**.
 - Click **Browse** and select the preferred.
 - Click **Add**.
- 7 Click **Save** to save the backup selection.

Configuring exclude lists

You can exclude the volumes from the backup selection list that you do not want to back up. For example, if `/prodVol*` is the backup selection, there may be a volume `/prodVol-Scratch`, which you do not want to back up.

To exclude volumes:

- 1 On the left, click **Policies**, under **Protection**. Click the **Schedules** tab. Under **Backup schedules**, click **Add**. Click the **Exclude volumes** tab.
- 2 On the **Exclude Volumes** tab, in the **Volume to exclude** field, add the preferred volumes that you do not want to backup.

The list of excluded volumes appears in the table below. You can edit or delete the entries in the table.

Ordering of backup from snapshot jobs

With the NetBackup 9.1 release, all SLP-initiated backup from snapshot jobs for Policy, Client, or Backup selection are scheduled sequentially. One scheduled backup from snapshot job must be complete before the subsequent job can start. This behavior applies to the NAS-Data-Protection policy also. For example, if there are two scheduled snapshot jobs T1 and T2, and T1 is scheduled before T2. The ordering ensures that the backup from snapshot job for T1 must be complete before the backup from snapshot job for T2 is started.

For the NAS-Data-Protection policy, if checkpoint restart is enabled and the backup from snapshot job is in a suspended or an incomplete state, then that job must be resumed first, so that the next backup from snapshot job can run.

With the NetBackup 10.4 release, you can configure a multi-stream policy containing a backup host pool with multiple hosts. This way, you can distribute the stream among all available backup hosts equally and on the best available resource hosts.

About mixed mode volumes

Mixed mode volumes are the volumes having multi-protocol access. Storage array vendors allow both NFS and SMB access to a NAS volume. D-NAS policy allows backup of volumes having multi-protocol access. The protocol used for backup of these volumes depends on the type of backup host pool specified in the policy. If a Linux backup host pool is specified in the policy, these volumes get backed up using NFS protocol. If a Windows backup host pool is specified in the policy, these volumes get backed up using SMB protocol.

This mechanism can be used to backup SMB share data using a Linux backup host. For this to happen, enable NFS and SMB access to the NAS volumes.

Note: When a Linux backup host is used to backup an SMB share, the backup of SMB ACLs does not happen. Only the SMB share data is backed up. Similarly, when a Windows backup host is used to backup an NFS share, the NFS ACLs are not backed up. Only the NFS share data is backed up.

Configuring include and exclude lists

With D-NAS backups, you can create include and exclude lists of the directories and the files that you want to protect in the client. NetBackup uses the include or exclude lists to skip or include files and directories during backups.

These lists are verified against the backup selections that you made for the client.

The exclude list indicates the files and directories to exclude from a backup.

The include list specifies the exceptions to the exclude list. This list indicates the excluded files that you want to back up from the client. You can use the include list when you want to backup only a few files from a large number of excluded files in a directory. Use the include list to add back the files that you eliminate with the exclude list.

Both the exclude and include lists must be configured on all backup hosts inside the backup host pool that you use for the D-NAS policy.

For syntax guidelines for the lists and more information, see *Exclude list properties* under the section *Configuring hosts* in the *NetBackup Web UI Administrator's Guide*.

Note the following:

- To exclude any folder, use the format: `\vol_name\dir`. Do not use a slash at the end of the path for the directory.
- To backup only the “dir” folder inside “vol”. In the exclude lists section add the path = `\vol\*`
In the include list, to create the exception to the exclude list, add path = `\vol\dir` [Do not add a slash at the end]
`\vol\dir\*`
This configuration backs up only the data of the `\vol\dir` folder. Note that if you add only one rule of the two to the include list, the rule does not work and everything on `\vol\dir` are excluded. You must add both rules to the include list.

Examples for NFS:

Suppose that we have six directories from d1 to d6 inside `volume1`.

```
/volume1/d1
```

```
/volume1/d2
```

```
/volume1/d3
```

```
/volume1/d4
```

```
/volume1/d5
```

```
/volume1/d6
```

Here is the exclude list:

```
/volume1/*
```

Here is the include list:

```
/volume1/d1
```

```
/volume1/d1/*
```

After successful back up the following directories are skipped from backup.

```
/volume1/d2
```

```
/volume1/d3
```

```
/volume1/d4
```

```
/volume1/d5
```

```
/volume1/d6
```

Only `/volume1/d1` is backed up successfully.

Examples for SMB:

Consider the directory structure:

```
\volume\d1\file1
```

```
\volume\d2\folder1\file1
```

```
\volume\d2\folder2\file2
```

```
\volume\d2\folder2\file3
```

```
\volume\d3\folder3\file3
```

```
\volume\d3\folder3\file2
```

Here is the exclude list:

```
file1
```

```
\volume\d2\folder2\file2
```

```
file3
```

Here is the include list:

```
\volume\d1\file1
```

```
\volume\d3\folder3\file3
```

After successful backup, the following directories are skipped from the backup:

```
\volume\d2\folder1\file1
```

```
\volume\d2\folder2\file2
```

```
\volume\d2\folder2\file3
```

```
\volume\d3\folder3\file2
```

These are the directories that are backed up:

```
\volume\d1\file1
```

```
\volume\d3\folder3\file3
```

Auto-resume backup for incomplete backup jobs

With this feature, whenever a backup job having Checkpoint restart enabled, goes to an incomplete state, the job auto-resumes after a configured time interval. If the resumed job fails again, it is marked as incomplete, until a scheduled retry job runs again. The job is marked as failed, when all the scheduled retry attempts have completed running. You can configure the number of retries for each job and the delay between two retry attempts.

To configure the number of retries and the interval between two retries:

- 1 On the left, click **Host properties**, under **Hosts**.
- 2 Select the host that you want to configure. If the host is not connected, click **Connect**. Once the host is connected, click **Edit primary sever**. Click **Global attributes**.
- 3 To set the interval between two retries, specify a value in minutes in the **Job retry delay** field.
- 4 To set the number of retries for each job, enter values in the **Schedule backup attempts** field. You can specify the number of retries that NetBackup should attempt for the specified time interval in hours.

For more details, see *NetBackup Administrator's Guide, Volume I*.

Using Accelerator

This chapter includes the following topics:

- [Accelerator for D-NAS](#)
- [About the track logs for Accelerator](#)
- [Track log sizing considerations](#)
- [Notes on accelerator for D-NAS](#)
- [Accelerator forced rescan option](#)

Accelerator for D-NAS

NetBackup accelerator provides faster full backups at the cost of incremental backups, eventually reducing the backup window for customers. With this solution, more data is protected in the specified backup window and less bandwidth consumption.

After an initial full backup that protects all data from the filer, NetBackup accelerator backs up only the changed data from the filer to the media server. The media server combines the changed data with any previous backup images to create a new full backup image. If a file or portion of a file is already in storage and has not been changed, the media server uses the copy in storage, rather than reading it from the filer to complete the backup image. The result is a faster NetBackup NAS backup.

To configure accelerator for D-NAS, select the **Use Accelerator** check box that is found on the policy **Attributes** tab.

Benefits of accelerator for D-NAS policy

Here are some benefits of using accelerator with D-NAS:

- Creates a compact backup stream that uses less network bandwidth between the filer and NetBackup servers.

- Reduces the I/O and CPU overhead on the media server and backup host.
- Independent of storage arrays. Works with all the supported NAS storage arrays.

About the track logs for Accelerator

NetBackup Accelerator uses track log to detect the new, change, and modify files in the subsequent Full and Increment backups. The track log is a binary file that you should not attempt to edit. For D-NAS policy each backup stream maintains its own track log. The number of backup streams depends on the policy attribute **Maximum number of streams per volume**.

Track log location on the backup host:

Windows:

```
install_path\NetBackup\track\primary_server\storage_server  
\client\policy_name\backup_selection\S1\
```

Linux:

```
install_path/netBackup/track/primary_server/storage_server  
/client/policy_name/backup_selection/S1/
```

Track log location on the primary server:

Windows:

```
install_path\NetBackup\db\track\primary_server\storage_server\  
client\policy_name\backup_selection\S1\
```

Linux:

```
install_path/NetBackup/db/track/primary_server/storage_server/  
client/policy_name/backup_selection/S1/
```

Where S1, S2... Sn are the number of streams.

You can manually delete track logs safely if any of the following situations occur:

- You can disable the **Use Accelerator** option.
- The backup selections are changed.
- The policy is renamed.
- The storage server that is used to perform the backup is changed.
- The primary server that is used to control the backups is changed.

Track log sizing considerations

The accelerator track log stores file system metadata, and the unique fingerprints of files (128KiB segments). The track log size is relative to the size of the file system, and the number of backup files. Different track logs are created for each policy, client, and stream combination.

Here are some general guidelines, but the requirements in a specific environment might be different. Environments with a high rate of data change may require a larger track log size.

For D-NAS policy, the track log is stored on the backup host, and transferred to the primary server in-line during the backup operation. You can use the following formula to calculate the approximate size:

Total Track log size in Bytes for a NAS volume backup job = $2 * ((\text{Number of files} * 200) + ((\text{Total used disk space in KiB} / 128\text{KiB}) * 20))$

For example, 1 TB NAS volume with one million files = ~ 701 MiB total track log size. If four streams are configured for backup and one million files are equally distributed amongst four streams, then each stream's track log can be ~175 MiB in size.

Notes on accelerator for D-NAS

In-line track log persistence on the primary server:

- The track log contents are synced in line with the primary server.
- If the backup host changes for subsequent backup, the track log is copied from the primary server to the current backup host.

If the number of backup streams is changed [policy attribute **Maximum number of streams per volume**] then in the next backup, the existing track logs are not used. A new baseline is created for the subsequent backups. After changing the number of backup streams the accelerator optimization becomes "0" in the next backup and all the contents of the volume are backed up.

Accelerator forced rescan option

The policy schedules tab contains an option called Accelerator forced rescan. This option creates a checksum of the content of each file during backup. It uses the checksums for change detection. It provides safety by establishing a new baseline for the next Accelerator backup.

The Accelerator forced rescan option detects the following events:

- The file's data changes, but the file's metadata does not change.
- The file's metadata becomes corrupted, such that it does not indicate that the file has changed.
- A malicious user or application changes the file's metadata such that it does not indicate that the file has changed.

Table 8-1 Required full-backup schedules for each Accelerator policy

Full backup schedules	Notes on schedule frequency
First schedule: Accelerator forced rescan disabled	Configure this schedule to run most of your Accelerator full backups.
Second schedule: Accelerator forced rescan enabled	Configure this schedule to run less often than the first full-backup schedule. For example: If the first full-backup schedule runs weekly, run the second schedule (with the Accelerator forced rescan option enabled) every few months. However, the best frequency for this schedule depends upon your environment. Note: If the policy has no schedule that enables the Accelerator forced rescan option, all full backups automatically enable that option and backup performance is reduced.

Note the following about the Accelerator forced rescan option:

- The Accelerator forced rescan option is grayed out if the Use Accelerator option on the Attributes tab is not selected.
- Because of the checksum processing on the backup host, this option reduces backup speed as compared to the Use Accelerator option on its own. The speed reduction depends on the backup host's configuration and its current processing load. If the backup host is busy with many jobs when Accelerator backup begins, checksum processing can reduce backup speed.

Using Vendor Change Tracking

This chapter includes the following topics:

- [About Vendor Change Tracking](#)
- [About NetApp SnapDiff support](#)
- [Using VCT with accelerator for D-NAS](#)
- [Using VCT for indexing](#)
- [Changing the number of backup streams when VCT and accelerator are enabled](#)
- [Index from snapshot for D-NAS](#)
- [Using VCT with NetBackup client exclude list](#)

About Vendor Change Tracking

Several NAS storage array vendors have a difference engine that identifies the list of changed files and directories between two snapshot copies of the same volume. When Vendor Change Tracking (VCT) is enabled for a D-NAS policy, NetBackup does not perform any file system tracking for backup of NAS volumes. Instead it relies only on the change-list from the difference engine of the storage array to perform backup of files and directories. This process optimizes the backup process.

To use this feature, ensure that the storage array provides this capability. D-NAS policy supports VCT-enabled backups and index operations for Dell EMC PowerScale (Isilon), NetApp, Nutanix Files, and Qumulo NAS arrays.

VCT is not applicable in the following conditions:

- Full schedule is only supported when accelerator is enabled along with VCT.

- The base snapshot is not available.
- Expire after copy retention option is selected for the snapshot in SLP.

About NetApp SnapDiff support

NetBackup integrates NetApp SnapDiff technology to enable indexing and backup of NetApp NAS volumes and shares. You can use NetApp SnapDiff v2 or v3 to perform VCT-enabled backups for NetBackup NAS-Data-Protection policy type. To enable NetApp SnapDiff, set the `USE_SNAPDIFF_FOR_NETAPP_DNAS_BACKUPS` parameter in the `bp.conf` file to 1, on the NetBackup primary server.

Refer to the NetApp communication [CPC-00352](#) regarding the effect on third-party applications.

Consider the following for SnapDiff-enabled backups of NetApp NAS volumes and shares:

- The destination storage that you use for the backup copies must reside on NetApp storage. This condition does not apply to the snapshot or replica copies.
- Use only NetApp storage for the NAS share backups, and any other copies that are created using duplication or Auto Image Replication (AIR).
- NetBackup supports the following storage units for storing backup, duplicate, and replicate copies:
 - MSDP storage residing on NetApp
 - NetApp StorageGRID (LAN)
 - NetApp StorageGRID (WAN)

Using VCT with accelerator for D-NAS

With NetBackup 10.2 onwards, you can enable accelerator along with VCT in the D-NAS policy for NAS backups. VCT along with accelerator technology is supported with Dell EMC PowerScale (Isilon), NetApp, Nutanix Files, and Qumulo NAS arrays.

With NetBackup 10.3, you can enable this feature for a full schedule also, which enables the forever-incremental backup capability. After the initial full backup, no more full backups are required.

During full or incremental backups, NetBackup leverages the storage array vendor's technology to get the change list (added, modified, and deleted files) between the two point-in-time snapshots. In the subsequent incremental or accelerator backups, NetBackup need not do a complete scan of the NAS volumes to determine the change list.

After an initial full backup that protects all data from the filer, NetBackup accelerator backs up only the changed data from the filer to the media server.

Combining both these functions in a single backup policy, the backup window is reduced to full and incremental backups.

- A regular full scan of the volume is performed for the full schedule with forced re-scan enabled in the schedule. NetBackup does not use the VCT information in this scenario.
- Irrespective of the schedule, the change list is obtained using VCT. This change list is used as the source for backup.

Using VCT for indexing

When Vendor Change Tracking (VCT) is enabled for a D-NAS policy, NetBackup does not perform any file system traversal for indexing a NAS volume. Instead, NetBackup uses the change list from the difference engine of the storage array to perform indexing of files and directories.

D-NAS policy supports the VCT-enabled index backups for NetApp, Dell EMC PowerScale (Isilon), Nutanix Files, and Qumulo NAS arrays.

Consider the following before using the index from snapshot with VCT for D-NAS policy:

- Index job for the first full schedule:
 - When the index from snapshot for the first full schedule runs for a NAS volume. Then, NetBackup performs the file system traversal and generates an image catalog file.
 - File system traversal happens as the array vendors do not have any previous snapshot to get the change list.
- Index job for the subsequent full schedule using VCT:
 - When VCT is enabled in the policy and a subsequent full schedule is run, NetBackup uses the following:
 - Previous full image catalog.
 - change list provided by the array vendor's difference engine from NBSM.
 - Using the previous catalogs, a synthetic catalog is created for the subsequent full schedule.
 - If number of streams per volume is changed from the last successful full index from snapshot, then NetBackup performs a file system traversal.

- Index job for an incremental schedule using VCT:
 When VCT is enabled in the policy and incremental schedule runs:
 - Differential incremental schedule: The change list is identified using the current snapshot and the snapshot of the last successful index job.
 - Cumulative incremental schedule: The change list is identified using the current snapshot and the snapshot of the last successful full index job.
- Following are the file entries that are added to the image catalog for index operation as per the schedule:
 - Full schedule: The full set of files.
 - Differential incremental schedule: Files that are added or modified after the last index job run for any schedule.
 - Cumulative incremental schedule: Files that are added or modified after the index job run for the last full schedule.
- For a VCT-enabled D-NAS policy, the Snapshot or Replication retention period in SLP must be greater than the policy's schedule frequency.
- VCT is not applicable if the base snapshot is not available.
- The expiration period of the previous full and subsequent incremental snapshots is postponed until the next full schedule is completed.
- For non-VCT supported arrays, the index job continues to use mount and file system crawl method.

Changing the number of backup streams when VCT and accelerator are enabled

When you enable both VCT and accelerator in a policy, new tracklogs are created based on the previous tracklogs and VCT data. If you change the number of backup streams in the policy attributes, then in the next backup, the existing tracklogs are discarded. In this case, NetBackup does not use the VCT mechanism for the backup, and performs regular incremental backup. The accelerator optimization is also discarded and all the contents of the volume are backed up.

Index from snapshot for D-NAS

The index from snapshot operation indexes the contents of the existing snapshots. When NetBackup indexes a snapshot, it creates an image catalog file in the NetBackup catalog for each snapshot. This image catalog file assists you when you restore a file from the snapshot.

Note: The index of a NAS share that uses NFS protocol on a Linux host is faster as compared to the SMB protocol on Windows hosts. It is recommended to enable mixed protocols for the NAS share on the storage array and use NFS protocol (using Linux host) for the index operation.

The backup from snapshot operation also creates an image catalog file. If a backup from snapshot occurs frequently for the restore job in the environment, then an index from snapshot is not required.

For example, if the backup from snapshot runs once per week but the file restores are required daily, consider using index from snapshot.

Starting with NetBackup version 10.4, you can perform index from snapshot operation of a NAS share using two streams.

Consider the following items before using the index from snapshot operation for D-NAS policy:

- The index from snapshot operation can run from a full or an incremental schedule.
- For index from snapshot operations, a volume uses one or two streams. This is unlike backup from snapshot operations.
- A single SLP can have either index from snapshot or backup from snapshot operation, but not both.
- Location of image catalog (. f) is: `<NetBackup Installation directory>/db/images/<StorageArrayFiler>/<directory>`
- To dump or check the contents of the image catalog, use the `cat_convert` utility as follows:

```
: /usr/opensv/netbackup/bin/cat_convert -dump <. f file name>
```
- NetBackup primary server, media server, and backup hosts must be of version 10.4 or higher to use multistream index from snapshot.
- By default, NetBackup uses one stream per NAS share for the index from snapshot operation.

Note: The index of a NAS share is faster when you use the NFS protocol as compared to using the SMB protocol. It is recommended to enable mixed protocols for the NAS share on the storage array and use the NFS protocol (using Linux host) for the index operation.

Using VCT with NetBackup client exclude list

You can configure the exclude lists for excluding files, directories, or patterns for VCT, from backup and index operations of D-NAS policy. See [“Configuring include and exclude lists”](#) on page 74.

In the policies using VCT, NetBackup uses array the vendor’s change list capabilities to avoid explicit file system traversal during backup or index operations. Once you run a backup or index operation with an exclude list, do not modify the exclude or include list for the subsequent backup or index operation. If you modify the include or exclude lists, the associated files may not get backed up or indexed in the subsequent operations.

To remove or modify the exclude list; remove VCT from the policy and perform the index or backup operation.

Using true image restore

This chapter includes the following topics:

- [About true image restore](#)
- [Configuring TIR information retention time](#)
- [Considerations for using TIR](#)

About true image restore

A True Image Restore (TIR) recovers the contents of a directory as they were at the time of the previous incremental or full backup. The files deleted before the backup are not restored.

About true image restore information

The **Collect true image restore information** option, in D-NAS policy, specifies whether the NAS policy can collect the information necessary to perform a true image restore.

NetBackup begins gathering TIR data for each incremental or full backup of the policy. The TIR information is collected for each volume regardless of whether any files were changed.

Collect true image restore is not supported along with Vendor Change Tracking (VCT) unless you have enabled accelerator. If the NetBackup admin wants to enable both TIR and VCT options in the backup policy, then accelerator must be enabled in the NAS policy.

NetBackup admin can perform TIR restore using single-stream or multi-stream restore.

An example of TIR

The table below shows the files that were backed up in the `/voll/dir/` directory between 12/01/2024 and 12/04/2024. The NAS policy that ran the backups had the **Collect true image restore information** option selected.

Table 10-1 Sample backups taken before a TIR

Day	Type of backup	Backed-up files
12/01/2024	Full	file1 file2 dirA/fileA dirB/fileB file3
12/02/2024	Incremental	file1 file2 dirA/fileA -----
12/02/2024	Incremental	file1 file2 dirA/fileA -----
12/03/2024	Incremental	file1 file2 dirA/fileA -----
12/04/2024	Incremental	file1 file2 ----- ----- ----- file4

Note: The dashes (-----) indicate the files were deleted before the backup.

If you restore the 12/04/2024 version of the `/voll/dir/` directory, you can see the following results:

After a regular restore

The restored directory contains all files and directories that ever existed in `/voll/dir/` from 12/01/2024 (last full backup) through 12/04/2024:

```
file1
file2
dirA/fileA
dirB/fileB
file3
dirC/fileC
file4
```

After a TIR

The restored directory contains only the files and directories that existed at the time of the incremental backup done on 12/04/2024:

```
file1  
file2  
file4
```

NetBackup does not restore any of the files that were deleted before the 12/04/2024 incremental backup. NetBackup did not restore these files or directories because they did not exist at the time of the incremental backup. The last selected incremental backup was the reference for the TIR.

Configuring TIR information retention time

You can specify how long NetBackup retains the true image restore information. This value defines the number of days to maintain true image restore information in catalog files (.f). After the specified number of days, the images are deleted. This setting is applicable to all policies for which NetBackup collects true image restore information. The default duration is one day.

To configure the TIR information retention time:

- 1 On the left, click **Host properties**, under **Hosts**.
- 2 Select the primary server that you want to edit, click **Edit primary server**.
- 3 Click **Clean up**, in the field **Keep true image restoration (TIR) information**, enter a value in days.
- 4 Click **Save**.

TIR information cleanup depends on two factors. The first is the TIR information expiration period specified in the UI, and the next full backup. NetBackup does not clean up the TIR information in the catalog files until the next full backup, even if the specified expiration period is over.

Considerations for using TIR

When you use TIR:

- NetBackup collects additional information for the incremental backups that to gather the TIR information.
- You can only list and select the directories in TIR mode. Display and restore of individual files are not supported.
- TIR is only supported from the backup copy (TAR). When you have a primary snapshot copy, the restore is done using the TAR copy.

Replication using D-NAS policy

This chapter includes the following topics:

- [Replication using D-NAS policy](#)

Replication using D-NAS policy

Using the NetBackup Snapshot Manager for Data Center you can replicate the hardware snapshots of NFS and SMB shares. The replicated snapshots are accessed on backup hosts and read by dynamic streams to create the point-in-time backup copies. The following procedure describes how to configure a NAS-Data-Protection policy to use hardware snapshots and replication of NAS volumes.

Note: For all the supported NAS storage arrays for replication, refer to the *NetBackup Snapshot Manager* section, under *Snapshot Solutions* in the *NetBackup Hardware and Cloud Storage Compatibility List (HCL)*.

Table 11-1 Configuration steps

Step	Description	Reference topic
1	Configure the NetBackup Snapshot Manager server in NetBackup.	For more details, refer to the <i>Installation and Upgrade</i> chapter of the <i>NetBackup™ Snapshot Manager for Data Center Administrator's Guide</i> .

Table 11-1 Configuration steps (*continued*)

Step	Description	Reference topic
2	Configure the NAS storage array plug-in.	For more details, see the <i>Configure NetBackup snapshot manager storage array plug-ins</i> chapter of the <i>NetBackup™ Snapshot Manager for Data Center Administrator's Guide</i> .
3	Add the backup hosts to a backup host pool. The backup hosts are responsible for data streaming.	See “Configuring a backup host pool ” on page 35.
4	Configure the SLP to use snapshot and replication	For more details about replication, refer to these chapters in the <i>NetBackup™ Snapshot Manager for Data Center Administrator's Guide</i> : ■ Storage array replication ■ Supported storage arrays in the data center See “About storage lifecycle policies” on page 37.
5	Configure a NAS-Data-Protection policy to perform the operations that are specified in the SLP.	See “About policies for NAS backups” on page 58.

Restoring from D-NAS backups

This chapter includes the following topics:

- [Considerations for restoring from D-NAS backups](#)
- [About the Overwrite existing file option during restore](#)
- [Multi-stream restores from D-NAS backups](#)
- [RBAC role for D-NAS restores](#)
- [Scanning for malware](#)
- [Restore everything to a different location](#)
- [Restore individual files and folders to different locations](#)
- [Original location restores for D-NAS Policy](#)
- [Restore Azure Files backups to the original SMB volume](#)
- [Point-in-time rollback](#)

Considerations for restoring from D-NAS backups

You can use the NetBackup web UI to restore individual files or directories, or a volume.

Points to remember before restoring:

- NAS-Data-Protection policy does not support original location restore.
- The destination client for the restoration must be a NetBackup host. For example, a media server or backup host.

- If you select either of the following rename options, ensure that you change the destination path:
 - Rename hard links.
 - Rename soft links.
- NetBackup version 10.5 onwards supports multi-stream restore from NetBackup-created backup images, snapshots, and replica copies. To use the multi-stream restore feature, upgrade the primary, media, mount host, and restore destination clients to version 10.5 onwards.
 - To use the multi-stream restore feature for backup copies, all the images selected in the selected time interval must be backup copies and set as primary copies.
 - Similarly, to use the multi-stream restore feature from the snapshot or replica copies, all the images selected in the selected time interval must be snapshot or replica copies and set as primary copies.
- NetBackup version 10.5 onwards supports checkpoint restarts for single and multi-stream restores.
 - The checkpoint restore is supported for the backup copy only. It is not supported for restoring from snapshots or replica copies.
 - NetBackup automatically maintains the checkpoint interval intervals.
- For a NAS volume, if there are multiple copies of NetBackup-created backup images, then NetBackup restores data from the first non-snapshot or non-replica copy. To restore from a specific backup copy, set that copy as the primary copy in the NetBackup catalog.

About the Overwrite existing file option during restore

For D-NAS restores, it is recommended to enable the **Allow overwrite of existing files** option to restore the directory metadata.

In D-NAS backups, the files and folders of the NAS share are backed up in different backup images. During a restore, a file inside a directory may get restored first, and then the directory gets restored from another backup image. For example, the file `/vol/dir/file` is in the image: `image1`, and the directory `/vol/dir` is in the image: `image2`. During a restore, the `image1` contents may get restored first, when the directory is not present. At this point, NetBackup creates the directory, but without the attributes of the original directory. During the restore from `image2`, the restore process finds that the directory is already created, and it must be replaced to restore

the original directory attributes. So, to restore directory metadata, the option **Allow overwrite of existing files** must be enabled during restore.

Multi-stream restores from D-NAS backups

Starting with NetBackup version 10.2, you can restore a NAS volume using multiple restore streams from NetBackup-created backup images. Each restore stream runs in parallel, and the restored files are dynamically distributed to each of these restore streams. This helps in achieving optimal performance during restore job. The result is a faster NAS volume restores. Ensure that the primary and media servers, along with the NetBackup client are upgraded to 10.2 to use multi-stream restore.

NetBackup version 10.5 onwards supports multi-stream restore from NetBackup-created backup images, snapshots, and replica copies.

Each NAS volume restore has a separate parent-child job hierarchy. The parent job is a controller job for the NAS volume and one or more child jobs perform the actual restore of the data. Each child restore job represents one restore stream.

RBAC role for D-NAS restores

Starting with NetBackup 10.2, you can use the NetBackup web UI to perform D-NAS restores. There is a default NAS Administrator role that you can use to perform restores of the NAS volume backups.

For more information, see *NetBackup Web UI Administrator's Guide*.

Scanning for malware

NetBackup lets you scan the backed-up images for malware and determines the last good-known image that is malware-free. If any malware is detected during the scanning, you can see a notification in the web UI.

If you try to recover a malware-affected backup image, NetBackup shows you a warning message and confirmation is required for proceeding. You need special user privileges to restore from malware-affected images.

For more information about malware scanning, see the *Malware detection* chapter in the *NetBackup™ Security and Encryption Guide*.

Restore everything to a different location

You can restore the entire backup to a different location, or restore individual files and folders to different locations.

Restoring from D-NAS backups

- 1 On the left, click **Recovery**. On the Recover page, under **Regular recovery**, click **Start recovery**.
- 2 In the **Basic properties** tab, select the policy type as **NAS-Data-Protection**. Select **Restore type** as **Normal Backups** or **True image backups**. Select the **Source client** from where you want to recover, and click **Next**.
- 3 In the **Recovery details** tab, select a volume on the left to recover. You can click the volume on the left to see the contents of that volume on the right side, and select the required folder(s) or file(s) on the right to restore. You can click a folder on the left to see the individual files and folders inside, on the right. Select any file(s) or folder(s) to recover.

Click **Edit** to change the date range of the displayed images. Click **Use date picker** to provide the start and end time of the required interval. Click **Use backup history**, to see the entire backup history of the image. Select the required image(s) and click **Apply**.

- 4 In the **Recovery options** tab, select **Restore everything to a different location**. Select the NetBackup host for the target location. Specify the **Target location** for the restore in the host. In the **Target location** dialog, click the drive on the left to see the locations on the right. Select a location and click **Add**.
- 5 (Optional) **Select Allow overwrite of existing files, Restore directories without crossing mount points, Rename hard links, and Rename soft links** as required.
- 6 Specify the number of simultaneous data streams that you want to use during restore, in the **Number of restore streams per volume** field. You can specify a value from 1 to 20. A higher number might affect network performance.

Note: If you specify the number of restore streams as 1, then all the backup streams of a NAS volume are restored sequentially.

- 7 Use the default media server for the restore, or specify a new one. Specify a job priority and click **Next**.
- 8 In the **Review** tab, review all the parameters. To go back and change a parameter, click **Previous**. Click **Start recovery**.

Restore individual files and folders to different locations

You can restore the individual files and folders in the backup to different locations.

Restoring files and folders from D-NAS backups

- 1 On the left, click **Recovery**. On the Recover page, under **Regular recovery**, click **Start recovery**.
- 2 In the **Basic properties** tab, the select policy type as **NAS-Data-Protection**. Select **Restore type** as **Normal Backups** or **True image backups**. Select the **Source client** from where you want to recover, and click **Next**.
- 3 In the **Recovery details** tab, select a volume on the left to recover. You can click the volume on the left to see the contents of that volume on the right side, and select the required folders(s) or files on the right to restore. You can click a folder on the left to see the individual files and folders inside, on the right. Select any file(s) or folder(s) to recover.

Click **Edit** to change the date range of the displayed images. Click **Use date picker** to provide the start and end time of the required interval. Click **Use backup history**, to see the entire backup history of the image. Select the required image(s) and click **Apply**.

- 4 In the **Recovery options** tab, select **Restore individual directories and files to different locations**. Select the NetBackup host for the target location. In the **Specify destinations for the selected item(s)** table, click **Browse** in the **Destination** column to specify destinations to the items that you want to recover.
- 5 (Optional) Select **Allow overwrite of existing files**, **Restore directories without crossing mount points**, **Rename hard links**, and **Rename soft links** as required.
- 6 Specify the number of simultaneous data streams that you want to use during restore, in the **Number of restore streams per volume** field. You can specify a value from 1 to 20. A higher number might affect network performance.

Note: If you specify the number of restore streams as 1, then all the backup streams of a NAS volume are restored sequentially.

- 7 Use the default media server for the restore, or specify a new one. Specify a job priority and click **Next**.
- 8 In the **Review** tab, review all the parameters. To go back and change a parameter, click **Previous**. Click **Start recovery**.

Original location restores for D-NAS Policy

Even though the **Restore everything to its original location** option is disabled for the D-NAS policy, it is possible to restore data to the original location. Use the following methods:

- **NFS Shares:** Manually mount the NFS share to one of the NetBackup hosts. Use that host as the destination client and the mount path as the destination location.
- **SMB Shares:** Specify the UNC path of the SMB share as the destination and one of the NetBackup hosts as the destination client. For example: \\<IP or FQDN>\<SMB_Share_Name>

Subsequently, you can perform a point-in-time rollback. See [“Point-in-time rollback”](#) on page 100.

Restore Azure Files backups to the original SMB volume

To restore backed-up data along with metadata such as ACLs to the original Azure Files SMB volume, Microsoft recommends accessing the volume using the storage account name and access key.

To restore Azure Files data to the original SMB volume using NetBackup, create a Credential Management entry with the storage account name and access key.

To create the Credential Management entry:

- 1 On the left, click **Credential management**. Under **Named credentials**, click **Add**.
- 2 In the **Add credential** dialog, select the **NetBackup** option, and click **Start**.
- 3 In the **Basic properties** tab, enter the **Credential name**. This name must be the same as the Azure Files Storage account name. Optionally, enter a tag and description for the credential. Click **Next**.

- 4 In the **Category** tab, from the **Category** dropdown, select **Storage key**.
Under **Access Details**, enter the storage account in the format *<Domain name>\<Storage account name>*. Enter the **Access Key** for the Storage account.
- 5 In the **Permissions** tab, add the roles that you want to assign to this credential.
- 6 In the **Review** tab, verify the configuration. Edit if needed. When done, click **Finish**.

Restoring to the original SMB volume

- 1 On the left, click **Recovery**. On the Recover page, under **Regular recovery**, click **Start recovery**.
- 2 In the **Basic properties** tab, select the policy type as **NAS-Data-Protection**. Select **Restore type** as **Normal Backups** or **True image backups**. Select the **Source client** from where you want to recover, and click **Next**.
- 3 In the **Recovery details** tab, select a volume on the left to recover. You can click the volume on the left to see the contents of that volume on the right side, and select the required folder(s) or file(s) on the right to restore. You can click a folder on the left to see the individual files and folders inside, on the right. Select any file(s) or folder(s) to recover.

Click **Edit** to change the date range of the displayed images. Click **Use date picker** to provide the start and end time of the required interval. Click **Use backup history**, to see the entire backup history of the image. Select one or more required images and click **Apply**.
- 4 In the **Recovery options** tab, select **Restore everything to a different location**. Select the NetBackup host for the target location. This must be a Windows computer. Enter the UNC path of the Azure Files SMB volume, for example, *\\<storage account name>\<volume name>* for the **Target location**.
- 5 (Optional) **Select Allow overwrite of existing files, Restore directories without crossing mount points, Rename hard links, and Rename soft links** as required.
- 6 Specify the number of simultaneous data streams that you want to use during restore, in the **Number of restore streams per volume** field. You can specify a value from 1 to 20. A higher number might affect network performance.

Note: If you specify the number of restore streams as 1, then all the backup streams of a NAS volume are restored sequentially.

- 7 Use the default media server for the restore, or specify a new one. Specify a job priority and click **Next**.
- 8 In the **Review** tab, review all the parameters. To go back and change a parameter, click **Previous**. Click **Start recovery**.

Point-in-time rollback

You can also restore a snapshot of an entire file system, volume, or share with minimal I/O. This type of restore is called point-in-time rollback. All the data in the snapshot is restored. Single file restore is not available in a rollback.

Warning: Rollback deletes all files that were created after the creation date of the snapshot that you restore. Rollback returns a file system or volume to a given point in time. Any data changes or snapshots that were made after that time are lost.

Also, if there are multiple logical volumes on a single disk or volume group and if you perform a Point in Time Rollback of a specific logical volume, the entire disk or volume group is restored to the point in time.

Rollback is available only when you restore the file system, volume, or share to the original location on the client.

Performing rollback using snapshot:

- 1 On the left, click **Recovery**. On the Recover page, under **Regular recovery**, click **Start recovery**.
- 2 In the **Basic properties** tab, select the policy type as **NAS-Data-Protection**. Select **Restore type** as **Point In Time Rollback**. Select the **Source client** from where you want to recover, and click **Next**.
- 3 **Recovery details** tab, the backups are displayed in the **Backup History** table, select the image for restore. Click **Edit** to search for the list of snapshot images, for all dates (you cannot set a date range).
- 4 Select an image from the list, and click **Next**.

Under **Restore target options**, select **Restore everything to original location**. You need to specify a NetBackup host.

- 5 (Optional) Under **Recovery options**, select **Force rollback even if it deletes the snapshot(s) taken after that backup point**. If you do not select this option, recovery does not run, if any snapshots taken after the selected backup point exists.

- 6** If you do not want to use the default media server for recovery, select the required media server. Set a priority for the recovery job.
- 7** In the **Review** tab, review all the selections that you made. To change any setting, click **Previous**. Click **Start recovery** to start the recovery.

Troubleshooting

This chapter includes the following topics:

- [Troubleshooting](#)
- [Setting the log level](#)
- [Logging directories for Linux platforms](#)
- [Logging folders for Windows platforms](#)
- [Logging folders for multi-stream restore](#)
- [Exclude list is not working during backup](#)
- [Restore from a snapshot fails with status 133](#)
- [Backup from snapshot jobs do not start after the snapshot job completes successfully](#)
- [Backup from snapshot fails with error 50](#)
- [Backup from snapshot parent job fails with error 4213: Snapshot import failed](#)
- [Backup host pool creation fails with the error "Failed to fetch host list"](#)
- [Snapshot job fails and the snapshot command does not recognize the volume name](#)
- [Accelerator enabled incremental backup of NetApp NAS volume](#)
- [Snapshot method: Auto](#)
- [Backup from snapshot jobs for NAS-Data-Protection policy fail with error 4213](#)
- [A full VCT-enabled indexing job runs, when followed by a non-VCT indexing job with a backup host prior to version 10.3](#)

- Backup from snapshot jobs for NAS data protection policy fail with error 927
- Error code: 930: No supported media server is available in the All_Media_Server_Pool to use to backup the NAS shares.
- Restore from NAS array volume fails with the status: 174 Media manager – system error occurred.
- NAS job fails with the error: Crawler process timed out after 600 seconds waiting for streams to attach with shared memory.
- D-NAS backup fails with the error: The file system crawler process timed-out waiting for streams to attach with shared memory. (3003)
- Isilon backup from snapshot failed with the Snapshot cannot be mounted error.
- Discovery and snapshot operations fail with the errors 156 and 1542

Troubleshooting

You can resolve many problems on your own by creating logging directories, reproducing the problem, and checking the logs. For an in-depth description of NetBackup logs, refer to the *NetBackup Troubleshooting Guide*.

For explanations of NetBackup job status codes, refer to the *NetBackup Status codes Reference Guide*.

Setting the log level

To create detailed log information, place a *VERBOSE* entry in the `bp.conf` file on the NetBackup primary and client server. Alternatively, set the Global logging level to a high value in the **Logging** dialog, under both **Primary Server Properties** and **Client Properties**.

These directories can eventually require a lot of disk space. Delete them when you are finished troubleshooting and remove the *VERBOSE* option from the `bp.conf` file. Alternatively, reset the Global logging level to a lower value.

Logging directories for Linux platforms

To create logging directories use the `/usr/opensv/netbackup/logs/mklogdir` script. You can also create the directories using an access mode of 755 so NetBackup can write to the logs.

Table 13-1 Linux logging directories for a snapshot operation

Path of the log directory	Where to create the directory?
/usr/opensv/netbackup/logs/bprd	NetBackup primary server
/usr/opensv/logs/nbjm	NetBackup primary server
/usr/opensv/netbackup/logs/bpbm	NetBackup media server
/usr/opensv/netbackup/logs/bpfis	NetBackup backup host client

Table 13-2 Linux logging directories for a backup operation

Path of the log directory	Where to create the directory?
/usr/opensv/netbackup/logs/bprd	NetBackup primary server
/usr/opensv/logs/nbjm	NetBackup primary server
/usr/opensv/logs/nbstserv	NetBackup primary server
/usr/opensv/netbackup/logs/bpdm	NetBackup primary server
/usr/opensv/netbackup/logs/bptm	NetBackup media server
/usr/opensv/netbackup/logs/bpbm	NetBackup media server
/usr/opensv/netbackup/logs/bpfis	NetBackup backup host client
/usr/opensv/netbackup/logs/bppfi	NetBackup backup host client
/usr/opensv/netbackup/logs/bpbkar	NetBackup backup host client
/usr/opensv/logs/ncfnbs	NetBackup backup host client

Table 13-3 Linux logging directories for an index from operation

Path of the log directory	Where to create the directory?
/usr/opensv/netbackup/logs/bprd	NetBackup primary server

Table 13-3 Linux logging directories for an index from operation (*continued*)

Path of the log directory	Where to create the directory?
/usr/opensv/logs/nbjm	NetBackup primary server
/usr/opensv/logs/bpdbm	NetBackup primary server
/usr/opensv/netbackup/logs/bptm	NetBackup primary server
/usr/opensv/netbackup/logs/bpbrm	NetBackup media server
/usr/opensv/netbackup/logs/bpcd	NetBackup backup host client
/usr/opensv/netbackup/logs/bppfi	NetBackup backup host client
/usr/opensv/logs/ncfnbcs	NetBackup backup host client

Table 13-4 Linux logging directories for single file restore from a snapshot copy

Path of the log directory	Where to create the directory?
/usr/opensv/netbackup/logs/bprd	NetBackup primary server
/usr/opensv/logs/bpbrm	NetBackup primary server
/usr/opensv/netbackup/logs/bpcd	Restore host client
/usr/opensv/netbackup/logs/bpbkar	Restore host client
/usr/opensv/netbackup/logs/bpfis	Restore host client
/usr/opensv/netbackup/logs/bppfi	Restore host client
/usr/opensv/logs/tar	Destination client where the files are restored.

Table 13-5 Linux logging directories for a point-in-time rollback

Path of the log directory	Where to create the directory?
/usr/opensv/netbackup/logs/bprd	NetBackup primary server

Table 13-5 Linux logging directories for a point-in-time rollback *(continued)*

Path of the log directory	Where to create the directory?
/usr/opensv/netbackup/logs/bpbm	NetBackup primary server
/usr/opensv/netbackup/logs/bpcd	Restore host client
/usr/opensv/netbackup/logs/bpbkar	Restore host client
/usr/opensv/netbackup/logs/bpfis	Restore host client
/usr/opensv/netbackup/logs/bppfi	Restore host client

Table 13-6 Linux logging directories for a create replication operation

Path of the log directory	Where to create the directory?
/usr/opensv/logs/nbjm	NetBackup primary server
/usr/opensv/logs/nbstserv	NetBackup primary server
/usr/opensv/logs/nbrb	NetBackup primary server
/usr/opensv/netbackup/logs/bpdm	NetBackup media server

Table 13-7 Linux logging directories for a delete replication operation

Path of the log directory	Where to create the directory?
/usr/opensv/netbackup/logs/bpdm	NetBackup media server
/usr/opensv/netbackup/logs/admin	NetBackup media server (for bppficorr logs)

Logging folders for Windows platforms

Table 13-8 Windows logging directories for a snapshot operation

Path of the log directory	Where to create the directory?
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\nbjm	NetBackup primary server

Table 13-8 Windows logging directories for a snapshot operation (*continued*)

Path of the log directory	Where to create the directory?
install_path\NetBackup\logs\bpbrm	NetBackup primary server if Instant Recovery backup is set to snapshot only; otherwise, on media server
install_path\NetBackup\logs\bpfis	Backup host client

Table 13-9 Windows logging directories for backup operation

Path of the log directory	Where to create the directory?
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\nbjm	NetBackup primary server
install_path\NetBackup\logs\nbstserv	NetBackup primary server
install_path\NetBackup\logs\bpdbm	NetBackup primary server
install_path\NetBackup\logs\bptm	NetBackup primary server
install_path\NetBackup\logs\bpbrm	NetBackup primary server
install_path\NetBackup\logs\bpfis	Backup host client
install_path\NetBackup\logs\bppfi	Backup host client
install_path\NetBackup\logs\bpbkarr	Backup host client
install_path\NetBackup\logs\ncfnbcs	Backup host client

Table 13-10 Windows logging directories for index from snapshot operation

Path of the log directory	Where to create the directory?
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\nbjm	NetBackup primary server
install_path\NetBackup\logs\bpdbm	NetBackup primary server
install_path\NetBackup\logs\bptm	NetBackup primary server

Table 13-10 Windows logging directories for index from snapshot operation
(continued)

Path of the log directory	Where to create the directory?
install_path\NetBackup\logs\bpbrm	NetBackup primary server
install_path\NetBackup\logs\bpcd	Backup host client
install_path\NetBackup\logs\bppfi	Backup host client
install_path\NetBackup\logs\ncfnbcs	Backup host client

Table 13-11 Windows logging directories for single file restore from snapshot copy

Path of the log directory	Where to create the directory?
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\bpbrm	NetBackup primary server
install_path\NetBackup\logs\bpcd	Remote host client
install_path\NetBackup\logs\bpbkar	Remote host client
install_path\NetBackup\logs\bpfis	Remote host client
install_path\NetBackup\logs\bppfi	Remote host client
install_path\NetBackup\logs\tar	Destination client where the files are restored.

Table 13-12 Windows logging directories for single file restore from point-in-time rollback

Path of the log directory	Where to create the directory?
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\bpbrm	NetBackup primary server
install_path\NetBackup\logs\bpcd	Remote host client
install_path\NetBackup\logs\bpbkar	Remote host client
install_path\NetBackup\logs\bpfis	Remote host client

Table 13-12 Windows logging directories for single file restore from point-in-time rollback *(continued)*

Path of the log directory	Where to create the directory?
<code>install_path\NetBackup\logs\bppfi</code>	Remote host client

Table 13-13 Windows logging directories for single file restore from create replication operation

Path of the log directory	Where to create the directory?
<code>install_path\NetBackup\logs\nbjm</code>	NetBackup primary server
<code>install_path\NetBackup\logs\nbstserv</code>	NetBackup primary server
<code>install_path\NetBackup\logs\nbrb</code>	NetBackup primary server Remote host client
<code>install_path\NetBackup\logs\bpdm</code>	NetBackup media server

Table 13-14 Windows logging directories for single file restore from delete replication operation

Path of the log directory	Where to create the directory?
<code>install_path\NetBackup\logs\bpdm</code>	NetBackup media server
<code>install_path\NetBackup\logs\admin</code>	NetBackup media server (for <code>bpffcorr</code> logs)

Logging folders for multi-stream restore

Table 13-15 Logging directories for multi-stream restore

Operation	VxUL logs	Non-VxUL logs	Hosts
Recovery API	<code>nbwebsservice</code>		Primary server
Recovery backend		<code>bprd</code> on primary server, <code>bpbrm</code> on media server, and <code>tar</code> on the client	Primary server, media server, and client

Exclude list is not working during backup

Explanation:

The exclude list may not work if the proper guidelines for creating the exclude lists are not followed.

Workaround:

Refer to the following guidelines to configure the exclude lists properly:

- See [“Configuring exclude lists”](#) on page 73.
- [Excluding files from backup.](#)

Restore from a snapshot fails with status 133

Restore from snapshot fails with status code 133 and displays the Invalid request message.

Explanation

The restore fails if you select a path other than the path mentioned in the backup selection.

For example, say that the backup selection contains `/ifs/voll/parent/dir1`. During a restore if you select only `/ifs/voll/parent`, which is the parent directory of the path mentioned for backup selection, the restore fails with status code 133.

Workaround

For a successful restore from the snapshot copy, you must select the original path mentioned in the **Backup selections** tab; that is `/ifs/voll/parent/dir1` or the subdirectory or file inside the backup selection.

Backup from snapshot jobs do not start after the snapshot job completes successfully

Explanation:

If the NAS-Data-Protection policy backup jobs fail for some backup IDs, then the image or SLP remains incomplete. This results in the backup from snapshot jobs not starting until the incomplete backup IDs are canceled.

Workaround:

Do the following:

- 1 To show the status of incomplete images, run:
`nbstlutil stlilist -image_incomplete`
- 2 To find information for an incomplete lifecycle image, run:
`nbstlutil stlilist -U`
- 3 To find the backup ID, refer to the following article:
<https://support.cohesity.com/s/article/article-100016129>
- 4 To cancel the lifecycle for a specific backup, run:
`nbstlutil cancel -backupid <backupid>`

Backup from snapshot fails with error 50

This error occurs when the NetBackup Client and NetBackup Legacy Network services are not restarted properly after configuration for the domain user.

Workaround

If you use primary or media as a backup host then follow these steps to troubleshoot:

- 1 Stop all NetBackup services using the `bpdown.exe`.
- 2 Log on to the NetBackup Client and NetBackup Legacy Network services as the domain user. But, do not start these services immediately after logon.
- 3 Start all the services together using `bpup.exe`.

Backup from snapshot parent job fails with error 4213: Snapshot import failed

The Job details section in the UI shows an error like:

"Snapshot export failed. Failed to export share: data_lif is not online. Check the data_lif status on vservers: VSERVER_1."

Where, VSERVER_1 is the vserver that is offline.

Explanation:

For a NAS-Data-Protection policy, all the vservers are listed in the client's section of the policy, irrespective of their state. So, you can include backup selection from offline SVM, and policy validation succeeds. However, the backup-from-snapshot export operation for those shares fails if the corresponding vserver is offline.

Workaround

To overcome this error, check the status of the vserver. The vserver must be reachable and the SLP retries must be successful.

Backup host pool creation fails with the error "Failed to fetch host list"

Explanation:

This issue appears if the NetBackup services are not started properly with the domain user.

Workaround:

- 1 Make sure that the NetBackup client service is running.
- 2 Log on as the domain user to the NetBackup client service.
- 3 Restart the NetBackup Client service.
- 4 Make sure that the NetBackup Legacy Network service is running.
- 5 Log on as the domain user to the NetBackup Legacy Network service.
- 6 Restart the NetBackup Legacy Network service.
- 7 Make sure that all NetBackup services are running.
- 8 Relaunch the NetBackup UI.

Snapshot job fails and the snapshot command does not recognize the volume name

Explanation:

A snapshot job fails if the volume name exceeds 15 characters.

When you create and name a volume, a prefix or a suffix is added to the volume name. If the volume name contains more than 15 characters, the addition of a prefix or suffix may make the volume name exceed the limit of 27 characters. When you run the `vxassist snapshot` command, it does not recognize the lengthy snapshot volume name and the snapshot job fails.

For example, if the primary volume name is **PFltest123456789vol** and the suffix **00043c8aaa** is added to it, the volume name exceeds the limit. The command `vxassist snapshot` does not recognize the name **PFltest123456789vol_00043c8aaa** and the snapshot job fails.

Workaround:

Cohesity recommended that you limit the primary volume names to up to 15 characters to create the VxVM mirror snapshots.

Accelerator enabled incremental backup of NetApp NAS volume

Accelerator enabled NAS-Data-Protection policy backups complete volume instead of only the incremental data. This also affects the run optimization.

This issue occurs under the following conditions:

- The policy type is NAS-Data-Protection.
- In the policy's Snapshot options, the value of Access Protocol is Default or NFS3.
- Backup selection has NetApp NAS volumes.

The Accelerator technology optimizes a backup by sending only changed blocks over a network for backup. A two-step process is used to identify the changed files and changed blocks in these files. File attributes and index node (inode) are the key parameters to identify a change. If the files are accessed over NFS version 3, a file on a NetApp NAS volume behaves differently because of the inode numbers. The same file has different inode numbers across snapshots of the volume if accessed over NFS3. All schedules of backup are based on the snapshot that is created for the run of the policy. New snapshots with different inode numbers than the previous ones help the accelerator to identify these files as new files. Because of this issue, all files are backed up instead of incremental data only.

To resolve this issue, avoid using NFS version 3 to access the snapshot for accelerator-enabled backups. You can change the Access Protocol to NFS4 for the affected policy. For more details, refer to the [NetApp documentation](#).

Snapshot method: Auto

Error scenario 1: Policy validation fails after a primary server upgrade, if you create a policy with VSO FIM for older clients and select Snapshot Method as Auto in the NetBackup 10.0 UI.

Error scenario 2: Snapshot jobs fail if you configure D-NAS policy with a backup host pool containing older version backup hosts and select Snapshot Method as Auto in the NetBackup 10.0 UI.

The Snapshot Method, Auto is supported only in NetBackup 10.0 onwards. If your environment contains older version backup hosts, select another snapshot method.

Backup from snapshot jobs for NAS-Data-Protection policy fail with error 4213

Backup from snapshot jobs for NAS-Data-Protection policy fail with error 4213.

```
---Activity monitor detailed status--- Oct 13, 2022 12:44:00 PM -
end SnapDupe Mount:Read File List; elapsed time 0:00:00 Oct 13, 2022
12:44:00 PM - begin SnapDupe Mount:Import Snapshot Oct 13, 2022
12:44:00 PM - Info RUNCMD (pid=13508) started Oct 13, 2022 12:44:14
PM - Info RUNCMD (pid=13508) exiting Operation Status: 4213 Oct 13,
2022 12:44:14 PM - end SnapDupe Mount:Import Snapshot; elapsed time
0:00:14 Oct 13, 2022 12:44:14 PM - Error nbjm (pid=3792)
ImportSnapshot failed, snapshotid=10.84.69.235@dsemc02dm_1665644972
Operation Status: 4213 Oct 13, 2022 12:44:14 PM - end Parent Job;
elapsed time 0:00:14 Snapshot import failed (4213)
```

Explanation:

This issue occurs if any one of your backup hosts in the backup host pool is at a lower version than 10.1.1, and the protected NAS volumes reside on Dell EMC Unity, Dell EMC PowerStore, or Hitachi NAS storage array.

Workaround:

Remove the backup hosts from the backup host pool that have a lower NetBackup version than 10.1.1. Alternatively, for these policies, use a different backup host pool that has only NetBackup 10.1.1 hosts.

A full VCT-enabled indexing job runs, when followed by a non-VCT indexing job with a backup host prior to version to 10.3

Workaround:

Do the following:

- For a VCT-enabled policy to support index from snapshot operation, ensure that the policy, primary server, media server, and backup hosts are of NetBackup version 10.3 or higher.

- For a non-VCT index operation, if you use a backup host prior to version 10.3 earlier; before running a VCT-based indexing job, run the non-VCT index job with full schedule using a NetBackup version 10.3 or higher client.

Backup from snapshot jobs for NAS data protection policy fail with error 927

Explanation:

This issue occurs if the backup host pool does not contain a host that is of the same or lower version of NetBackup than the media server.

Workaround:

Ensure that all media servers associated with the storage unit, specified in the Storage Lifecycle Policy (SLP), have a higher version of NetBackup than the lowest version of the backup host in the backup host pool.

To exclude a media server, go to the storage unit properties for the STU specified in the SLP. Select the **Only use the following media servers** option. Then select the media servers with a NetBackup version higher or equal to the lowest NetBackup version of the hosts in the backup host pool.

Error code: 930: No supported media server is available in the All_Media_Server_Pool to use to backup the NAS shares.

Explanation:

No media server with the supported NetBackup version and operating system is available. The server(s) may be busy, or down.

Workaround:

Ensure that you have:

- Windows media server(s) if you back up the NAS shares using the SMB protocol.
- Linux media server(s) if you back up the NAS shares using NFS protocol.

For Windows media servers, you must log on to the following services as a domain user:

- NetBackup Client Service
- NetBackup Legacy Network Service

Restore from NAS array volume fails with the status: 174 Media manager – system error occurred.

Ensure that the NetBackup version of the media server is equal to or higher than the minimum required version, based on these parameters of the NAS data protection policy. See [“Minimum supported backup host versions for different features”](#) on page 34.

Restore from NAS array volume fails with the status: 174 Media manager – system error occurred.

Explanation:

This error occurs during granular restoration of a specific combination of files and directories that were backed up from a NAS array.

- For NetBackup version 10.4, this issue occurs only for multi-stream restores.
- For NetBackup version 10.5 and above, this issue occurs for both single-stream and multi-stream restores.

Workaround:

Do any of the following:

- For NetBackup version 10.4, use a single stream for the restore job.
- For NetBackup version 10.5 and above, do only one selection at a time, in the restore selection, like a volume, a directory, or individual file(s).

NAS job fails with the error: Crawler process timed out after 600 seconds waiting for streams to attach with shared memory.

The job details has entries similar to this:

```
Info nbjm (pid=30970) Started child jobs 278, 279, 280, 281, 282,
283, 284, 285, 286, 287 on host Host1.domain.com
Info nbjm (pid=30970) Started child jobs 288, 289, 290, 291, 292,
293, 294, 295, 296, 297 on host Host2.domain.com
Error nbcs (pid=790643) Crawler process timed out after 600 seconds
waiting for streams to attach with shared memory
Info nbjm (pid=30970) Started child jobs 298, 299, 300, 301, 302,
303, 304, 305, 306, 307 on host Host3.domain.com
```

Explanation:

D-NAS backup fails with the error: The file system crawler process timed-out waiting for streams to attach with shared memory. (3003)

This error is encountered when the parent job initiates and the child stream jobs are queued sequentially. As a result, by the time all child stream jobs begin, the nbcs crawler process reaches its timeout period, leading to a failure in job execution.

Workaround:

If your scheduled configurations cause such timeout you can change the timeout value by using the configuration parameter

`DYNAMIC_STREAMING_START_CHILD_BACKUP_JOBS_TIMEOUT`. You can change the value of this variable by using the `bpsetconfig` command. Use the `bpgetconfig` CLI to view the value of this variable. You can set this configuration parameter on the NetBackup primary server. For more information, see *NetBackup™ Administrator's Guide, Volume I*.

D-NAS backup fails with the error: The file system crawler process timed-out waiting for streams to attach with shared memory. (3003)

Explanation:

The job fails after 20 minutes of waiting for the resources to be allocated, as the crawler is configured to time-out after 1200 seconds.

Workaround:

To resolve the issue, increase the Maximum concurrent jobs for the storage unit.

- 1 On the left, click **Storage units**, under **Storage**.
- 2 Click the storage unit that you want to edit.
- 3 Under **Basic properties**, click **Edit**. Enter a new value for the **Maximum concurrent jobs** parameter.

Isilon backup from snapshot failed with the Snapshot cannot be mounted error.

Explanation:

The required privileges are not assigned to the domain user for the SMB share backup.

Workaround:

To add privileges to the user:

- 1 Log on to the Windows host.
- 2 To open the local security policy settings, click Start, and enter: `secpol.msc`.
- 3 Add the service account (domain\username) to the following:
 - Act as part of the operating system
 - Adjust memory quotas for a process
 - Replace a process level token
- 4 Add the service account (domain\username) to the local administrators group.
- 5 Restart the NetBackup Legacy Network Service.

Discovery and snapshot operations fail with the errors 156 and 1542

Explanation:

Occurs when the NetBackup Snapshot Manager for Data Center services are not running properly.

Workaround:

You need to restart the services in the NetBackup Snapshot Manager for Data Center correctly so that your environmental data is preserved.

Run the following command using the flexsnap_configure CLI:

```
# flexsnap_configure restart
```

Retry the operation.

Using NDMP

- [Chapter 14. Introduction to NetBackup for NDMP](#)
- [Chapter 15. Installation Notes for NetBackup for NDMP](#)
- [Chapter 16. Configuring NDMP backup to NDMP-attached devices](#)
- [Chapter 17. Configuring NDMP backup to NetBackup media servers \(remote NDMP\)](#)
- [Chapter 18. Configuring NDMP DirectCopy](#)
- [Chapter 19. Accelerator for NDMP](#)
- [Chapter 20. Remote NDMP and disk devices](#)
- [Chapter 21. Using the Shared Storage Option \(SSO\) with NetBackup for NDMP](#)
- [Chapter 22. NAS appliance information for NDMP](#)
- [Chapter 23. Backup and restore procedures](#)
- [Chapter 24. Troubleshooting](#)
- [Chapter 25. Using NetBackup for NDMP scripts](#)

Introduction to NetBackup for NDMP

This chapter includes the following topics:

- [About NetBackup for NDMP](#)
- [About Network Data Management Protocol \(NDMP\)](#)
- [Types of NDMP backup](#)
- [About NDMP policies in NetBackup](#)
- [About NetBackup storage units](#)
- [About assigning tape drives to different hosts](#)
- [About the NDMP backup process](#)
- [About the NDMP restore process](#)
- [About Direct Access Recovery \(DAR\)](#)
- [Snapshot Client assistance](#)
- [About NDMP multiplexing](#)
- [About NDMP support for Replication Director](#)
- [Limitations of Replication Director with NDMP](#)
- [About NDMP support for NetApp clustered Data ONTAP \(cDOT\)](#)

About NetBackup for NDMP

NetBackup for NDMP is an optional NetBackup application. It enables NetBackup to use the Network Data Management Protocol (NDMP) to initiate and control backups and restores of Network Attached Storage (NAS) systems.

NetBackup for NDMP features

The following table describes the NetBackup for NDMP features.

Table 14-1 NetBackup for NDMP features

Feature	Description
Support for NDMP protocol	Supports the NDMP protocol versions V2, V3, and V4.
Centralized backup policy management	Scheduling, catalog management, and other backup tasks are managed from a NetBackup primary server. NetBackup for NDMP can be installed on a NetBackup primary or media server.
Accelerator for NDMP	NetBackup's Accelerator option makes NDMP backups for NetApp and Isilon filers run faster than normal NDMP backups. NetBackup Accelerator increases the speed of full backups by using the filer's change detection techniques to identify the modifications that occurred since the last backup. More information about the feature is available: See "About NetBackup Accelerator for NDMP" on page 181.
Support for NetApp cDOT filers	NetBackup for NDMP supports NetApp clustered Data ONTAP (cDOT) filers. More information about configuring NetBackup to work with NetApp cDOT filers is available: See "Configure an NDMP filer" on page 153.
Support for wildcards in NDMP backup policy selections	Wildcard characters in regular expressions or directives are valid for streaming and non-streaming NDMP backups.
Device and media management	NetBackup software provides complete management and control of the devices and media that are used for backups and restores of NDMP hosts. The NetBackup web UI discovers and configures the storage devices that are attached to an NDMP host. (This function requires NDMP protocol V3 or V4.) Note that wizard-based discovery depends upon a number of device-specific features, such as SCSI inquiry and serialization, which some NAS vendors may not support.
High-speed local backup of NDMP hosts	Backup data travels between the disk drives and tape drives that are directly attached to the same NDMP host. This transfer provides high-speed backup but does not slow network throughput.

Table 14-1 NetBackup for NDMP features (*continued*)

Feature	Description
Backup of network-attached NDMP hosts to a tape device on another NDMP host or to advanced tape libraries with an embedded NDMP server	Backup data travels across the network from a disk on an NDMP host to tape on another NDMP host. This backup is referred to as a three-way backup. This data movement option requires support from the NAS/NDMP host.
Backup of a network-attached NDMP host to a tape device on a NetBackup media server	Backup data travels across the network from a disk on an NDMP host to tape on a NetBackup media server. This backup is a form of three-way backup also known as remote NDMP. This feature supports NDMP V2, V3, and V4 on the NDMP hosts.
Shared tape libraries	Tape libraries can be shared between NDMP hosts and NetBackup servers or between multiple NDMP hosts. Robotic control can be on an NDMP host or on a NetBackup server.
Shared tape drives with the Shared Storage Option	Tape drives can be shared between servers (both NetBackup servers and NDMP hosts). This setup requires the Shared Storage Option (SSO) license. For a list of the features and software releases for each NAS vendor, for SSO support, and for the NetBackup versions that support these vendors, see the NetBackup Compatibility List.
Snapshots of data on NDMP hosts	NetBackup can take point-in-time data snapshots on an NDMP (NAS) host without interrupting client access to data, using the NDMP V4 snapshot extension. The snapshot is stored on the same device that contains the NDMP client data. From the snapshot, you can restore individual files or roll back a file system or volume by means of Snapshot Client Instant Recovery. A NetBackup Snapshot Client license is required, in addition to the NetBackup for NDMP license. This Snapshot Client feature uses the NAS_Snapshot method and the NDMP method. For more information about the NDMP snapshot method, refer to the NetBackup Replication Director Solutions Guide
NDMP DirectCopy	NetBackup can copy virtual tape library (VTL) images directly from the VTL to physical tape or to another VTL. This function occurs without using media server I/O resources or network bandwidth. NetBackup can copy NDMP backup images directly from one NDMP-attached tape drive to another NDMP tape drive that is attached to the same NDMP host. Note that the operation does not use media server I/O. Note: The VTL must have an embedded NDMP tape server.
Direct Access Recovery (DAR)	For NDMP hosts that support DAR, this feature greatly reduces the time to restore a directory, a single file, or a small number of files.

Table 14-1 NetBackup for NDMP features (*continued*)

Feature	Description
Path-based file history	The NDMP server can send catalog information consisting of complete path names to NetBackup. Some vendors do not support this feature. Up-to-date information is available on the vendors that support path-based history. For a list of the features and software releases for each NAS vendor, for SSO support, and for the NetBackup versions that support these vendors, see the NetBackup Compatibility List .
Support for NetBackup for NDMP servers in a NetBackup-clustered environment	The NetBackup for NDMP servers are supported in a NetBackup-clustered environment.
Enhanced ability to run customized scripts during a backup	The enhanced ability to run customized scripts during a backup, especially for relational databases residing on NAS devices.
NDMP multiplexing	NDMP multiplexing enables NDMP backups to be multiplexed to Media Manager storage units. Only remote NDMP multiplexing is supported.
NDMP to disk	NetBackup can write NDMP backups to disk storage units.
IPv6 support	NDMP supports 128-bit IPv6 address data connections in addition to the 32-bit IPv4 address data connections. NDMP data connections are made between filers or between a NetBackup media server and a filer that is used to transfer the backup image. By default the NetBackup media server is enabled for IPv6 data communication. Consider the following general items when using NDMP IPv6 address data connections. ■ The filer needs to be enabled for IPv6 data communication. ■ The filer vendor must support connection address extension or full IPv6.
NDMP support for Replication Director	NDMP support for Replication Director enables NetBackup to use NDMP for the following functions: backup from snapshots, restore from snapshot backups, live browse snapshots, and restore from snapshots (for copy back method). For more information about Replication Director, refer to the NetBackup Replication Director Solutions Guide.

NetBackup for NDMP terminology

The following table describes the NetBackup for NDMP terminology. For explanations of other NetBackup terms, consult the NetBackup online glossary in NetBackup Help.

Table 14-2 Terminology

Term	Definition
DAR (Direct Access Recovery)	DAR is an optional capability of NDMP data and tape services where only relevant portions of the secondary media are accessed during recovery operations. The NDMP host positions the tape to the exact location of the requested file(s), reading only the data that is needed for those files. Restore times can be reduced from hours to minutes.
NDMP (Network Data Management Protocol)	NDMP is a widely used protocol through which an NDMP-conformant backup application can control the backups and restores for an NDMP host.
NDMP client	An NDMP client is an NDMP-compliant backup application (also known as a Data Management Application or DMA) that is an NDMP server application client. An NDMP client sends commands to the NDMP server application to control the backups and restores on an NDMP host. NetBackup for NDMP allows NetBackup to act as an NDMP client.
NetBackup for NDMP server	A NetBackup for NDMP server is a NetBackup primary or media server on which NetBackup for NDMP software is installed.
NDMP host	An NAS system that serves files to clients using HTTP, FTP, CIFS, or NFS protocols. It also runs an NDMP server application that communicates with NDMP client backup software to configure and perform backup and restore tasks. NAS systems provide fast, multi-protocol file access and cost effective data storage to workstations and servers in the network or across the Internet. In a NetBackup configuration, the NDMP host is considered a client of NetBackup. However, NetBackup client software is never installed on an NDMP host.
NDMP multiplexing	NDMP multiplexing concurrently writes multiple backup streams to the same Media Manager tape storage device from the same client or different clients. NDMP multiplexing improves overall NetBackup performance by more efficient use of the storage unit drives. State of the art storage devices can typically stream data faster than client agents can create backup streams. Therefore, multiple data streams can be sent to and effectively processed by a given storage unit. Only remote NDMP multiplexing is supported.

Table 14-2 Terminology (*continued*)

Term	Definition
NDMP server application	An NDMP server application runs on an NDMP host and runs backup, restore, and device control commands that it receives from an NDMP-conformant backup application. The backup application (NetBackup) is considered an NDMP client. A separate instance of an NDMP server process exists for each connection to an NDMP client. That is, if two backups are in progress, an NDMP server process exists for each backup.
NDMP storage unit	An NDMP storage unit stores the backup data for an NDMP host. The tape drives in this storage unit attach directly to the NDMP host or can be configured on a SAN. Note that NDMP storage units cannot be used to store data for non-NDMP hosts, and NetBackup disk storage units cannot be used for NDMP tasks.
Redirected restore (to a different client)	In a redirected restore, files are restored to a client other than the one from which they were originally backed up. In NetBackup for NDMP, the restore data travels from an NDMP host (or NetBackup media server) with a locally attached storage device to another NDMP host on the network.
Remote NDMP	A form of three-way backup and restore also known as NDMP backup to Media Manager storage units. Data travels from an NDMP host to a tape drive that is attached to a NetBackup media server. See “Configuring NDMP backup to Media Manager storage units” on page 172.
Three-way backup and restore	In a three-way backup or restore, data travels between an NDMP host and a storage device that is attached to another NDMP host or to a NetBackup media server. This backup contrasts with local NDMP backup or restore where the data travels between an NDMP host’s disk and a storage device directly attached to the same NDMP host.
Virtual Tape Library (VTL)	A virtual tape library is a storage system that uses disk-based technology to emulate a tape library and tape drives. For secondary storage, NetBackup can copy VTL images directly to a physical tape or to another VTL by means of NDMP DirectCopy.

About Network Data Management Protocol (NDMP)


NDMP is a widely used protocol through which an NDMP-conformant backup application controls the backups and restores of any NDMP host that runs an NDMP server application.

NDMP architecture follows the client and server model:

- The NetBackup primary or media server where NetBackup for NDMP is installed is called a NetBackup for NDMP server.
- The host where the NDMP server application resides is called an NDMP host.
- The NetBackup software is a client of the NDMP server application. NetBackup for NDMP lets NetBackup act as an NDMP client. The NDMP hosts, on the other hand, act as NetBackup clients.

The following figure shows an example of NDMP and NetBackup hosts as clients of each other.

Figure 14-1 NDMP and NetBackup hosts as clients of each other



Types of NDMP backup

The NDMP server application on the NDMP host performs backups and restores of the NDMP host, directed by commands from an NDMP client (NetBackup). Backups can be conducted in any of the following ways:

- NDMP local backup
See [“NDMP local backup”](#) on page 127.


- NDMP three-way backup
See “[NDMP three-way backup](#)” on page 127.
- Backup to a Media Manager storage unit on the NetBackup server
See “[Backup to Media Manager storage units \(remote NDMP\)](#)” on page 128.

NDMP local backup

If you use the NDMP local backup, the NetBackup for NDMP server initiates the backup. The data travels from the NDMP host's disk to a storage device that is attached to the same host or is available on a SAN.

The following figure shows an example of an NDMP local backup and restore.

Figure 14-2 NDMP local backup and restore



Local NDMP backup

Data travels from disk to tape on same NDMP host, or from disk to tape device on SAN. *Backup data is NOT sent over local network.*


The tape drives must be in NDMP-type storage units.

NDMP three-way backup

If you use the NDMP three-way backup, the NetBackup for NDMP server initiates the backup. Data travels over the network by going from an NDMP host to a storage device that is attached to another NDMP host on the local network or is available on a SAN.

The following figure shows an example of an NDMP three-way backup and restore.

Figure 14-3 NDMP three-way backup and restore



Three-Way NDMP backup


Data travels from disk on an NDMP host to tape device on another NDMP host. *Backup data is sent over the local network.*

The tape drives must be in NDMP-type storage units.

Backup to Media Manager storage units (remote NDMP)

With this backup method, the data travels over the network by going from an NDMP host to a Media Manager-type storage device that is attached to a NetBackup media server or is available on the SAN. The NetBackup drives must be in Media Manager storage units not in NDMP storage units.

The following figure shows an example of an NDMP backup to a Media Manager device (remote NDMP).

Figure 14-4 NDMP backup to a media manager device (remote NDMP)

To NetBackup Server-Attached Media Manager Storage Units

Data travels from NDMP host to a drive on a NetBackup media server or on a SAN. *Backup data is sent over the local network.*

NOTE: The NetBackup drive(s) must be in Media Manager type storage units.

About NDMP policies in NetBackup

After you install and configure NetBackup for NDMP, you can schedule backups by creating an NDMP policy in NetBackup.

An NDMP policy can have one or more NetBackup clients. Each NetBackup client must be an NDMP host.

See [Figure 14-1](#) on page 126.

Note that you do not install any NetBackup software on the NDMP hosts.

The allowable backup types for schedules in an NDMP policy are: Full, Cumulative Incremental, or Differential Incremental. User-initiated backups and archives are not allowed because the NDMP protocol does not permit these tasks.

Restores of NDMP host backups can be initiated from any NetBackup media server that meets the following criteria:

- Resides within the same overall NetBackup storage domain
- Uses the same NetBackup primary server that the media server uses to perform the backup

The data can be restored to the NDMP host where it was backed up, or to another NDMP host.

NDMP policies can use either NDMP storage units or Media Manager storage units.

About NetBackup storage units

NetBackup uses either one of the following storage units:

- NDMP-type storage units (for local or three-way backup)

NetBackup requires NDMP-type storage units when you back up NDMP host data to the devices that are as follows:

- Attached to an NDMP host
- Available to the NDMP host on a SAN

An NDMP storage unit can contain standalone or robotic drives. Robotic controls can be in a TLD (tape library DLT) or ACS robot type.

- Media Manager storage units (for backup to devices that are attached to a NetBackup media server)

You can use the drives that were configured in Media Manager-type storage units when you back up NDMP host data to devices that are as follows:

- Attached to a NetBackup for NDMP server
- Available to the server on a SAN

For NDMP backup, drives in Media Manager-type storage units do not have to be dedicated to NDMP data. They can store backups of regular (non-NDMP) NetBackup clients as well as of NDMP clients.

About assigning tape drives to different hosts


The robotic tape drives can be divided up among NDMP hosts and NetBackup servers.

The following figure shows the NDMP and non-NDMP storage units with the following configuration:

- Tape drives 1, 3, and 5 are attached to NDMP hosts. They are in the NDMP storage units that can be used for NDMP backups (local or three-way). The commands that control these drives originate on the NetBackup for NDMP server and are sent through the NDMP connection on the network. The NDMP server application on each NDMP host translates the NDMP commands into SCSI commands for the local drives.

- Tape drives 2 and 4 are attached to a NetBackup server. They are in non-NDMP storage units and are controlled in the same way as other drives on NetBackup servers. Depending on the type of storage unit, these drives can be used for the following:
 - Non-NDMP clients of NetBackup
 - In the case of tape drives in Media Manager storage units, they can be used for both NDMP (local or three-way) and non-NDMP backups.
- In the following figure, all of the tape drives are used for NDMP backup except drive 4.

Figure 14-5 NDMP and non-NDMP storage units



- * In NDMP storage unit
- † In NetBackup Media Manager storage unit
- ‡ In another type of NetBackup storage unit (not NDMP or Media Manager)

Drives 1, 3, and 5 (in NDMP storage units) can be used for NDMP backups.

Drive 2 (in Media Manager storage unit) can be used for NDMP or non-NDMP backup.


Drive 4 (in different type of NetBackup storage unit) cannot be used for NDMP backup.

About robotics control

Robotics control can be attached to an NDMP host or to a NetBackup server.


The following figure shows how NetBackup sends commands over the network to the NDMP host, which in turn sends them to the robot.

Figure 14-6 Robotics control that is attached to an NDMP host



The following figure shows how the robot is controlled in the same way as the other robots on NetBackup servers.

Figure 14-7 Robotics control that is attached to a NetBackup server




About the NDMP backup process

During a backup, the following events occur in this order:

- From the Enterprise Media Manager (EMM), NetBackup obtains a media ID for the tape that is used for the backup. It then sends a tape-mount request to `ltid`.
- `ltid` on the NetBackup for NDMP server sends the necessary NDMP (SCSI robotic) commands to mount the requested tape on the storage device.
- NetBackup sends the NDMP commands that are necessary to have the NDMP server application perform a backup to the tape. The backup data travels in one of two ways:
 - Between the local disk and tape drives on an NDMP host.
 - Over the network, data travels from an NDMP host without its own storage device to an NDMP host (or NetBackup media server) with a locally attached storage device (three-way back up).
- The NDMP server application sends information to the NetBackup for NDMP server about the files that were backed up. This information is stored in the NetBackup file database.
- The NDMP server application sends status about the backup operation to the NetBackup for NDMP server.

The following figure shows the NetBackup processes that are involved in the NDMP backups.

Figure 14-8 NetBackup backup processes



About the NDMP restore process


Because of the design of the NDMP protocol, only an administrator on a NetBackup server (primary or media) can restore files from NDMP backups. During a restore,

the administrator browses the file catalog and selects files from NDMP images in the same manner as for standard backup images.

The following events occur during a restore, in this order:

- The NetBackup for NDMP server looks in its Enterprise Media Manager (EMM) database for the tape that contains the backup, and asks `ltid` to mount that tape.
- `ltid` on the NetBackup for NDMP server sends the necessary NDMP commands to load the requested tape on the storage device.
- NetBackup sends the NDMP commands that are necessary to have the NDMP server application perform a restore operation to the disk. The restore data travels in one of two ways:
 - From a tape drive to a local disk (tape drive and disk are on the same NDMP host)
 - Over the network, from an NDMP host (or NetBackup media server) with a locally attached storage device to another NDMP host (three-way backups or restores)
- The NDMP server application sends status about the restore operation to the NetBackup for NDMP server.

The following figure shows the NetBackup processes involved in NDMP restores.

Figure 14-9 NetBackup restore processes

About Direct Access Recovery (DAR)

NetBackup uses Direct Access Recovery (DAR) to restore a directory or individual files from a backup image. DAR can greatly reduce the time it takes to restore files and directories. DAR is enabled by default. No configuration is required.

DAR enables the NDMP host to position the tape to the exact location of the requested files. It reads only the data that is needed for those files. For individual file restore, NetBackup automatically determines whether DAR shortens the duration of the restore. It activates DAR only when it results in a faster restore.

The following prerequisites are necessary for using DAR with NetBackup for NDMP:

- The NDMP host must support DAR where the NDMP server application resides.
- NetBackup 4.5 GA or later, with the catalog in binary format (binary format is the default).

Further details are available as to when DAR is used and how to disable it.

See [“About enabling or disabling DAR”](#) on page 167.

Snapshot Client assistance

The Snapshot Client Configuration document includes the following information:

- An up-to-date list of supported operating systems and peripherals
- A list of NAS vendors that are supported for the NAS_Snapshot method
- Sections on SAN device configuration and on setting up NetBackup for off-host data mover backups (including instructions on creating `3pc.conf` and `mover.conf` files).

About NDMP multiplexing

NDMP multiplexing concurrently writes multiple backup streams to the same tape storage device from the same client or different clients. NDMP multiplexing supports only remote NDMP and improves overall NetBackup performance by better using tape storage devices. State-of-the-art tape storage devices can typically stream data faster than client agents can create backup streams. Therefore multiple data streams can be sent to and effectively processed by a given tape storage unit.

A network-attached storage (NAS) device with an NDMP server is an agent that produces a backup stream that is similar to a NetBackup client. Multiplexing is desired for NDMP backups because NAS devices are limited in the rate at which they create backup streams. These backup streams are often much slower than the tape storage device consuming and writing the stream.

NDMP multiplexing provides the following benefits:

- Several backups can be run at the same time writing to the same tape. This process can reduce the need for many tape devices.
- Backup time is reduced by writing concurrent backups to a single tape storage device.
- Many tape storage devices require that data is streamed to them at high transfer rates. When data is not streamed fast enough, they do not work efficiently and are subject to possible excessive wear.

Consider the following general items when implementing NDMP multiplexing:

- Only media manager tape storage units can be used for NDMP multiplexing.
- Multiplexing of NDMP backups and restores supports only remote NDMP. The remote NDMP processes backup streams by going through the media server.
- NDMP local and NDMP three-way backups and restores are not supported for NDMP multiplexing. Each method processes backup streams without going through the media server.
- Synthetic backups are not supported.
- Only tape devices are supported.
- Disk storage devices are not supported.
- A mix of NDMP and non-NDMP backups can be present in the same MPX backup group.
- File and directory DAR are allowed.
- NDMP multiplexing works with both VTL and PTL. However, VTL users typically do not use NDMP multiplexing because they can add more virtual tape devices to accommodate additional streams.
- For NDMP multiplexed backups the storage unit and policy schedule multiplex value must be set to a value greater than one.

About NDMP support for Replication Director

NDMP can be used to back up, browse, and restore snapshots. The advantage to using Replication Director and creating a backup policy that uses NDMP is that NetBackup needs to mount only the primary data to perform these actions.

For additional information about NDMP with Replication Director, see the [NetBackup Replication Director Solutions Guide](#).

Limitations of Replication Director with NDMP

Consider the following limitations before configuring NDMP to be used with Replication Director:

- The Solaris_x86 operating system is not supported.
- The **Multiple copies** NetBackup policy option is not supported for image copies in the NDMP data format.
- The **Restore the file using a temporary filename** restore option is not supported on Windows clients.

- Restores to a local file system are not supported with an **MS-Windows** or a **Standard** policy that has the NDMP **Data Mover** enabled.
- Do not include both the qtree and the volume on which the qtree resides in the same **Backup Selection** list.
- Only one NDMP backup of a snapshot per `backupid` is allowed.
- The **Index From Snapshot** operation is supported only in a Replication Director configuration, however, a Standard or MS-Windows policy with NDMP Data Mover enabled is also not supported.

Note: The **Index From Snapshot** operation is not supported for NetApp ONTAP 7-mode.

- When you make changes to the NDMP policy after the last full or incremental schedule (for example, if you add or delete a backup selection), the content for the next incremental retrieves the entire content of the snapshot rather than retrieving only the content that has changed. The next incremental schedule however, after only retrieves content that has changed as expected.
- If IPv6 is enabled on a primary server running Linux, the NDMP Index-from-snapshot job may fail with an error: (2113) Invalid or no disk array credentials are added for vserver. You may see the error even after providing the correct credentials for the vserver disk array host. To resolve this issue, do any one of the following:
 - Disable IPv6 on the primary server.
 - Add a mapping of the disk array host's IP and the FQDN short name in the `/etc/hosts` file of the primary server.

About NDMP support for NetApp clustered Data ONTAP (cDOT)

The following table describes the terminology that is used in this topic.

Table 14-3 NetApp cDOT terminology

Term	Definition
CAB	Specifies the Cluster Aware Backup (CAB) NDMP API extension. The CAB enables support of a NetApp cDOT system for optimal, node-transparent backups.

Table 14-3 NetApp cDOT terminology (*continued*)

Term	Definition
cDOT	Specifies the clustered Data ONTAP (cDOT); the NetApp clustered filer storage solution.
Cluster-management LIF	Specifies a single management interface for the entire cluster. This is the only logical interface (LIF) that NetBackup supports for device configuration.
Data LIF	Specifies the data logical interface (LIF) that is associated with the Vserver.
Intercluster LIF	Specifies a logical interface (LIF) that is used for intercluster communication.
LIF	Specifies a logical interface (LIF); an IP address and port that is hosted on a node of a NetApp cDOT system.
Node-management LIF	Specifies a dedicated IP address that is used to manage a node.
SVM	Specifies the Storage Virtual Machine (SVM); a NetApp clustered Data ONTAP construct that is a virtualization layer that includes volumes and LIFs. This allows for non-disruptive user and NDMP operations when the physical cluster resources change. Multi-tenancy is achieved by multiple SVMs (see the data LIF). The cluster itself is also an SVM (see cluster-management LIF).
Vserver	Specifies the virtual storage server; contains data volumes and one or more LIFs through which it serves data to the clients.

It is recommended to run a NetApp cDOT cluster in SVM-scoped NDMP mode (also called Vserver aware mode).

NetBackup supports optimal backup, restore, and duplication of NetApp cDOT FlexVol volumes using the CAB extension. The NetApp cDOT server (that runs in Vserver aware mode) provides unique location information (affinity) about volumes and tape drives. Using this affinity information, NetBackup performs a local backup instead of a three-way or remote backup if a volume and a tape drive share the same affinity. If multiple volumes that are hosted on different nodes are backed up or restored using the same job, NetBackup may switch drive paths if necessary (and possible) to perform the local backup.

Note: The NetApp Infinite volumes can be backed up and restored by using the standard policy types.

Note: There should be at least one intercluster LIF for each node of the cluster that does not host a cluster-management LIF. This is required for three-way and remote backups. If you do not specify an intercluster LIF, all of the three-way and remote backups for volumes that are not hosted on the same node as the cluster-management LIF fail. NetBackup does not access these LIFs directly, so it does not need credentials for them.

Installation Notes for NetBackup for NDMP

This chapter includes the following topics:

- [NetBackup for NDMP installation prerequisites](#)
- [Adding the NetBackup for NDMP license](#)
- [About existing NetApp cDOT configurations before you upgrade](#)

NetBackup for NDMP installation prerequisites

Note the following items before installing NetBackup and adding the NetBackup for NDMP license:

- NetBackup for NDMP functionality installs when the NetBackup server software is installed. No separate installation procedure is required. However, you must enter a valid license to use NetBackup for NDMP.

Note: If your NetBackup for NDMP server is not your primary server, install your NDMP license on the primary server.

In a clustered environment, perform the steps to add the license on each node in the cluster. First, freeze the active node so that migrations do not occur during installation. Unfreeze the active node after the installation completes. For information about freezing or unfreezing a service group, see the clustering section in the [NetBackup High Availability Administrator's Guide](#) for the cluster software you are running.

For more information about administering licenses, see the [NetBackup Administrator's Guide, Volume I](#).

Note: NetBackup for NDMP cannot be uninstalled separately from the full NetBackup product.

If you uninstall the full NetBackup product, make sure that no NetBackup for NDMP backups are active or running for the client. On the primary server, check the Activity Monitor in the **NetBackup web UI**. If the **Job State** for the backups indicates *Done*, you can then perform the uninstall procedure that is described in the [NetBackup Installation Guide](#).

- For lists of supported operating systems, hardware platforms, and NAS vendor features and software releases, see the [NetBackup Compatibility List for all Versions](#).

For a list of NAS platforms that NetBackup for NDMP supports, see the [NetBackup for NDMP: NAS Appliance Information](#) document.

- The drives and robots that are attached to the NDMP host must be the types that the NDMP host and NetBackup support. A list of supported robot types is available.

See “[About robotics control](#)” on page 131.

For more information about storage devices, see the [NetBackup Administrator's Guide, Volume I](#).

Adding the NetBackup for NDMP license

NetBackup for NDMP installs when the NetBackup server software is installed. No separate installation procedure is required. However, you must enter a valid license to use NDMP. Use the following procedure on the host that you want to be the NetBackup for NDMP server.

Note: If you install in a clustered environment, first freeze the active node so that migrations do not occur during installation. For information about freezing a service group, see the clustering section in the [NetBackup High Availability Administrator's Guide](#) for the cluster software you are running.

To add the NetBackup for NDMP license

- 1 Install NetBackup server and client software as explained in the [NetBackup Installation Guide](#).
- 2 NetBackup for NDMP is part of the core NetBackup product. To make sure a valid license for NetBackup for NDMP is registered, do the following:
 - Open the NetBackup web UI.

- Click **Settings > License management**.
 - Click **Add license**.
- 3 If this NetBackup for NDMP server is not your primary server, install your NDMP license on the primary server.
 - 4 In a clustered environment, perform these steps on each node in the cluster.
 - 5 If you install in a clustered environment, unfreeze the active node after the installation completes.

For information about unfreezing a service group, see the clustering section in the [NetBackup High Availability Administrator's Guide](#) for the cluster software you are running.

About existing NetApp cDOT configurations before you upgrade

This topic describes how to upgrade NetBackup with a NetApp cDOT system. If you use a NetApp cDOT system, review the following information before you upgrade to NetBackup 7.7 or later.

If your NetApp cluster is set to **node-scope-mode** and you have not yet installed NetBackup, your environment should be set up as follows before the upgrade:

- The client name that is used in the backup policy is the node-management LIF.
- Only the volumes that are hosted by the node that hosts the LIF are available for backup or restore. Each node must have a node-management LIF in the client list of the policy.
- Tape devices that are attached to a node are available for backup or restore.
 - The NDMP host name that is used for the device configuration is the node name (node-management LIF).
 - The tape devices are available only to the nodes to which they are connected.

After you upgrade to NetBackup, everything works as it did before the upgrade until you enable the NetBackup cDOT capabilities by disabling node-scope mode.

To start using the NetBackup cDOT capabilities, do the following:

1. Back up the catalog.
2. (Optional) Create a detailed image catalog report that provides the following:
 - Collects information, such as NDMP host names, policies, and backup selections, that can be used when you create the new cDOT backup policies.

- Determines the client names to search for when you restore the pre-cDOT backups in the new cDOT environment.
- 3. Upgrade all of the NetBackup media servers that are authorized to access the cluster. Upgrades do not have to occur at the same time, but must be done before the following step.
- 4. Enable the Vserver aware mode on the cluster by disabling node-scope-mode. Please see your specific cluster documentation.
- 5. If there are tape devices attached to the cluster, you must reconfigure your tape devices to use the cluster-management LIF as the NDMP host for the device configuration. See [“About Media and Device Management configuration”](#) on page 152.

Caution: NetBackup only supports the use of the cluster-management LIF for device configurations.

Note: For each node in the cluster that will have tape devices, be sure to configure all of the tape devices available to the cluster on that node. Any node that has access to a tape device should also have access to all of the tape devices.

- 6. Enable the NDMP service on the cluster for each data LIF that will be used for backups. See the NetApp documentation for more information.
- 7. Authorize the data LIF as needed for NetBackup access. See [“Authorizing NetBackup access to a NAS \(NDMP\) host”](#) on page 150.
- 8. Add, delete, or update the old storage units that are using the node names of the cluster.
- 9. Add, delete, or update the old policies that back up the cluster.
 - You must use either the data LIF or the cluster-management LIF as the client name. NetBackup does not support use of the node name for the client name.
 - Backup selections may also need to be adjusted.

Note: The use of the data LIF as a client will protect and catalog all volumes associated with the data LIF's Vserver under this client. The use of the cluster-management LIF as a client will protect and catalog all volumes on the entire cluster under this client.

10. To read the old images, you may have to use alternate client restore. For more information about alternate client restores, see the [NetBackup Administrator's Guide, Volume I](#)

If your NetApp cluster is set to **Vserver aware mode**, and you have not yet installed NetBackup, your environment should be set up as follows before the upgrade:

- The cluster is in Vserver aware mode. The Cluster Aware Backup (CAB) extension is enabled on the filer. NetBackup does not use the CAB extension.
- The client name used in the backup policy is the data LIF associated with a Vserver or the cluster-management LIF.
- Only volumes (that belong to the Vserver) hosted by a node that hosts the data LIF are available for backup or restore.
- The tape devices that are attached to the cluster are not available for backup or restore.

After you upgrade to NetBackup, the behavior is different and you need to make some changes. NetBackup now uses the CAB extension and enables it by default. Because of this, the following occurs:

- NetBackup uses all of the volumes that belong to the Vserver.
- NetBackup uses the volume affinities.

As a result of this change, the following occurs:

- When the `ALL_FILESYSTEMS` directive is in use by multiple policies for the same Vserver, NetBackup may back up the same volume multiple times under different policies. And further incremental backups may not be reliable.
- Multi-streamed backup jobs will start failing with status code 99. The following message is displayed in the job details for the failed jobs:

```
12/10/2014 14:42:11 - Error ndmpagent (pid=29502) NDMP backup failed,
path = /vs02/voll:PARAMETER:AFFINITY=4ac6c4b6-7e99-11e4-b3b6-1779f43af917
```

This happens because some components of NetBackup are not told to use the cluster in the Vserver aware mode. It is highly recommended to upgrade and enable the cDOT capabilities as soon as possible.

To start using the cDOT capabilities, you must do the following:

1. Back up the catalog.
2. Create a detailed image catalog report (it can be referenced later for read operations).

3. Upgrade all of the NetBackup media servers that are authorized to access the cluster. All media servers should be upgraded at the same time to avoid inconsistent behavior.
4. Run the `tpautoconf -verify ndmp_host` command for each pre-existing LIF that is configured in NetBackup. This command must be run from the media servers that have credentials to the LIF. After the command is successfully run, the `nbemmcmd` command should display output similar to the following example:

```
servername1@/>nbemmcmd -listsettings -machinename machinename123 -machinetype ndmp
NBEMMCMD, Version: 7.7
The following configuration settings were found:
NAS_OS_VERSION="NetApp Release 8.2P3 Cluster-Mode"
NAS_CDOT_BACKUP="1"
Command completed successfully.
```

NAS_OS_VERSION displays the NetApp Version.
 NAS_CDOT_BACKUP tells us if NetBackup uses the new cDOT capabilities.

Note: The `tpautoconf -verify ndmp_host` command is not required when a new Vserver is added.

5. You can now add devices to the NDMP cluster and access them using the cluster-management LIF. If you add devices, you must discover the devices.
6. Add storage units for the newly discovered devices.
7. Add, delete, or update the policies that reference the cluster as needed. Start using the cluster in Vserver aware mode.

If you do not want to enable the cDOT functionality immediately; for example, you want to upgrade the media servers in phases, you can disable the cDOT capabilities by doing the following:

1. Create the following touch file on all of the media servers that are authorized to access the NDMP host. This causes NetBackup to disable the CAB extension for all of the NDMP hosts for that media server.
 - On Windows: `install_path\NetBackup\db\config\DISABLE_NDMP_CDOT`
 - On UNIX: `/usr/opensv/netbackup/db/config/DISABLE_NDMP_CDOT`
2. You can disable the CAB extensions for specific NDMP hosts by creating the following file on the media servers with one or more NDMP host names (one per line):

- On Windows:

`install_path\NetBackup\db\config\DISABLE_NDMP_CDOT_HOST_LIST`

- On UNIX:

`/usr/opensv/netbackup/db/config/DISABLE_NDMP_CDOT_HOST_LIST`

An example of the content of the file is as follows. NetBackup disables the CAB extension only for Filer_1 and Filer_2.

Filer_1

Filer_2

To enable the cDOT functionality, these files must be deleted and you must follow all of the steps explained in the previous upgrade procedure.

Configuring NDMP backup to NDMP-attached devices

This chapter includes the following topics:

- [About configuring NDMP-attached devices](#)
- [Authorizing NetBackup access to a NAS \(NDMP\) host](#)
- [About access for three-way backups and remote NDMP](#)
- [About Media and Device Management configuration](#)
- [Configure an NDMP filer](#)
- [About creating an NDMP policy](#)
- [About enabling or disabling DAR](#)
- [Configuring NetBackup for NDMP in a clustered environment](#)

About configuring NDMP-attached devices

This topic explains how to configure backups on the storage devices that are attached to NDMP hosts. Only the NDMP-specific steps are described.

You can also use the NetBackup web UI to discover and configure the robots and drives that are attached to an NDMP host. NetBackup requires NDMP protocol versions V3 or V4.

To configure and use the NAS_Snapshot method, see the [NetBackup NAS Administrator's Guide](#).

See [“Authorizing NetBackup access to a NAS \(NDMP\) host”](#) on page 150.

See [“About Media and Device Management configuration”](#) on page 152.

See [“About adding volumes”](#) on page 154.

See [“About verifying NDMP password and robot connection”](#) on page 154.

See [“Adding NDMP storage units”](#) on page 155.

See [“About creating an NDMP policy”](#) on page 156.

See [“About enabling or disabling DAR”](#) on page 167.

See [“Configuring NetBackup for NDMP in a clustered environment”](#) on page 169.

Authorizing NetBackup access to a NAS (NDMP) host

Before NetBackup can perform backups using NDMP, it must have access to the NAS (or NDMP) host.

Note: Perform the following procedure on the primary server (not media server) if you plan to create snapshots using Replication Director.

To authorize NetBackup access to the NDMP host

- 1 Open the NetBackup web UI.
- 2 On the left, click **Credential management**. Then click the **Client credentials** tab.
- 3 Click **Add**. Select **NDMP host** and click **Next**.
- 4 Enter the name of the NDMP server for NetBackup to back up.

If you use NetApp's Clustered Data ONTAP, the NDMP host must be a Storage Virtual Machine (SVM).

The NDMP host name is case-sensitive. The name must match the name that is entered here whenever this host name is used.

Note: If you do not plan to use Replication Director and you add NDMP host credentials using the fully qualified domain name (FQDN), you must also indicate the fully qualified domain name on the client for lookups. The server list in the **Backup, Archive, and Restore** client interface must list the NDMP host by the FQDN as well.

If you add NDMP host credentials using a short name, you can use either the short name or the FQDN in the client's server list.

5 Specify the following:

(The term *credentials* refers to the username and password that NetBackup uses to access the NDMP host.)

Use the following credentials for this NDMP host on all media servers

Enables all NetBackup media servers that are connected to the NDMP host to access the NDMP host using the logon you specify:

- **Username:** The username under which NetBackup accesses the NDMP server. This user must have permission to run NDMP commands.
You can find out whether your NDMP host vendor requires a particular username or access level.
- **Password:** Enter the password for this user.

Use different credentials for this NDMP host on each media server

Specifies NDMP logons for specific NetBackup servers.

- Click **Add**.
- Select a NetBackup server and specify the username and password it uses to access the NDMP host.
- Click **Add**. NetBackup validates the username and password.
- If necessary, click **Add** again to specify other servers and credentials.

6 Click **Add**.

7 Repeat this procedure for each NDMP host that NetBackup backs up.

About access for three-way backups and remote NDMP

To perform three-way backups, you must authorize access to the NDMP host as described in the previous section.

Note the following points:

- Three-way backups; for the **NDMP host name**, specify the NDMP host that has no attached tape drive.
- NDMP to Media Manager storage units (remote NDMP); for the **NDMP host name**, specify the NDMP host to back up to the Media Manager storage unit that is defined on the NetBackup server.

See [“About remote NDMP”](#) on page 171.

See [“About configuring NDMP-attached devices”](#) on page 149.

About Media and Device Management configuration

On the **NetBackup web UI**, use **Storage > Devices** to add drives and robots.

Note: It is recommended to connect any tape drive that is attached to a NetApp cDOT system to all of the cluster nodes. If you do not follow this recommendation, NetBackup may not be able to find the optimal path for data transfer.

The following procedures and examples treat NDMP configuration issues only.

- See [“Configure an NDMP filer”](#) on page 153.
- See [“Adding a robot or tape drive directly attached to an NDMP host”](#) on page 152.
- See [“Checking the device configuration”](#) on page 153.

See the [NetBackup Administrator's Guide for UNIX, Windows, and Linux, Volume I](#), for general information on configuring NetBackup media.

More information on configuring storage devices for specific NDMP hosts is available.

- See [“About NAS appliances support”](#) on page 198. for information about supported NDMP operating systems and NAS vendors.
- For a list of the features and software releases for each NAS vendor, for SSO support, and for the NetBackup versions that support these vendors, see the [NetBackup Compatibility List for all Versions](#).

These procedures do not apply to setting up the devices that are attached to the NetBackup media server. To back up NDMP data to media servers, you must configure storage units in the same way as ordinary NetBackup (non-NDMP) devices. More information is available:

See [“About remote NDMP”](#) on page 171.

See [“About adding volumes”](#) on page 154.

See [“About configuring NDMP-attached devices”](#) on page 149.

Adding a robot or tape drive directly attached to an NDMP host

This procedure describes how to configure a robot that is attached to an NDMP host.

To add a robot or tape drive directly attached to an NDMP host

- 1 In the **NetBackup web UI**, select **Storage > Tape storage**.
- 2 On the **Actions** menu, click **Configure storage devices** from the pop-up menu.
- 3 Select the primary server for the list. Click **Edit**.
- 4 In the **Edit device host** dialog, select **NDMP server**, click **Change**. Click **Next**.
- 5 Select the host that has the tape or robot configured. Click **Next**.

After the scan is complete, you can see the list of tapes and robots attached to the device. Select the required tapes and robots and click **Next**. Click **Commit changes**.

See [“About configuring NDMP-attached devices”](#) on page 149.

Checking the device configuration

On the NetBackup for NDMP server, use the following procedure to check the device configuration.

To check the device configuration

- ◆ On UNIX:
 - Execute `/usr/opensv/volmgr/bin/vmps`.
 - Verify that `ltid`, `vmd`, `avrd`, and any required robotic daemons are active.
- On Windows:
 - From the **NetBackup web UI**, select **Activity Monitor**.
 - In the right pane, select the **Processes** tab.
 - Verify that `ltid`, `vmd`, `avrd`, and any required robotic daemons processes are active.

See [“About configuring NDMP-attached devices”](#) on page 149.

Configure an NDMP filer

This procedure shows how configure NetBackup to an NDMP filer.

To configure an NDMP filer:

- 1 On the left, select **Storage > Tape storage**.
- 2 On the **Actions** menu, click **Configure storage devices** from the pop-up menu.

- 3 Select the primary server for the list. Click **Edit**.
- 4 In the **Edit device host** dialog, select **NDMP server**, click **Change**. Click **Next**.
On the next page, click **Add**.
- 5 In the Add NDMP host dialog, enter the NDMP host name. Select **Use the following credentials for this NDMP host on all media servers**. Enter the User name and password for the desired NDMP filer.
- 6 Click **Add** to add the filer.

See [“About NAS appliances support”](#) on page 198. for information about supported NDMP operating systems and NAS vendors.

About verifying NDMP password and robot connection

When you authorize NetBackup access to the NDMP host and configure robots using the **NetBackup web UI**, NetBackup automatically verifies the NDMP credentials and the robotic configuration. If you want, you can re-verify them. For example:

```
tpautoconf -verify ndmp_host_name
```

A successful verification looks like the following:

```
Connecting to host "stripes" as user "root"...
Waiting for connect notification message...
Opening session--attempting with NDMP protocol version n...
Opening session--successful with NDMP protocol version n
  host supports MD5 authentication
Getting MD5 challenge from host...
Logging in using MD5 method...
Host info is:
  host name "stripes"
  os type "NetApp"
  os version "NetApp Release n.n.n.n"
  host id "0033625811"
Login was successful
Host supports LOCAL backup/restore
Host supports 3-way backup/restore
```

About adding volumes

Use the NetBackup **Media and Device Management** utility to add the volumes that you plan to use for the NDMP host backups.

See the [NetBackup Administrator's Guide, Volume I](#), for instructions.

When you specify the **Robot control host** for a volume that is in a robot, specify the host name for the NetBackup for NDMP server. Do not specify the NDMP host.

See [“About configuring NDMP-attached devices”](#) on page 149.

Adding NDMP storage units

On the NetBackup primary server, add an NDMP-type storage unit for the devices that contain the backup data. Most of the requirements are the same as for adding a Media Manager storage unit. The following procedure explains how to add an NDMP storage unit.

See the [NetBackup Administrator's Guide, Volume I](#), for more information on storage units.

The NDMP-type storage units are not used for backups to devices that are attached to NetBackup media servers. Use a non-NDMP storage unit instead.

See [“About remote NDMP”](#) on page 171.

To add NDMP storage units

- 1 On the left, click **Storage**, then **Storage units**.
- 2 Click Add. In the **Add storage unit** dialog, select **Drives and robots** from the **Storage type** drop-down. Click **Start**.
- 3 In the **Add Tape storage unit** page, enter the following:

Storage unit name	Enter a unique name for the storage unit.
Maximum concurrent write drives	Select the maximum number of drives for concurrent writing.
Reduce fragment size to	Enter the minimum fragment size for this storage unit.
Enable multiplexing	Enter 1 as multiplexing is not allowed with NDMP storage units.
On demand only	Specify whether the storage unit is available only when a policy or schedule specifically requests it. If this option is not used, the storage unit is available to any NDMP policy or schedule.

Click **Next**.

- 4 In the **Storage devices** page, select the storage device to attach to the storage unit. Click **Next**.
 - 5 In the **Media server** page, select the media server associated with this storage unit. Click **Next**.
 - 6 In the **Review** page, check the entire configuration, and click **Save**.
- See [“About configuring NDMP-attached devices”](#) on page 149.

About creating an NDMP policy

On the NetBackup primary server, you must create an NDMP policy to configure backups of the NDMP host.

Creating an NDMP policy is very similar to creating other NetBackup policy types. The following topics explain the differences when creating NDMP policies.

- See [“Attributes tab options for an NDMP policy”](#) on page 157.
- See [“Schedules tab options for an NDMP policy with Accelerator for NDMP enabled”](#) on page 158.
- See [“Clients tab options for an NDMP policy”](#) on page 159.
- See [“Backup selection options for an NDMP policy”](#) on page 159.
- See [“About appropriate host selection for NetApp cDOT backup policies”](#) on page 156.

See the [NetBackup Administrator’s Guide, Volume I](#), for more information on NetBackup policies.

To configure an NDMP policy for the NDMP Snapshot and Replication method, see the [NetBackup Replication Director Solutions Guide](#).

To configure a policy for the NAS_Snapshot method, see the [NetBackup NAS Administrator’s Guide](#).

About appropriate host selection for NetApp cDOT backup policies

When configuring a backup policy to protect NetApp cDOT systems, use either the cluster-management LIF or the data LIF. Consider the following when using the cluster-management LIF as the backup policy client.

Advantages:

- Everything is cataloged under the cluster-management LIF.
- You only have to validate the cluster-management LIF.
- It is easier to back up everything using a few policies.

Disadvantages:

- If the cluster is in use by multiple departments in the same organization, it may be difficult to isolate the data between divisions. This may also be a security concern for some organizations if they want to share data between divisions.
- There is a limited granularity in the choice of volume pools and destination storage.
- Finding the appropriate data may be more difficult at the time of restore.

Consider the following when using the data LIF as the backup policy client.

Advantages:

- Everything is cataloged under the data LIF.
- If the cluster is in use by multiple departments in the same organization, it is very easy to isolate data between divisions.
- Data from different divisions can go to different volume pools and destination storage.
- Finding the appropriate data is easier at the time of restore.

Disadvantages:

- You need to add credentials for each data LIF.
- You need multiple policies to backup up the entire cluster.

Attributes tab options for an NDMP policy

The following policy attributes are applicable when you create an NDMP policy:

Policy Type: NDMP Do not select any other policy type.

Policy Storage Unit

- To direct backups for this policy to a specific storage unit if the NDMP host has multiple storage units, specify that storage unit name.
- For policies that use Accelerator for NDMP, the storage unit groups are supported only if the storage unit selection in the group is Failover. See the **Use Accelerator** attribute.
- For a three-way backup, specify a storage unit that was defined for the target NDMP host with attached tape.
- For NDMP backup to Media Manager storage units, specify a Media Manager storage unit that is defined for a device that is connected to a NetBackup media server.
See [“About remote NDMP”](#) on page 171.

Use Accelerator	Select Use Accelerator to enable Accelerator for NDMP. See the Policy Storage Unit attribute. See “About NetBackup Accelerator for NDMP” on page 181. for more information.
Replication Director	Select the Replication Director to configure an NDMP policy for Replication Director.
Allow multiple data streams	Set the value to a number greater than 1.

Schedules tab options for an NDMP policy with Accelerator for NDMP enabled

In the schedules list under the **Attributes** tab, the following parameter is optional for an NDMP policy with Accelerator for NDMP enabled.

Accelerator forced rescan	Select this option to enable an Accelerator forced rescan. This option is available only for the NDMP policies that use Accelerator for NDMP. An Accelerator forced rescan provides a safety net by establishing a new baseline for the next Accelerator backup. When you include this option, all the data on the filer is backed up. This backup is similar to the first full Accelerator backup: it provides a new baseline for the backups that follow. If you set up a weekly full backup schedule with the Use Accelerator option, you can supplement the policy with another schedule that enables Accelerator forced rescan. You can set the schedule to run every 6 months or whenever it is appropriate for your environment. Expect backups with Accelerator forced rescan to run slightly longer than accelerated full backups. More information about Accelerator for NDMP is available: See “About NetBackup Accelerator for NDMP” on page 181.
----------------------------------	--

About backup types in a schedule for an NDMP policy

You can specify any of the following backup types in a schedule for an NDMP policy:

- Full
- Cumulative Incremental
- Differential Incremental

Specify **Override policy storage unit** only if the client of NetBackup (the NDMP host) has more than one storage unit and you want to use a specific storage unit for this schedule. In this case, the client must be the only client in this NDMP policy.

See [“About configuring NDMP-attached devices”](#) on page 149.

Clients tab options for an NDMP policy

In the client list, the following options are required for each client in an NDMP policy:

Hostname	<i>Name of the NDMP host.</i> If you use a NetApp cDOT system, the NDMP host name can only be a Vserver (a data LIF or a cluster-management LIF). NetBackup does not support any other LIF type as the NDMP host name.
Hardware and operating system	NDMP NDMP. If you use a NetApp cDOT system, NetBackup changes the operating system name from NDMP to cDOT.

Backup selection options for an NDMP policy

The backup selections list must specify directories from the perspective of the NDMP host.

For example:

```
/vol/home/dir1/  
/vol/voll
```

If you have a Windows primary server or media server, you cannot specify a directory that contains unsupported characters in its name. For example, Windows does not support the following characters in file and folder names and therefore they cannot be used in backup selection specifications:

- ~ (tilde)
- # (number sign)
- % (percent)
- & (ampersand)
- * (asterisk)
- [] (braces)
- / (backslash)
- : (colon)
- < > (angle brackets)

- ? (question mark)
- \ (slash)
- | (pipe)
- " (quotation mark)

Refer to your Windows documentation for a complete list of unsupported characters.

You can also use wildcard characters in regular expressions or the directive `ALL_FILESYSTEMS` to specify path names in NDMP policy backup selections.

See [“Wildcard characters in backup selections for an NDMP policy”](#) on page 160.

See [“ALL_FILESYSTEMS and VOLUME_EXCLUDE_LIST directives”](#) on page 163.

See [“About environment variables in the backup selections list”](#) on page 165.

See [“About configuring NDMP-attached devices”](#) on page 149.

Wildcard characters in backup selections for an NDMP policy

You can use wildcard characters in regular expressions or the directive `ALL_FILESYSTEMS` to specify path names in NDMP policy backup selections.

Wildcard characters in regular expressions or directives are valid for streaming and non-streaming NDMP backups.

Note: Directory-level expansion is not supported for some NDMP servers. Some NDMP filer vendors do not have the APIs that NetBackup uses to support wildcard characters lower than the volume level.

If you specify a backup selection using wildcard characters lower than the volume level for these filers, status code 106 is generated. The following message is displayed: **Invalid file pathname found, cannot process request.**

Currently, only NetApp filers support wildcard characters for backup selections lower than the volume level. This support is not available in NetApp clustered Data ONTAP version 8.2.

To see the versions of NetApp Data ONTAP that support wildcard characters for backup selections lower than the volume level, refer to the [NetBackup Compatibility List for all Versions](#).

You cannot use any wildcard characters that also match file names. For example, a backup selection might include `/vol/vol_archive_01/autoit*`. This specification might match a path name such as `/vol/vol_archive_01/autoit_01/`. However, if this specification also matches a file name like `/vol/vol_archive_01/autoit-v1-setup.exe`, the backup job fails with status code 99 because wildcards can specify only path names. The following message is displayed: **NDMP backup failure (99).**

Table 16-1 Valid wildcard characters for NDMP policy backup selections

Wildcard character	Description
*	Specifies a string match. For example: <pre>/vol/vol_archive_*</pre> This form of the path specification matches all paths that begin with the literal characters <code>/vol/vol_archive_</code> and end with any characters. The string match wildcard can also specify multiple variable characters between literal characters as in the following examples: <pre>/vol/ora_*archive or /vol/ora_*archive*</pre> <pre>/vol/ora_vol/mtree_*archive or /vol/ora_vol/mtree_*archive*</pre>
?	Specifies a single-character match. <pre>/fs?</pre> This path specification matches all paths that begin with the literal characters <code>/fs</code> and end with any single character. For example, <code>/fs1</code>, <code>/fs3</code>, <code>/fsa</code>, <code>/fsd</code> and so on match the specified pattern <code>/fs?</code>.

Table 16-1 Valid wildcard characters for NDMP policy backup selections
(continued)

Wildcard character	Description
[...]	Specifies an alphanumeric pattern match. For example: <pre>/fs[1-9]</pre> This path specification matches all paths that begin with the literal characters <code>/fs</code> and end with any single numeric character from 1 through 9. For example, <code>/fs1</code>, <code>/fs2</code>, and so on up to <code>/fs9</code> match the specified pattern <code>/fs[1-9]</code>. However, <code>/fs0</code> and <code>/fsa</code> do not match the specified pattern; 0 is out of the specified numeric range, and <code>a</code> is a non-numeric character. The pattern match wildcard can also specify alphanumeric patterns such as <code>/fs[1-5a]</code>. This specification matches <code>/fs1</code>, <code>/fs2</code>, and so on up to <code>/fs5</code> as well as <code>/fsa</code>. Similarly, the pattern match wildcard can also specify patterns like <code>/fs[a-p4]</code>. This specification matches <code>/fsa</code>, <code>/fsb</code>, and so on up to <code>/fsp</code> as well as <code>/fs4</code>. You must use multiple backup selection specifications if the pattern can match more than 10 volume names in a numeric series. For example, you may want to back up 110 volumes that begin with the literal characters <code>/vol/ndmp</code> and are numbered 1 through 110. To include these volumes in a backup selection with wildcards, specify three backup selections with the following wildcard patterns: ■ <code>/vol/ndmp[0-9]</code> This pattern matches any volume name that begins with <code>/vol/ndmp</code> and ends with a single numeric character 0 through 9. ■ <code>/vol/ndmp[0-9][0-9]</code> This pattern matches any volume name that begins with <code>/vol/ndmp</code> and ends with the two-digit numeric characters 00 through 99. ■ <code>/vol/ndmp[0-9][0-9][0-9]</code> This pattern matches any volume name that begins with <code>/vol/ndmp</code> and ends with the three-digit numeric characters 000 through 999. Do not specify <code>/vol/ndmp[1-110]</code> in this example. This pattern produces inconsistent results.
{...}	Curly brackets can be used in the backup selection list and the <code>VOLUME_EXCLUDE_LIST</code> directive for NDMP policies. A pair of curly brackets (or braces) indicates multiple volume or directory name patterns. Separate the patterns by commas only; no spaces are permitted. A match is made for any or all entries. For example: <pre>{*volA,*volB} or {volA*,volB*}</pre>

Note the following restrictions and behaviors regarding wildcard expressions:

- It is not recommended that you use a single forward-slash character (/) in an NDMP policy backup selection. This method of including all the volumes on an NDMP filer in the selection is not supported. Instead, use the `ALL_FILESYSTEMS` directive:

See [“ALL_FILESYSTEMS and VOLUME_EXCLUDE_LIST directives”](#) on page 163.

- Nested wildcard expressions can result in recursive path name expansion operations that can impact performance, especially for directories that have a very large number of files or directories. An example of nested wildcard expansion is as follows:

```
/vol/fome06/*/*private
```

- Wildcard expressions do not span or include a path separator (/).
- All backup selections that contain a wildcard expression must start with a path separator (/). An example of a correct wildcard expression is as follows:

```
/vol/archive_*
```

An example of an incorrect wildcard expression is as follows:

```
vol/archive_*
```

ALL_FILESYSTEMS and VOLUME_EXCLUDE_LIST directives

The `ALL_FILESYSTEMS` directive provides a method to include all file systems and volumes on an NDMP filer in an NDMP backup policy.

You can exclude specific volumes from an `ALL_FILESYSTEMS` backup selection if you do not want to back up every volume on an NDMP filer. Use the `VOLUME_EXCLUDE_LIST` directive for this purpose. You may use valid wildcard characters in the `VOLUME_EXCLUDE_LIST` statement.

Note: The following examples use selections that are specific to NetApp Data ONTAP 7-mode. For specific examples of backup selections for other configurations, refer to the appropriate documentation.

The `VOLUME_EXCLUDE_LIST` statements must precede `ALL_FILESYSTEMS` statement. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Hr_allfiles_vol01  
ALL_FILESYSTEMS
```

or

```
VOLUME_EXCLUDE_LIST=/vol/testvol*  
ALL_FILESYSTEMS
```

To specify multiple values in a `VOLUME_EXCLUDE_LIST` statement, separate the values with a comma. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Hr_allfiles_vol01,/vol/testvol*  
ALL_FILESYSTEMS
```

You can also specify more than one `VOLUME_EXCLUDE_LIST` statement with an `ALL_FILESYSTEMS` directive. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Hr_allfiles_vol01  
VOLUME_EXCLUDE_LIST=/vol/testvol*  
ALL_FILESYSTEMS
```

A `VOLUME_EXCLUDE_LIST` statement may include a maximum of 256 characters. Create multiple `VOLUME_EXCLUDE_LIST` statements if necessary to avoid exceeding the limit of 256 characters. If you specify more than 256 characters, the volume list is truncated. A truncated statement may result in a backup job failure, and the error message `Invalid command parameter(20)` is displayed.

If the backup selection includes read-only volumes or full volumes, an NDMP backup job fails with the status code 20 (`Invalid command parameter(20)`). If you encounter a similar NDMP backup job error, review the `ostfi` logs to identify the volumes for which the failure occurred. You can use `VOLUME_EXCLUDE_LIST` statements with the `ALL_FILESYSTEMS` statement to exclude the read-only volumes and the volumes with insufficient space.

In a NetBackup Replication Director environment where snapshots are replicated to a secondary filer, it is recommended that you use storage lifecycle policies to control backups on the secondary filer.

On NetApp 7-mode storage systems, it is generally not recommended for users to store files in `/vol/vol0` because the volume contains filer system files. For this reason, `vol0` should be excluded from the backup if the `ALL_FILESYSTEMS` directive is used in the backup policy. The following is a backup selection list that excludes `/vol/vol0`:

```
VOLUME_EXCLUDE_LIST=/vol/vol0  
ALL_FILESYSTEMS
```

- Do not use `ALL_FILESYSTEMS` to backup all volumes on a secondary filer. Inconsistencies may occur when automatically created NetApp FlexClone volumes are backed up or restored. Such volumes are temporary and used as

virtual copies or pointers to actual volumes and as such do not need to be backed up.

- If you must back up all volumes on a secondary filer, it is recommended that you exclude the FlexClone volumes as well as replicated volumes. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Clone_*
VOLUME_EXCLUDE_LIST=/vol/*_[0-9]
VOLUME_EXCLUDE_LIST=/vol/*_[0-9][0-9]
VOLUME_EXCLUDE_LIST=/vol/*_[0-9][0-9][0-9]
ALL_FILESYSTEMS
```

This example assumes all FlexClone volumes and only FlexClone volumes begin with `/vol/Clone_`. Adjust the volume specifications appropriately for your environment.

- `VOLUME_EXCLUDE_LIST` applies only to `ALL_FILESYSTEMS`. It does not apply to explicit backup selections or wildcard-based backup selections. If you use the `ALL_FILESYSTEMS` directive in an NDMP policy for Clustered Data ONTAP, you must exclude each selected SVM's root volume using the `VOLUME_EXCLUDE_LIST` directive. Otherwise the backups fail.

Backups from snapshots for NDMP policies fail when the import of a snapshot fails for volumes where logical unit numbers (LUNs) reside with status code 4213 (Snapshot import failed). To avoid this error, use the `VOLUME_EXCLUDE_LIST` directive to exclude any volumes that are used to create LUNs accessed through a storage area network (SAN).

About environment variables in the backup selections list

NDMP lets you use environment variables to pass configuration parameters to an NDMP host with each backup. NDMP environment variables can be one of the following types:

- Defined as optional by the NDMP protocol specification.
You can set these variables.
- Specific to an NDMP host vendor.
You can set these variables.
See [“About NAS appliances support”](#) on page 198. for up-to-date information on environment variables relating to particular NAS vendors. The topic also contains configuration and troubleshooting help for particular NAS systems.

For Isilon filers only, note the following behaviors with environmental variables:

- With Isilon filers, if you set the `HIST` environment variable in a NetBackup NDMP backup policy with Accelerator enabled, you may specify only the

value `D` (that is, `SET HIST=D`). `D` specifies a directory/node file history format. If you specify any other value for the `HIST` variable, NetBackup generates a message that asks you to change the value to `D`. If you do not use a `HIST` variable in the policy, the backup should complete successfully.

- If you change any of the variables in a NetBackup NDMP backup policy with Accelerator enabled, the Accelerator optimization will be 0% until you run a second full backup with the same variables. When the policy's variables change, a new baseline image is created with the first full backup. You will see Accelerator optimization only after the second full backup with the same variables.

- Reserved for use by NetBackup:

```
FILESYSTEM
DIRECT
EXTRACT
ACL_START
```

In NetBackup, environment variables can be set within the backup selections list by specifying one or more `SET` directives.

Note: In the backup selections list, the `SET` directive must be the first in the list, followed by the file systems or volumes to back up.

In general, the syntax of a `SET` directive is as follows:

```
SET variable = value
```

Where *variable* is the name of the environment variable and *value* is the value that is assigned to it. The value can be enclosed in single or double quotes, and must be enclosed in quotes if it contains a space character. For example:

```
SET ABC = 22
SET DEF = "hello there"
```

Setting a variable equal to no value removes any value that was set previously for that variable. For example:

```
SET ABC =
SET DEF =
```

Variables accumulate as the backup selections list is processed. For example, a backup selection may contain the following entries:

```
/vol/vol1
SET HIST = N
```

```

/vol/vol2
SET DEF = 20
SET SAMPLE = all
/vol/vol3

```

In this example, directory `/vol/vol1` is backed up without any user-specified environment variables. The second directory (`/vol/vol2`) is backed up with the variable `HIST` set to `N`. The third directory (`/vol/vol3`) is backed up with all three of the environment variables set (`HIST = N`, `DEF = 20`, and `SAMPLE = all`).

Note: You cannot restore a single file if `HIST = N` is set. Only full volume restores are available when the `HIST` variable is set to `N`.

If an environment variable appears again later in the list, the value of this variable overrides the previous value of the variable.

The values that each backup uses are saved and provided to subsequent restores of the directory. The NDMP host may have some environment variables that are set internally and these are also saved for restores.

See [“About configuring NDMP-attached devices”](#) on page 149.

About enabling or disabling DAR

By default, NetBackup for NDMP is configured to use Direct Access Recovery (DAR) to restore files or directories. DAR is used somewhat differently for file restore than for directory restore.

The following table describes how DAR is used for file and directory restores.

Table 16-2 How DAR is used for file and directory restores

Type of restore	Description
File restore	For each restore of files (not of directories), NetBackup automatically determines if the use of DAR speeds up the restore. NetBackup uses DAR only when it results in a faster restore.

Table 16-2 How DAR is used for file and directory restores (*continued*)

Type of restore	Description
Directory restore	For restore of directories, by default DAR is always used to restore a subdirectory but never used to restore the directory containing an entire image. For example, if <code>/vol/vol0</code> contains the entire image, and <code>/vol/vol0/dir1</code> is a subdirectory, DAR is used by default to restore <code>/vol/vol0/dir1</code>. But it is not used to restore <code>/vol/vol0</code>. For restore of subdirectories, NetBackup does not attempt to gauge the effectiveness of using DAR. Unless DAR is manually disabled, NetBackup always uses DAR to restore subdirectories. See “Disabling DAR for file and directory restores” on page 168.

Note: You may have to disable DAR if you have problems with DAR and your NDMP host is an older computer or is not running the latest NAS OS version.

See [“About configuring NDMP-attached devices”](#) on page 149.

Disabling DAR for file and directory restores

This procedure disables DAR for both file and directory restores, for all NDMP policies.

To disable DAR for file and directory restores

- 1 In the **NetBackup web UI**, select **Hosts > Host properties**.
- 2 Select the server name and click **Edit media server**.
- 3 Select the General server.
- 4 Uncheck the **Use direct access recovery for NDMP restores** box.
This action disables DAR on all NDMP restores.
- 5 Click **Save**.

See [“About configuring NDMP-attached devices”](#) on page 149.

Disabling DAR for directory restores only

This procedure disables DAR for directory restores only. It leaves DAR enabled for individual file restores.

To disable DAR on restores of directories only, for all NDMP policies

- 1 Enter the string NDMP_DAR_DIRECTORY_DISABLED in the following file:

```
/usr/opensv/netbackup/db/config/ndmp.cfg
```

- 2 To turn on directory DAR, remove (or comment out) the NDMP_DAR_DIRECTORY_DISABLED string from the `ndmp.cfg` file.

See [“About configuring NDMP-attached devices”](#) on page 149.

Configuring NetBackup for NDMP in a clustered environment

The following must be installed on each node of the cluster before you can configure NetBackup for NDMP in a clustered environment:

- The NetBackup server
See the [NetBackup Installation Guide](#).
- NetBackup for NDMP software.
See [“NetBackup for NDMP installation prerequisites”](#) on page 142.
For Windows servers, only the NetBackup for NDMP license has to be installed.

To configure NetBackup for NDMP in a clustered environment

- 1 Configure NDMP-attached robots and drives. Then configure storage units and policies as in a normal, non-clustered environment:
 - You can use the NetBackup web UI, or configure the devices manually.
See [“Authorizing NetBackup access to a NAS \(NDMP\) host”](#) on page 150.
 - To use the same robotic libraries throughout a cluster, the robot numbers must be consistent. NetBackup attempts to ensure this configuration. If you configure robots manually, be sure to use the same robot number for a given robot, from one host to another in the cluster.
- 2 When you finish configuring devices and policies for NetBackup for NDMP, failover to the next node in the cluster and configure the drives and robots.

Select the same robot number that you used when configuring the robot for the first node.

After NetBackup is configured in a clustered environment, most configuration information is available to all nodes in the cluster. The information is available by means of a shared hard drive. However, in the **NetBackup web UI**, if you make changes to **Host > Host properties**, they are not available on the shared drive. Such changes apply only to the active node. You must manually duplicate on each

node the changes to **Host Properties** that are made on the active node. This action lets NetBackup perform exactly the same way in case of failover to another node.

Refer to the [NetBackup High Availability Guide](#) for further assistance.

See “[About configuring NDMP-attached devices](#)” on page 149.

Configuring NDMP backup to NetBackup media servers (remote NDMP)

This chapter includes the following topics:

- [About remote NDMP](#)
- [Configuring NDMP backup to Media Manager storage units](#)


About remote NDMP

This topic describes how to configure NetBackup for NDMP to make backups to Media Manager storage units (remote NDMP). Only NDMP-specific steps are described.

Using remote NDMP, you can back up NDMP data to a configured drive in a Media Manager storage unit on a NetBackup media server. The drive can be used for both NDMP backups and for non-NDMP backups.

An added feature to remote NDMP is NDMP multiplexing. NDMP multiplexing works with remote NDMP. It concurrently writes multiple backup streams to the same storage device from the same client or different clients.

Figure 17-1 NDMP backup to a Media Manager storage unit



Configuring NDMP backup to Media Manager storage units

This section describes how to configure NDMP backups to Media Manager storage units.

To configure NDMP backups to Media Manager storage units

- 1 Authorize the NetBackup server to access the NDMP hosts you want to back up.

Perform the following steps on the primary server (not media server) if you plan to create snapshots using the Snapshot Client NAS_Snapshot method:

- Create an NDMP filer. See [“Configure an NDMP filer”](#) on page 153.
- Fill in the values.
See [“Authorizing NetBackup access to a NAS \(NDMP\) host”](#) on page 150.
- Repeat these steps for each NDMP filer that the NetBackup server backs up.

- 2 Use NetBackup web UI to configure the drive(s) and robot(s).

Note the following:

- Do not use the "Configuring NDMP backup to NDMP-attached devices" topic in this guide. Configure the robots and drives as ordinary NetBackup devices, not as NDMP-attached devices.

- Drives can be shared using the NetBackup Shared Storage Option (SSO). The drives can be shared as both NDMP drives and non-NDMP drives. See [“About the Shared Storage Option \(SSO\) with NetBackup for NDMP”](#) on page 195.
- 3** Create a Media Manager storage unit for the drive(s). The storage unit type must be Media Manager, not NDMP.
- For NDMP multiplexing, do the following steps:
- Select the **Enable Multiplexing** check box on the **New Storage Unit** menu.
 - Set the **Maximum streams per drive** entry to a value greater than one.
- For details on storage units, refer to the [NetBackup Administrator's Guide, Volume I](#).
- 4** Create an NDMP-type policy. On the **New/Change Policy** display, be sure to specify the storage unit that was created in the previous step.
- Note the following for NDMP multiplexing:
- Set the **Media multiplexing** attribute on the **Add New Schedule** menu to a value greater than one.

Configuring NDMP DirectCopy

This chapter includes the following topics:

- [About NDMP DirectCopy](#)
- [Configuring NDMP DirectCopy](#)
- [Using NDMP DirectCopy to duplicate a backup image](#)

About NDMP DirectCopy

NetBackup supports virtual tape libraries (VTLs). A virtual tape library uses disk-based technology to emulate a tape library (robot) and drives. The backup image is written to one or more disks in the VTL. The VTL allows the image to be treated as though it resides on tape, but with the access speed of a disk.

For additional storage (such as for disaster recovery), NetBackup copies backup images from the VTL disk to a physical tape in an NDMP storage unit. It copies without using media server I/O or network bandwidth. NetBackup can also copy NDMP images directly between NDMP tape drives attached to an NDMP host.

In both cases, this function is called NDMP DirectCopy. This function also enables NetBackup to restore data directly from either the image in the VTL or from the physical NDMP tape. NDMP DirectCopy supports backup to tape and restore from tape for NDMP data as well as non-NDMP data. Tape-to-tape duplications of backup images are also supported.

NDMP DirectCopy does not support multiplexed backup, synthetic backup, or multiple copies. It also does not support storage unit groups for the destination device. If you select a storage unit group, NDMP DirectCopy is disabled. The data transfer takes place over the network by means of the NetBackup server.

To initiate the NDMP DirectCopy, you can use the NetBackup duplication feature in the **NetBackup web UI**, the `bpduplicate` command, or NetBackup Vault.

NDMP DirectCopy operates in the following environments:

- A NetBackup media server that is connected to a VTL that has access to a physical tape library. The steps for configuring NDMP DirectCopy are described in this topic.
- A NetBackup for the NDMP server that is connected to an NDMP host that has access to a tape library (no VTL). This NDMP backup environment is described in other topics of this guide. In this environment, no additional configuration is required for NDMP DirectCopy.

If your NDMP host and storage devices are correctly configured, NetBackup uses NDMP DirectCopy when you duplicate an NDMP backup that NetBackup had created.

Prerequisites for using NDMP DirectCopy

Note the following prerequisites for using NDMP DirectCopy:

- NetBackup for NDMP software must be installed. NetBackup for NDMP is enabled by the Enterprise Disk Option license. It requires the NDMP protocol version V4 or higher.
- The [NetBackup Compatibility List for all Versions](#) indicates which VTL software supports this functionality.
- If your environment includes a VTL, the VTL must be installed and set up according to the vendor's instructions. The NetBackup Enterprise Disk Option license(s) are required. The Enterprise Disk Option license enables NDMP DirectCopy functionality.
- The VTL must have the NDMP capabilities needed to support NDMP DirectCopy.
- To make direct copies from one NDMP tape drive to another (no VTL), the NetBackup for NDMP license is required.


NDMP DirectCopy with VTL

The NDMP DirectCopy feature uses a VTL that has an embedded NDMP tape server using the NDMP protocol. The embedded NDMP tape server moves the image from the VTL disk directly to a physical tape. The image does not pass through the NetBackup media server or travel over the network.


Note: In a VTL environment, a NAS appliance is not required. The VTL emulates a NAS (NDMP) host. The VTL requires NDMP tape server functionality.

The following figure represents a VTL from two perspectives. It shows the actual hardware present in a VTL configuration and the configuration from the perspective of NetBackup.

Figure 18-1 Overview of NDMP DirectCopy with VTL



The following figure shows the data flow and control for a VTL.


Figure 18-2 NDMP DirectCopy with VTL data flow and control

1. NetBackup media server sends the backup over a direct device path (SCSI or SAN) to the VTL.
2. NetBackup selects an NDMP device path to the VTL and creates an NDMP control session for the device.
3. NetBackup selects a tape volume from the physical tape library. It then selects an NDMP device path from the library and creates a second NDMP control session for the device.
4. By means of the NDMP protocol, the backup image in the VTL is copied directly to the physical tape library (not sent over the network).
5. The image can be restored directly to the media server from either the VTL or the physical tape.

NDMP DirectCopy without VTL

By means of the NetBackup duplication feature, NetBackup can copy NDMP images between tape drives attached to an NDMP host. A typical usage is to copy images between tape drives within the same tape library. (Images can also be copied between tape libraries.) Like NDMP DirectCopy with a VTL, the copied data does not pass through the NetBackup media server or travel over the network.

Figure 18-3 NDMP DirectCopy between tape drives accessible to an NDMP host



Configuring NDMP DirectCopy

Use the following procedure to configure NDMP DirectCopy from the backups that were made to a VTL.

To configure NDMP DirectCopy from the backups that were made to a VTL

- 1** Configure the VTL as an NDMP host. You can do it from the web UI, **Storage > Tape storage**.
- 2** Authorize NetBackup access to the VTL. Note that the VTL emulates an NDMP host.

See [“Authorizing NetBackup access to a NAS \(NDMP\) host”](#) on page 150.

- 3** Configure the VTL as a robot, then configure one or more tape drives in a Media Manager storage unit.

- 4 Configure one or more tape drives in the VTL as Network Attached Storage devices, and create one or more NDMP storage units for the drives.

See [“Adding NDMP storage units”](#) on page 155.

The drives can be the same as those that were selected in the previous step. NetBackup supports sharing of drives among media servers and NDMP hosts.

- 5 Configure one or more NDMP tape drives in the physical tape library, and add the drives to NDMP storage units. Use the same procedures as those mentioned in the previous step.

You can also use these drives in Media Manager storage units, if they are shared on a SAN.

Using NDMP DirectCopy to duplicate a backup image

NetBackup uses NDMP DirectCopy when you duplicate a backup image. To run a duplication, you can use any of the following methods:

- Initiate the duplication from the **NetBackup web UI**. In the **NetBackup web UI**, select **Catalog**.
Select **Duplicate** option.
See [“Initiating NDMP DirectCopy with the NetBackup web UI”](#) on page 180.
- NetBackup Vault
Refer to the [NetBackup Vault Administrator’s Guide](#) for more information.
- The `bpduplicate` command
Refer to the [NetBackup Commands Guide](#) for detailed information about this command.
- A storage lifecycle policy (SLP)
In the **NetBackup web UI**, select **Storage > Storage lifecycle policies**.
See [“About storage lifecycle policies”](#) on page 37.

If you use a NetApp cDOT system in SVM-scoped NDMP mode, NetBackup tries to match the affinity for the source and the destination tape drive path, if possible, so the duplication can be performed optimally.

Requirements to use NDMP DirectCopy for image duplication

When NetBackup uses NDMP DirectCopy to duplicate an image, note the following:

- For the destination for the duplication, you must designate an NDMP storage unit in a VTL or in a physical tape library.

- An NDMP tape drive must be available to mount the source image. The NDMP tape drive can be one that was defined in the VTL, or it can be a physical tape drive in a tape library.

Setup instructions are available.

See [“About NDMP DirectCopy”](#) on page 174.

If these two requirements are met, NDMP DirectCopy is enabled. NetBackup copies the image directly to the designated storage unit without using media server I/O or network bandwidth.

NetBackup policy type for image duplication

You can duplicate an image that any NetBackup policy created. The policy need not be an NDMP policy.

See [“About NDMP DirectCopy”](#) on page 174.

The backup can be made to a storage unit in the VTL or to a storage device that is attached to an NDMP host. You can then copy the backup directly to a tape drive using the NetBackup Duplicate feature, as follows.

Initiating NDMP DirectCopy with the NetBackup web UI

Use the following procedure to initiate NDMP DirectCopy.

To initiate NDMP DirectCopy

- 1 In the **NetBackup web UI**, select **Catalog**.
- 2 Set up the search criteria for the image that you want to duplicate. Click **Search**.
- 3 Select the want to duplicate and select **Duplicate** from the shortcut menu.

You must designate an NDMP storage unit as the destination for the duplication. Use the **Storage unit** field in the **Setup Duplication Variables** dialog box.

Accelerator for NDMP

This chapter includes the following topics:

- [About NetBackup Accelerator for NDMP](#)
- [About the track log for Accelerator for NDMP](#)
- [Accelerator messages in the NDMP backup job details log](#)
- [NetBackup logs for Accelerator for NDMP](#)

About NetBackup Accelerator for NDMP

Note: Currently only NetApp filers and Isilon filers are supported with the NetBackup Accelerator for NDMP option. (See the [NetBackup Compatibility List for all Versions](#) for the most recent list of supported versions of each NAS vendor.)

For NetApp filers, Accelerator for NDMP supports only the DUMP format. Consult your NetApp documentation for specific details about its DUMP format.

NetBackup's Accelerator option makes NDMP backups for NetApp and Isilon filers run faster than normal NDMP backups. NetBackup Accelerator increases the speed of full backups by using the filer's change detection techniques to identify the modifications that occurred since the last backup. After an initial full backup that protects all data from the filer, NetBackup Accelerator backs up only the changed data from the filer to the media server. The media server combines the changed data with any previous backup images to create a new full backup image; if a file or portion of a file is already in storage and has not been changed, the media server uses the copy in storage rather than reading it from the filer to complete the backup image. The end result is a faster NetBackup NDMP backup.

Note: For NetApp filers, you can expect to see Accelerator optimization in both full backups (regular and forced rescan) and incremental backups. For Isilon filers, you can expect to see Accelerator optimization only in full backups (regular – not forced rescan).

Accelerator for NDMP has the following advantages:

- Supports all NetBackup NDMP features, such as replication, DAR restores, and multiplexing.
- Creates a compact backup stream that uses less network bandwidth between the filer and NetBackup servers.
- Reduces the I/O and CPU overhead on the media server.

To configure Accelerator for NDMP, select the **Use Accelerator** check box that is found on the NDMP policy **Attributes** tab. No change to the filer is required.

Note: For Isilon filers only, note the following behaviors with environmental variables:

With Isilon filers, if you set the `HIST` environment variable in a NetBackup NDMP backup policy with Accelerator enabled, you may specify only the value `D` (that is, `SET HIST=D`). `D` specifies a directory/node file history format. If you specify any other value for the `HIST` variable, NetBackup generates a message that asks you to change the value to `D`. If you do not use a `HIST` variable in the policy, the backup should complete successfully.

If you change any of the variables in a NetBackup NDMP backup policy with Accelerator enabled, the Accelerator optimization will be 0% until you run a second full backup with the same variables. When the policy's variables change, a new baseline image is created with the first full backup. You will see Accelerator optimization only after the second full backup with the same variables.

More information about environmental variables in NDMP policies is available:

See [“About environment variables in the backup selections list”](#) on page 165.

Note: If you include the `smtape` environment variable for NetApp filers in an NDMP backup policy, no optimization is seen with Accelerator for NDMP enabled. The `smtape` environment variable always backs up an entire volume as if it is a full backup of a single file. Consult your NetApp filer documentation for specific details about `smtape`. See [“About NAS appliances support”](#) on page 198. for information about `smtape` in a NetBackup backup policy is available in the NetApp section.

If your NDMP policies include combinations of filers from NetApp, Isilon, and filers from other vendors, only the NetApp and Isilon filers use the Accelerator option. Messages in the job details identify which filers use the Accelerator option and when the option is used. More information about these job detail messages is available:

See [“Accelerator messages in the NDMP backup job details log”](#) on page 187.

Note: Unlike non-accelerated NDMP backups, accelerated NDMP backups do not use NDMP dump levels 0-9 to determine changed files. Instead, `BASE_DATE` and `DUMP_DATE` are used to determine changed files. `BASE_DATE` provides the timestamp of the most recent full or incremental backup. `DUMP_DATE` provides the timestamp of the currently running backup. Only the data that has changed between the `BASE_DATE` and the `DUMP_DATE` is backed up when Accelerator for NDMP is enabled.

Dump level messages from the filer continue to be included in the job detail log. However, the message `please ignore references to LEVEL in future messages` also appears in the job details as a reminder that dump levels are not used with Accelerator for NDMP.

How Accelerator works with NDMP backups:

- **First full backup with Accelerator**
The first full NDMP backup job with the Accelerator option enabled is similar to a normal full backup. It may run slightly longer than a non-Accelerator backup. It backs up all of the data from the filer, provides a baseline backup image, and creates an initial track log.

Note: If you first enable Accelerator when the next scheduled backup is an incremental backup, NetBackup does not automatically trigger a full backup image, as is the case with NetBackup Accelerator for non-NDMP policies. With Accelerator for NDMP, incremental backups continue to run as scheduled. An initial track log is also created after the **Use Accelerator** option is enabled, and with NetApp filers, you should see faster incremental backups. The next full backup runs only when it is scheduled.

- **Incremental backups with Accelerator**
Subsequent incremental backup jobs back up only the data that changed since the last backup job.
- **Next full backups with Accelerator**

Subsequent full backup jobs back up only the data that changed since the last backup job. The track log is used to determine what data can be included from previous backups, including the previous full backup and all of the incremental backups that follow it. NetBackup then creates a full backup image that includes all of the filer's data.

- Forced rescan full backups with Accelerator
The **Accelerator forced rescan** option provides a safety net by establishing a new baseline for the next Accelerator backup. When you include this option, which is found on the policy's **Schedules** tab, all the data on the filer is backed up. This backup is similar to the first full backup with Accelerator; it provides a new baseline for the backups that follow. If you set up a weekly full backup schedule with the **Use Accelerator** option, you can supplement the policy with another schedule that enables **Accelerator forced rescan**. You can set the schedule to run every 6 months or whenever it is appropriate for your environment. With NetApp filers, expect backups with **Accelerator forced rescan** to run slightly longer than accelerated full backups. With Isilon filers, backups with **Accelerator forced rescan** may run as long as a first full backup with Accelerator. More information about these options is available:
 - See [“Attributes tab options for an NDMP policy”](#) on page 157.
 - See [“Schedules tab options for an NDMP policy with Accelerator for NDMP enabled”](#) on page 158.

About the track log for Accelerator for NDMP

The track log is a binary file that you should not attempt to edit. On occasion, Cohesity Technical Support may request the track log for troubleshooting purposes. Two copies of the track log exist in the following locations:

- Primary server:
UNIX: `/usr/openv/netbackup/db/track`
Windows: `install_path\NetBackup\db\track`
- Media server:
UNIX: `/usr/openv/netbackup/track`
Windows: `install_path\NetBackup\track`

You can manually delete track logs safely if any of the follow situations occur:

- You disable the **Use Accelerator** option.
- The backup selections are changed.
- The policy is renamed.

- The NDMP filer is removed from the policy.
- The storage server that is used to perform the backup is changed.
- The primary server that is used to control the backups is changed.

Navigate to the following locations to manually delete track logs for specific backup selections:

- Primary server:

UNIX:

```
/usr/opensv/netbackup/db/track/primary_server/storage_server/filer_name/  
policy/backup_selection
```

Windows:

```
install_path\NetBackup\db\track\primary_server\storage_server\filer_name\  
policy\backup_selection
```

- Media server:

UNIX:

```
/usr/opensv/netbackup/track/primary_server/storage_server/filer_name/  
policy/backup_selection
```

Windows:

```
install_path\NetBackup\track\primary_server\storage_server\filer_name\  
policy\backup_selection
```

How to redirect track logs for Accelerator for NDMP

Track log size is relative to the size and number of files in a backup. In some cases, you may need to relocate the track logs to a different volume because of space issues. In these cases, it is recommended that you "redirect" the track logs to a volume where there is sufficient disk space.

One copy of the track log exists on the primary server and another copy exists on a media server in the following directories:

- Primary server:

UNIX: `/usr/opensv/netbackup/db/track`

Windows: `install_path\NetBackup\db\track`

- Media server:

UNIX: `/usr/opensv/netbackup/track`

Windows: `install_path\NetBackup\track`

To redirect these directories, complete the appropriate procedures in this topic. After completion, the next Accelerator-enabled backup that is executed redirects the track logs it creates to the directory you specified.

To redirect the track log directories on UNIX systems:

1 Rename the track log directories to make backup copies:

- On the primary server:

```
# mv /usr/opensv/netbackup/db/track  
/usr/opensv/netbackup/db/track.sv
```

- On the media server:

```
# mv /usr/opensv/netbackup/track /usr/opensv/netbackup/track.sv
```

2 Copy the backup to a new location:

- On the primary server:

```
# cp -rp /usr/opensv/netbackup/db/track.sv/* <path to new  
destination directory for track logs>
```

- On the media server:

```
# cp -rp /usr/opensv/netbackup/track.sv/* <path to new  
destination directory for track logs>
```

3 Create symbolic links from track log directories to the desired locations. For example, if the desired directory is /voll/track, enter the following command:

- On the primary server:

```
# ln -s /voll/track /usr/opensv/netbackup/db/track
```

- On the media server:

```
# ln -s /voll/track /usr/opensv/netbackup/track
```

4 After you have verified that everything works properly, you can remove the backup track.sv directory to free up space on the original volume.

To redirect the track log directories on systems with Windows Server:

1 Rename the track log directories to make backup copies:

- On the primary server:

```
> move "install_path\NetBackup\db\track"  
"install_path\NetBackup\db\track.sv"
```

- On the media server:

```
> move "install_path\NetBackup\track"  
"install_path\NetBackup\track.sv"
```

2 Copy the backup to a new location:

- On the primary server:

```
> xcopy /e "install_path\NetBackup\db\track.sv" "<path to new destination directory for track logs>"
```
- On the media server:

```
> xcopy /e "install_path\NetBackup\track.sv" "<path to new destination directory for track logs>"
```
- 3 Before performing an Accelerator-enabled backup, use `mklink` to link the `<install_dir>\NetBackup\track` directory to the desired directory. For example, if the desired directory is `E:\track`, enter the following command:

```
> mklink /D "<install_dir>\NetBackup\track" E:\track
```
- 4 After you have verified that everything works properly, you can remove the backup `track.sv` directory to free up space on the original volume.

More information about Accelerator for NDMP is available:

See [“About NetBackup Accelerator for NDMP”](#) on page 181.

See [“About the track log for Accelerator for NDMP”](#) on page 184.

Accelerator messages in the NDMP backup job details log

This topic provides explanations of some specific messages that appear in an NDMP job details log when Accelerator for NDMP is enabled.

The messages in the NetBackup job details include messages that are generated directly from the filer. To find the messages from the filer, look for the NDMP host name in the message following the PID number as in the following example:

```
mm/dd/yyyy hh:mm:ss - Info ndmpagent (pid=10780) [NDMP_host_name]:  
Filetransfer: Transferred 146841088 bytes in 2.855 seconds  
throughput of 50231.929 KB/s
```

Note: Some messages that are generated directly from the filer, such as `filer volume is full`, may require your immediate attention. Consult the documentation for the filer to determine how to resolve any issues with the filer that are indicated by a message from the filer in the job details.

First Accelerator-enabled full backup

Messages similar to the following appear in the job details log for the first full NDMP backup that uses Accelerator for NDMP.

```
mm/dd/yyyy 1:28:47 PM - Info bpbrm(pid=3824) accelerator enabled
...
...
mm/dd/yyyy 1:28:53 PM - Info ndmpagent(pid=10556) accelerator
optimization is <off>, unable to locate accelerator tracklog
...
...
mm/dd/yyyy 1:29:05 PM - Info ndmpagent(pid=10556) accelerator sent
1310720 bytes out of 1310720 bytes to server, optimization 0.0%
```

Note the following items about messages for the first Accelerator-enabled full backup:

- `accelerator enabled`
This message indicates that the Accelerator option is being used.
- `accelerator optimization is <off>, unable to locate accelerator tracklog`
Because this is the first full backup, NetBackup creates a new track log. More information about the locations of the track log is available:
See [“NetBackup logs for Accelerator for NDMP”](#) on page 190.
- `accelerator sent 1310720 bytes out of 1310720 bytes to server, optimization 0.0%`
Because this is the first full backup, all data is backed up and no optimization occurs yet.

Subsequent Accelerator-enabled incremental backup

Messages similar to the following appear in the job details log for subsequent incremental NDMP backups that use Accelerator for NDMP.

```
mm/dd/yyyy 2:01:58 PM - Info ndmpagent(pid=8652) accelerator
optimization is <on>
mm/dd/yyyy 2:01:58 PM - Info ndmpagent(pid=8652) BASE_DATE will be
used to determine changed files for accelerator
mm/dd/yyyy 2:01:58 PM - Info ndmpagent(pid=8652) please ignore
references to LEVEL in future messages
...
...
mm/dd/yyyy 2:14:14 PM - Info ndmpagent(pid=10044) accelerator sent
1104896 bytes out of 100310720 bytes to server, optimization 15.7%
```

Note the following items about messages for the subsequent incremental accelerator backups:

- `accelerator optimization is <on>`

This message indicates that a track log exists and the backup shall perform with the Accelerator option.

- `BASE_DATE` will be used to determine changed files for accelerator and please ignore references to `LEVEL` in future messages

These messages are a reminder that Accelerator for NDMP uses `BASE_DATE` and `DUMP_DATE` rather than dump levels to identify changed data. Messages that refer to dump levels come from the filer. However, the message to ignore references to `LEVEL` also appears in the job detail logs as a reminder that dump levels are not used with Accelerator for NDMP.

- `accelerator sent 1104896 bytes out of 100310720 bytes to server, optimization 15.7%`

This message provides the amount of data that was sent to the server and the percentage of optimization that was realized.

Next Accelerator-enabled full backups

Messages similar to the following appear in the job details log for subsequent full NDMP backups that use Accelerator for NDMP.

```
mm/dd/yyyy 2:01:58 PM - Info ndmpagent(pid=8652) accelerator
optimization is <on>
mm/dd/yyyy 2:01:58 PM - Info ndmpagent(pid=8652) BASE_DATE will be
used to determine changed files for accelerator
mm/dd/yyyy 2:01:58 PM - Info ndmpagent(pid=8652) please ignore
references to LEVEL in future messages
...
...
mm/dd/yyyy 1:40:27 PM - Info ndmpagent(pid=12244) accelerator sent
887296 bytes out of 1159725056 bytes to server, optimization 99.9%
```

Note the following items about messages for the subsequent incremental accelerator backups:

- `accelerator optimization is <on>`
This message indicates that a track log exists and the backup shall perform with the Accelerator option.
- `BASE_DATE` will be used to determine changed files for accelerator and please ignore references to `LEVEL` in future messages

These messages are a reminder that Accelerator for NDMP uses `BASE_DATE` and `DUMP_DATE` rather than dump levels to identify changed data. Messages that refer to dump levels come from the filer. However, the message to ignore references to `LEVEL` also appears in the job detail logs as a reminder that dump levels are not used with Accelerator for NDMP.

- accelerator sent 887296 bytes out of 1159725056 bytes to server, optimization 99.9%
 This message provides the amount of data sent to the server and the percentage of optimization that was realized.

Accelerator-enabled forced rescan full backup

Messages similar to the following appear in the job details log for full NDMP backups that use Accelerator for NDMP with the **Accelerator forced rescan** option.

```
mm/dd/yyyy 2:13:43 PM - Info bpbrm(pid=8628) Accelerator enabled
backup with "Accelerator forced rescan", all data will be scanned and
processed.Backup time will be longer than a normal Accelerator enabled
backup.
...
...
mm/dd/yyyy 2:13:46 PM - Info ndmpagent(pid=10044) accelerator
optimization is <on> but 'forced rescan' is enabled
```

Note the following items about messages for accelerator forced rescan backups:

- Accelerator enabled backup with "Accelerator forced rescan", all data will be scanned and processed. Backup time will be longer than a normal Accelerator enabled backup **and** accelerator optimization is <on> but 'forced rescan' is enabled

These messages indicate that a forced rescan is enabled and that the job shall run longer than a normal Accelerator full backup. Though accelerator optimization is on, the job may run slightly longer than accelerated full backups.

NetBackup logs for Accelerator for NDMP

Accelerator for NDMP does not require its own log directory. Instead, messages appear in standard NetBackup log files. [Table 19-1](#) lists the standard NetBackup log files in which messages for Accelerator for NDMP appear.

Table 19-1 NetBackup logs that may contain Accelerator for NDMP information

Log directory	Resides on
UNIX: /usr/opensv/netbackup/logs/ndmpagent Windows: <i>install_path</i> \NetBackup\logs\ndmpagent	NetBackup media server

Table 19-1 NetBackup logs that may contain Accelerator for NDMP information (*continued*)

Log directory	Resides on
UNIX: /usr/opensv/netbackup/logs/bpbrm Windows: <i>install_path</i> \NetBackup\logs\bpbrm	NetBackup media server
UNIX: /usr/opensv/netbackup/logs/bptm Windows: <i>install_path</i> \NetBackup\logs\bptm	NetBackup media server
UNIX: /usr/opensv/netbackup/logs/bpfis Windows: <i>install_path</i> \NetBackup\logs\bpfis	NetBackup media server
UNIX: /usr/opensv/netbackup/logs/bpcd Windows: <i>install_path</i> \NetBackup\logs\bpcd	NetBackup primary server
UNIX: /usr/opensv/netbackup/logs/bprd Windows: <i>install_path</i> \NetBackup\logs\bprd	NetBackup primary server
UNIX: /usr/opensv/netbackup/logs/bpdbm Windows: <i>install_path</i> \NetBackup\logs\bpdbm	NetBackup primary server

To create the log directories, run the following command on the NetBackup servers and backup host:

On Windows:

```
install_path\NetBackup\logs\mklogdir.bat
```

On UNIX/Linux:

```
/usr/opensv/netbackup/logs/mklogdir
```

Remote NDMP and disk devices


This chapter includes the following topics:

- [About remote NDMP and disk devices](#)
- [Configuring remote NDMP](#)

About remote NDMP and disk devices

This remote NDMP feature involves backing up NAS data (Network Attached Storage) to a storage device that is configured on a NetBackup media server. NetBackup supports disk devices on the media server.

The following figure shows the main components for NDMP backup to disk storage.

Figure 20-1 NDMP backup to a storage unit on media server (remote NDMP)

Configuring remote NDMP

Configure NetBackup to back up data to either disk storage or tape storage units that are attached to a NetBackup media server. Only NDMP-specific steps are described.

To configure NDMP backups to disk storage or tape storage units

- 1 Authorize the NetBackup server to access the NDMP hosts that you want to back up.

Do the following on the NetBackup media server:

- Enter the name of the NDMP server (NAS filer) to back up. This NDMP host name is case-sensitive.
- Repeat the previous step for each NDMP host that the NetBackup server backs up.

- If you plan to create snapshots using the Snapshot Client NAS_Snapshot method, do the previous step on the primary server (not on the media server).
- 2** In a similar way, you can configure devices for remote NDMP (disks, or tape drives and robots, on the media server).

Note the following items:

- Do not use the device configuration procedure that is described for configuring NDMP-attached devices. Instead, configure the disk, robots, and drives the same way as the ordinary NetBackup devices are configured. See the [NetBackup Administrator's Guide, Volume I](#).
 - Tape drives can be shared using the Shared Storage Option (SSO) of NetBackup. The drives can be shared as both NDMP drives and non-NDMP drives. See [“About the Shared Storage Option \(SSO\) with NetBackup for NDMP”](#) on page 195.
- 3** Create a disk or Media Manager storage unit for the drive(s). The storage unit type must be Disk or Media Manager, not NDMP.
- For details on storage units, refer to the [NetBackup Administrator's Guide, Volume I](#).
- 4** Create an NDMP-type policy.
- See [“About creating an NDMP policy”](#) on page 156.

Using the Shared Storage Option (SSO) with NetBackup for NDMP


This chapter includes the following topics:

- [About the Shared Storage Option \(SSO\) with NetBackup for NDMP](#)
- [Setting up SSO with NetBackup for NDMP](#)

About the Shared Storage Option (SSO) with NetBackup for NDMP

The following figure shows a robotic library on a SAN that can share its drives between two NetBackup for NDMP servers and two NDMP hosts. Drive sharing requires a license for the Shared Storage Option. A SAN is not required.

Figure 21-1 NDMP backup using Shared Storage Option



For each robot, either a NetBackup media server or an NDMP server (not both) can handle robotic control.

Setting up SSO with NetBackup for NDMP

This topic describes the steps for setting up access to a drive that is shared between NDMP and NetBackup servers.

For a more complete discussion of SSO, refer to the [NetBackup Administrator's Guide, Volume II](#).

This procedure assumes that the following conditions are true:

- The prerequisites for SSO have been met, as described in the [NetBackup Administrator's Guide, Volume II](#).
- All physical devices, including the NDMP host, are correctly connected to the network.
- NetBackup for NDMP supports the NDMP host.
 See [“About NAS appliances support”](#) on page 198. for information about supported NDMP operating systems and NAS vendors. The topic also contains configuration and troubleshooting help for particular NAS systems.
 The [NetBackup Compatibility List for all Versions](#) indicates which versions of vendor software support SSO for NDMP. The NAS systems (hardware) do not

provide the support; the proper software version provides it. For a list of the features and software releases for each NAS vendor, for SSO support, and for the NetBackup versions that support these vendors, refer to the *NetBackup Compatibility List for all Versions*.

To set up an SSO with NetBackup for NDMP

- 1** Configure NetBackup access to the NDMP host.

See [“Authorizing NetBackup access to a NAS \(NDMP\) host”](#) on page 150.

- 2** Verify that the NDMP host can access the required robots and drives.

To verify NDMP host access to the required devices, run the following commands on a NetBackup media server that is authorized to access the host:

```
tpautoconf -verify ndmp_host_name  
tpautoconf -probe ndmp_host_name
```

The `-verify` option verifies that the NetBackup server can access the NDMP host. The `-probe` option lists the devices that are visible to the NDMP host.

- 3** Use web UI to configure the devices and storage units.

You must define an NDMP storage unit for each NDMP host that shares a drive. If all hosts have access to the shared drive(s), web UI can create these storage units automatically.

NAS appliance information for NDMP

This chapter includes the following topics:

- [About NAS appliances support](#)
- [Non-vendor-specific information](#)
- [Vendor-specific information](#)

About NAS appliances support

Here you can find information on the Network Attached Storage (NAS) appliances supported by NetBackup for NDMP. Also included are configuration tips and restrictions for each supported NAS appliance.

Note: Device configuration is described in the *NetBackup Device Configuration Guide*. Special notes on some configurations are included, but full device configuration is not included here.

Non-vendor-specific information

Supported operating systems

The NetBackup operating systems compatibility lists provide the most up-to-date list of supported operating systems, with notes on supported hardware platforms. NDMP compatibility is listed as a column in the NetBackup Server section for each operating system in the compatibility list:

<https://support.cohesity.com/s/article/article-100040093>

Note: NetBackup for NDMP is installed on a master or media server, not on a client.

NAS appliance versions

The information in this document may apply to many different NetBackup releases. To find if your device and its features are supported at a particular NetBackup level, refer to the NetBackup Hardware Compatibility List (HCL) for the version of NetBackup that you are running:

<https://support.cohesity.com/s/article/article-100040093>

Authorizing NetBackup access to the NDMP host

Use the `tpconfig` and `tpautoconf` commands as described under “Example Configuration Sequence” in each of the following sections of this document for particular NDMP hosts.

For NAS hosts that support Direct Access Recovery (DAR), NetBackup enables DAR by default.

About file restores to a different location

When restoring files from NetBackup backups of NAS filers that use older versions of the NDMP protocol (V2 and V3), the destination path for the restore must end with the original folder and file name, even if you are restoring the files to a different location. (This restriction does not apply to NAS filers that use the NDMP protocol V4.)

If the original backup path was `/vol/vol1/mydir/myfile`, the destination path for restore to a different location must end with `/mydir/myfile`. Otherwise, NetBackup appends `/mydir/myfile` to the end of the destination path.

For example, to restore `/vol/vol1/mydir/myfile` to a folder under `/vol/vol2/`, specify `/vol/vol2/mydir/myfile` as the destination.

For NAS filers that use NDMP V4, you can specify a different subfolder or file name. For example, `/vol/vol1/mydir/myfile` can be restored to `/vol/vol2/mydir2/myfile_restored`.

For any restrictions unique to a particular NAS filer, see the appropriate section in this document for your filer.

About NDMP Environment Variables

The NDMP protocol specification uses environment variables to control backup and restore operations. These variables are defined by each vendor individually, with some amount of adherence to a common, defined set. NetBackup controls the setting of some of those variables without the ability for a user to change them (such as `LEVEL` and `USER`), and some of them are modified by NetBackup settings and

policy configuration (such as `DIRECT` and `FILESYSTEM`). It is possible to change some of the variables in the include list of a NetBackup policy, such as variables for which NetBackup just passes a default value (`HIST`, `TYPE`) and for those variables that are vendor-specific.

With such a large set of possibly vendor-specific implementations of environment variables, it is impossible to provide an exhaustive list in this document. An attempt has been made to document those variables that were found during testing to have the most impact. Refer to vendor documentation for more information.

Vendor-specific information

This section includes the following topics:

- [Dell EMC Isilon](#)
- [Dell EMC VNX](#)
- [Dell EMC Unity](#)
- [EMC Celerra](#)
- [Hitachi HDI/VFP](#)
- [Hitachi NAS \(HNAS\)](#)
- [HP X9000 NAS](#)
- [Huawei OceanStor V3](#)
- [IBM System Storage Nxxxx](#)
- [NEC Storage NV series](#)
- [NetApp](#)
- [Nexenta](#)
- [Nexsan](#)
- [Oracle Axiom Series](#)
- [Oracle Solaris Server](#)
- [Stratus V Series](#)

Dell EMC Isilon

General Information

This information is provided to help you use NetBackup for NDMP with an Isilon system.

For further details, contact Dell.

Access Configuration

To enable NDMP

- 1 Access the management browser to sign in to the system.
- 2 Select **Backup > Configuration**, then select the username and password, and enable the NDMP service state.

NetBackup configuration

OneFS 7.1 introduced two new features for NDMP backups: Snapshot-based incrementals and unlimited incremental backups. These features are enabled and disabled with environment variables that are set in the **Backup selections** list in a NetBackup policy:

- `set BACKUP_MODE=SNAPSHOT`
Enables snapshot-based incrementals.
- `set LEVEL=10`
In OneFS 7.1, Isilon implemented a feature enabling unlimited incremental backups. This feature is enabled by setting the `LEVEL` environment variable to 10.
You can set the `LEVEL` environment for Differential Incremental backup schedules only. For Full or Cumulative Incremental schedules, the value is ignored.

Note: In OneFS 8.0, Isilon introduced NDMP restartable backups (checkpoints). NetBackup does not support this feature.

Troubleshooting

The NDMP logs can be found at `/var/log/isi_ndmp_d` on each node.

To monitor system status from the management browser, go to **Alerts** to view alert activity.

Dell EMC VNX

General Information

This information is provided to help you use NetBackup for NDMP with a EMC VNX Network Server.

- Documentation

For more information on your VNX Network Server, refer to the *EMC VNX series - Configuring NDMP Backups on VNX* guide, which can be downloaded from Dell's support site.

Access Configuration

To assign a user account name and password to one or more data movers

- 1 Log in to the Celerra Network Server Control Station as `nasadmin` and switch user to `root` by typing the following command:

```
$ su
```

You must use the `su` command. The `su -` command will fail.

Type the root password when prompted.

- 2 Choose one for the following methods to assign a user account and password to a data mover. Replace `<movernamename>` with the name of the data mover for which you want to assign a user account and password.

- Text method

```
# /nas/sbin/server_user <movernamename> -add -password ndmp
```

For example:

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

- MD5 password encryption method

```
# /nas/sbin/server_user <movernamename> -add -md5 -password ndmp
```

For example:

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

The output from the command that you entered should look similar to the following example, in which the data mover is `server_2`:

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

Enter a new password when prompted, then retype the new password to confirm it. The password you assign to the data mover can contain between six and eight characters. The user name must be `ndmp`. You can accept the default values for the other settings.

In the output, the two mandatory fields, `User ID (UID)` and `Group ID (GID)`, are integers. The Celerra Network Server uses UNIX-style UIDs and GIDs to record the ownership of files and directories. The UID of the root user is 0.

3 Repeat the preceding steps for each NDMP-host data mover.

Device configuration

Tips for control and configuration:

- After logging on to the Celerra Network Server Control Station, you can use the following commands:
 - `nas_version`
(Displays the Celerra version number.)
 - `server_devconfig`
(Queries the device configuration of the specified data mover.)
- Make sure that the backup is created from a snapshot. See the following section: See [“Specifying snapshot-based backups in an NDMP policy”](#) on page 208.

Tips for robot and media discovery:

- Check that each Data Mover is recognizing the robot/media devices by entering the `server_devconfig` command from the Control Station.

For example, the following command queries the device configuration of the specified data mover (`server_2`):

```
server_devconfig server_2 -list -probe -scsi -nondisks
```

Example output:

```
server_2 :
SCSI non-disk devices :
chain= 0, scsi-0
symm_id= 0 symm_type= 0
tid/lun= 15/15 type= disk val= -99 info= 52658653C310 diskerr=
-1
chain= 1, scsi-1
symm_id= 0 symm_type= 0
tid/lun= 15/15 type= disk val= -99 info= 52686043C320 diskerr=
-1
chain= 2, scsi-2 : no devices on chain
chain= 3, scsi-3
symm_id= 0 symm_type= 0
tid/lun= 0/0 type= jbox info= HP C5173-7000 3.04
tid/lun= 1/0 type= tape info= QUANTUM DLT7000 2560q`
tid/lun= 2/0 type= tape info= QUANTUM DLT7000 2560q`
```

Troubleshooting

EMC VNX logs are located on each data mover. For example, to access the server_2 data mover log file, enter the following at the VNX Network Server Control Station:

```
server_log server_2
```

Other

Known restrictions

- The user name used with the `tpconfig` command must be defined as `ndmp` for each data mover.
While `tar`, `dump`, and `vbb` are all supported data types, Veritas recommends the use of `dump` or `vbb` instead of `tar`. Refer to the following tech note for more information:
<https://support.cohesity.com/s/article/article-100031122>
- If you specify an incorrect path name in the NetBackup policy's Backup Selections list (file list), the entire backup fails with status code 99, "NDMP backup failure."

Dell EMC Unity

General Information

This information is provided to help you use NetBackup for NDMP with the Dell EMC Unity system.

Access Configuration

After a NAS server is created, or during the NAS server creation, the **Protection & Events** tab on the **Edit NAS Server** page allows the user to **Enable NDMP** and **Change Password**.

Device configuration

1 Log in to command processor over SSH as service.

This will show the attached devices.

```
> svc_nas [NAS_server|ALL] -devconfig -probe -nondisks -all
```

An example:

```
> svc_nas dsunity001dm01 -devconfig -probe -nondisks -all
dsunity001dm01 :
SCSI devices :
chain=0, scsi-0 : no devices on chain
chain=1, scsi-1 : no devices on chain
chain=2, scsi-2
tid/lun= 0/0, type= tape, info=
tid/lun= 0/1, type= jbox, info=
chain=3, scsi-3 : no devices on chain
chain=4, scsi-4 : no devices on chain
chain=5, scsi-5 : no devices on chain
chain=6, scsi-6 : no devices on chain
chain=7, scsi-7 : no devices on chain
chain=8, scsi-8 : no devices on chain
chain=9, scsi-9 : no devices on chain
chain=10, scsi-10 : no devices on chain
chain=11, scsi-11 : no devices on chain
chain=12, scsi-12 : no devices on chain
chain=13, scsi-13 : no devices on chain
chain=14, scsi-14 : no devices on chain
chain=15, scsi-15 : no devices on chain
```

2 Using the information shown, configure the devices to 1|ALL NAS_servers:

```
> svc_nas {<NAS_server_name> | ALL} -devconfig -create
-scsi [<chain_number>] {-nondisks|-all}
```

3 It is necessary to enable scsi reservations, as it is not enabled by default (as of version 4.4):

```
> svc_nas {<NAS_server_name> | ALL} -param -f NDMP -m
scsiReserve -v 1
```

4 The NAS_server will need to be rebooted to have the setting take effect:

```
> svc_shutdown -r now
```

EMC Celerra

General Information

This information is provided to help you use NetBackup for NDMP with an EMC Celerra Network Server.

- Documentation
For more information on your Celerra Network Server, refer to the Celerra Network Server Version 5.5 Documentation CD, which can be downloaded from EMC's Powerlink website.

Device configuration

Tips for control and configuration:

- After logging on to the Celerra Network Server Control Station, you can use the following commands:
 - `nas_version`
(Displays the Celerra version number.)
 - `server_devconfig`
(Queries the device configuration of the specified data mover.)
- Make sure that the backup is created from a snapshot. See the following section: See [“Specifying snapshot-based backups in an NDMP policy”](#) on page 208.

Tips for robot and media discovery:

- Check that each Data Mover is recognizing the robot/media devices by entering the `server_devconfig` command from the Control Station.
For example, the following command queries the device configuration of the specified data mover (`server_2`):

```
server_devconfig server_2 -list -probe -scsi -nondisks
```

Example output:

```
server_2 :  
SCSI non-disk devices :  
chain= 0, scsi-0  
symm_id= 0 symm_type= 0  
tid/lun= 15/15 type= disk val= -99 info= 52658653C310 diskerr=  
-1  
chain= 1, scsi-1  
symm_id= 0 symm_type= 0  
tid/lun= 15/15 type= disk val= -99 info= 52686043C320 diskerr=  
-1  
chain= 2, scsi-2 : no devices on chain
```

```
chain= 3, scsi-3
symm_id= 0 symm_type= 0
tid/lun= 0/0 type= jbox info= HP C5173-7000 3.04
tid/lun= 1/0 type= tape info= QUANTUM DLT7000 2560q`
tid/lun= 2/0 type= tape info= QUANTUM DLT7000 2560q`
```

Troubleshooting

EMC Celerra logs are located on each data mover. For example, to access the server_2 data mover log file, enter the following at the Celerra Network Server Control Station:

```
server_log server_2
```

Other

Known restrictions

- The user name used with the `tpconfig` command must be defined as `ndmp` for each data mover.
While `tar`, `dump`, and `vbb` are all supported data types, Veritas recommends the use of `dump` or `vbb` instead of `tar`. Refer to the following tech note for more information:
<https://support.cohesity.com/s/article/article-100031122>
- If you specify an incorrect path name in the NetBackup policy's Backup Selections list (file list), the entire backup fails with status code 99, "NDMP backup failure."

Information for Celerra Network Server version 5.5 software and later

EMC Celerra Network Server version 5.5 and later software supports file and directory exclude lists with wild cards, in the NetBackup policy's Backup Selections list (file list).

Celerra Network Server version 5.5 and later software also supports NDMP Volume Backup.

File and directory exclusion statements

In the policy's Backup Selections list, file and directory exclusion statements can be used with the `set` directive. The statements are named `EMC_EFILE[01-05]` for file exclusion, and `EMC_EDIR[01-05]` for directory exclusion, in the following form (see examples below):

```
set EMC_EFILExx=file_exclusion_statement
set EMC_EDIRxx=directory_exclusion_statement
```

where *xx* is a two-digit number. For restrictions in the use of these statements and other details such as the use of wildcards, refer to the Celerra Network Server Version 5.5 Documentation CD.

In the following examples of file and directory exclude statements in a NetBackup Backup selections list, the backup of */fs2* will not include the files and directories specified by the `EMC_EDIR` and `EMC_EFILE` statements:

```
set HIST=y
set TYPE=tar
set EMC_EDIR01=/fs2/l*
set EMC_EDIR02=/fs2/Ndmp*
set EMC_EDIR03=/fs2/NAS*
set EMC_EDIR05=/fs2/j*
set EMC_EFILE01=*tar
set EMC_EFILE03=*dat
set EMC_EFILE02=*dat
set EMC_EDIR04=/fs2/Millions
set UPDATE=y
/fs2
```

Specifying snapshot-based backups in an NDMP policy

To make sure that the backup is based on data that is transactionally consistent, the following statement should be the first entry in the NetBackup policy Backup Selections list:

```
set snapsure=yes
```

The Celerra Server creates a snapshot, and the backup is created from the snapshot. The snapshot is managed by the EMC Celerra, not by NetBackup.

Specifying NDMP volume backups (VBB)

The NetBackup policy's Backup Selections list can specify that the backup is performed using EMC's NDMP volume backup.

In the following example of NDMP volume backup entries in the Backup Selections list, */testfs* will be backed up using NDMP Volume Backup.

```
set snapsure=yes
set type=vbb
/testfs
```

For restrictions and other information on NDMP Volume Backup, refer to the “Configuring NDMP Backups on Celerra” technical module on the Celerra Network Server Version 5.5 Documentation CD. This CD is downloadable from the EMC Powerlink website.

Hitachi HDI/VFP

General Information

This information is provided to help you use NetBackup for NDMP with the Hitachi HDI/VFP system.

Access Configuration

Using SSH to access one of the processing nodes using the service account, run the following command:

```
sudo ndmppasswd root oldpasswd newpasswd newpasswd
```

Device configuration

- For a list of the devices that are attached to the system, run the following command:

```
sudo tapelist
```
- For a list of tape drives that are not yet configured for NDMP, run the following command:

```
sudo tapelist -D
```
- To add those drives for NDMP access, run the following command:

```
sudo tapeadd -a
```

NetBackup configuration

Note: The HDI/VFP system only supports the tar backup type. However, the default type that NetBackup uses is dump. Therefore, you must add the following variable to the NetBackup policy:

```
set TYPE=tar
```

Hitachi NAS (HNAS)

General Information

This information is provided to help you use NetBackup for NDMP with the Hitachi NAS (HNAS) system.

Access Configuration

For access configuration options, including setting the NDMP username and password, and enabling and disabling NDMP access, go to **Credential management** in web UI.

Device configuration

Once devices are attached to the HNAS, use the following sequence to configure them for use with NDMP:

- 1 Allow access to the devices:

```
backup-device-allow-access all
```

- 2 Assign the devices to an EVS:

```
•backup-device-set-evs <device #> [<EVS_Name|Any]
```

- 3 Refresh the list of devices that are available through NDMP:

```
ndmp-devices-update
```

- 4 If the drives are shared between multiple hosts, enable SCSI reservations on the devices with the following command:

```
ndmp-option reserve_devices all
```

NetBackup configuration

You may use the following supported environment variables in the file list:

- EXCLUDE

For example: `set EXCLUDE="*mp3,core"`

- FUTURE_FILES

For example: `set FUTURE_FILES=y`

- HIST

For example: `set HIST=n`

- `set LEVEL=i`

For Hitachi NAS, setting `LEVEL=i` instructs the device to take an incremental based off the most recent previous backup of any level.

You can set the `LEVEL` environment for Differential Incremental backup schedules only. For Full or Cumulative Incremental schedules, the value is ignored.

HP X9000 NAS

General Information

This information is provided to help you use NetBackup for NDMP with an HP StorageWorks X9000 NAS system.

- Documentation
For further details, refer to the following documentation:
 - *X9000 File Serving CLI Reference*
 - *X9000 Administration Guide*

Access Configuration

Once the tape libraries are detected and listed, the HP X9000 should be configured to be the NDMP server.

To configure NDMP parameters on the management console GUI

- 1 Select **Cluster Configuration** from the **Navigator**.
- 2 Select **NDMP Backup**.

The **NDMP Configuration Summary** shows the default values for the parameters.
- 3 Click **Modify** on the **Configure NDMP** dialog box to configure the parameters for your cluster. See online Help for a description of each field.

To configure NDMP parameters from the CLI, use the following command:

```
ibrix_ndmpconfig -c [-d IP1,IP2,IP3,...] [-m MINPORT] [-x MAXPORT]  
[-n LISTENPORT] [-u USERNAME] [-p PASSWORD] [-e {0=disable,1=enable}]  
-v {0=10}] [-w BYTES] [-z NUMSESSIONS]
```

The NDMP server starts automatically if NDMP sessions are enabled on the cluster. You can use the following command to start, stop, or restart the NDMP server on one or more file serving nodes:

```
ibrix_server -s -t ndmp -c { start | stop | restart} [-h SERVERNAMES]
```

Device configuration

Once the connection between HP X9000 and tape library is completed, it is essential that HP X9320 detect and list the tape libraries connected to it.

To view the tape and media changer devices currently configured for backups

- 1** Select **Cluster Configuration** from the **Navigator**.
- 2** Select **NDMP Backup > Tape Devices**.
- 3** If you add a tape or media changer device to the SAN, click **Rescan Device** to update the list. If you remove a device and want to delete it from the list, you will need to restart all of the servers to which the device is attached.

To view tape and media changer devices from the CLI, use the following command:

```
ibrix_tape -l
```

To rescan for devices, use the following command:

```
ibrix_tape -r
```

For more information, refer to the *HP StorageWorks X9320 Network Storage System Administration Guide*.

Troubleshooting

All X9000 IBRIX commands on CLI can be ran in the following path:

```
/usr/local/ibrix/bin.
```

The following logs are available for troubleshooting:

- Errors warning and configuration events:
`/usr/local/ibrix/log/fusionserver.log`
- Cluster events:
`/usr/local/ibrix/log/events.log`
- Configuration messages from IAD and statistic reporting:
`/usr/local/ibrix/log/iad.log`
- Kernel messages from IDE:
`/var/log/messages`
- NDMP logs
`/usr/local/ibrix/logs/ndmp/tracelog`

Huawei OceanStor V3

General Information

This information is provided to help with using NetBackup for NDMP with a Huawei OceanStor V3 system. For further details contact Huawei.

Access Configuration

NDMP Settings are found in the Device Manager for the Huawei system: **Settings > Storage Settings > File Storage Service > NDMP Settings**.

Device configuration

Once devices are connected to the system, run the following from the command line on the Huawei system to rescan for devices:

```
admin:/> change service ndmp_scanbus
```

Then, restart the NDMP service from the NDMP Settings window in the Device Manager.

NetBackup configuration

File systems are presented over NDMP as /fs?, where "?" is the file system ID. To determine the available file systems use the following command:

```
admin:/>show file_system general
```

				Available	
ID	Name	...	Capacity	...	Capacity
--	-----	...	-----	...	-----
0	NFS100G1	...	100.000GB	...	79.632GB
1	CIFS100G1	...	100.000GB	...	79.576GB
2	NFS100G2	...	100.000GB	...	79.632GB

Note: The `ALL_FILESYSTEMS` feature that was introduced in NetBackup 7.6 uses the `NDMP_CONFIG_GET_FS_INFO NDMP` command to get a list of file systems that the system presents over NDMP. The Huawei system supports this command, but it reports "/" as the only available file system. This is not a valid file system for use with NDMP. This means that the `ALL_FILESYSTEMS` and `VOLUME_EXCLUDE_LIST` file list directives are not supported with this system.

IBM System Storage Nxxxx

General Information

This information is provided to help you use NetBackup for NDMP with an IBM System Storage Nxxxx filer.

For further details, refer to the following documentation:

- *Data ONTAP Command Reference Guide*
- *Data ONTAP System Administrator's Guide*

Device configuration

Tips for robotic devices

- To display the robot device file, sign on to the IBM Nxxxx host and enter the following command:

```
sysconfig -m
```

The device names in the output are in the format `mcN`, where `N` is 0 or higher.

Example `sysconfig` output:

```
Medium changer (6a.4) HP C6280-7000
mc0 - medium changer device
```

Tips for tape drives

- To display tape device files, sign on to the IBM Nxxxx host and enter the following command:

```
sysconfig -t
```

Always use the drive names that begin with `nr` (such as `nrst0a`) because these are the non-rewinding devices.

Example `sysconfig` output:

```
Tape drive (6a.5) Quantum DLT7000
rst0l - rewind device,          format is: 81633 bpi 40 GB (w/comp)
nrst0l - no rewind device,      format is: 81633 bpi 40 GB (w/comp)
urst0l - unload/reload device, format is: 81633 bpi 40 GB (w/comp)
rst0m - rewind device,          format is: 85937 bpi 35 GB
nrst0m - no rewind device,      format is: 85937 bpi 35 GB
urst0m - unload/reload device, format is: 85937 bpi 35 GB
rst0h - rewind device,          format is: 85937 bpi 50 GB (w/comp)
nrst0h - no rewind device,      format is: 85937 bpi 50 GB (w/comp)
urst0h - unload/reload device, format is: 85937 bpi 50 GB (w/comp)
rst0a - rewind device,          format is: 85937 bpi 70 GB (w/comp)
nrst0a - no rewind device,      format is: 85937 bpi 70 GB (w/comp)
urst0a - unload/reload device, format is: 85937 bpi 70 GB (w/comp)
```

Troubleshooting

The logs on the IBM Nxxxx filer must be viewed through an NFS or CIFS mount point. On the IBM filer, general messages appear in `/etc/messages`.

Other

- The NDMP service is controlled by the Data ONTAP administrative interface or the following commands:

`ndmpd on`(Starts the NDMP service.)

`ndmpd off`(Stops the NDMP service.)

`ndmpd status`(Displays the status of the NDMP service including any active NDMP sessions.)

`ndmpd probe session-number`(Displays details about the specified session.)

- By default, the NDMP service is not started at boot time. To start it, add the following line to the end of the `/etc/rc` file on the IBM system:

`ndmpd on`

- To determine the number of objects in a volume, enter the following command:

`maxfiles`

Known restrictions

- The user name used with the `tpconfig` command must be defined as **root** for each data mover.
- If you eject a tape from an IBM Nxxxx-attached drive and then try to open the device, it will reload the tape. This happens when the device is still UP and the NetBackup automatic-volume-recognition daemon (`avrd`) polls it.

NEC Storage NV series

General Information

This information is provided to help you use NetBackup for NDMP with an NEC Storage NV Series file server.

For more information on the NEC Storage NV Series, refer to the following documentation:

- *NEC Storage NV Series Software - Users Guide*
- *NEC Storage NV Series Software - Maintenance Manual*

For further details, contact NEC Corporation.

Access Configuration

To enable the NDMP option Program Package (PP), use a browser to start the package installer. See the *NEC Storage NV Series Software - Maintenance Manual* for more details.

Device configuration

- Robot

To find the robot device name, use the `telnet` command to log in to the NEC Storage NV system. Then run the following command:

```
dmesg | grep "scsi generic"
```

Example output:

```
Attached scsi generic sg0 at scsi0, channel 0, id 0, lun 0, type  
8
```

The robot device is `/dev/sg0`.

- **Drives**

To find the tape device names, log in to the NEC Storage NV system. Then enter the following:

```
dmesg | grep "scsi tape"
```

Example output:

```
Attached scsi tape st0 at scsi0, channel 0, id 1, lun 0  
Attached scsi tape st1 at scsi0, channel 0, id 2, lun 0
```

The tape device names, to be entered in the NetBackup web UI are `/dev/nst0` for tape drive 1 and `/dev/nst1` for tape drive 2. Always use the drive names that begin with "n" because these are the non-rewinding devices.

NetBackup configuration

The following directive must be placed at the start of the NetBackup policy's **Backup Selections** tab (file list):

```
set XFS=yes
```

This directive must be specified for all NetBackup backups of the NEC Storage NV Series, otherwise the backup will fail. Note that the `set XFS=yes` directive must be specified for both XFS and XFSFW file systems.

To use snapshots for NDMP backups, add the following directive to the file list:

```
set SANPSHOT=y
```

Troubleshooting

To enable debugging for NDMP, log in to the NEC Storage NV Series and add the following lines to the `/etc/sysconfig/ndmpd` file:

```
LOGFILE=/var/dumpfile/ndmpd  
DEBUG=yes  
LEVEL=65535
```

Debug logs are located in the `/var/dumpfile/ndmpd` directory.

Other

Known restrictions

- The NEC Storage NV Series supports the NDMP protocol version V2 only.
- The NEC Storage NV Series can back up only file systems, not subdirectories.
- Only one backup or restore can be running per file system. For example, if a backup job is currently backing up `/export/sxfs/vol1`, another attempt to back up or restore `/export/sxfs/vol1` at the same time will fail.
- A second backup of the same file system could fail if started too soon after the first backup of that file system. This is because a backup job needs time to delete the snapshot after completion of the backup. Until the snapshot is deleted, the second backup of the same file system cannot start. The same is true for restores: a restore of a file system could fail if started too soon after a previous restore of that file system.

NetApp

General Information

This information is provided to help you use NetBackup for NDMP with a NetApp Network Attached Storage (NAS) filer.

For further details, refer to the following documentation or contact NetApp.

- *Data ONTAP Command Reference Guide*
- *Data ONTAP System Administrator's Guide*
- Models
- Documentation

Device configuration

Tips for control and configuration

- For NDMP devices to share tape drives, tape reservation must be enabled in the ONTAP software on the filer and on NetBackup. You can use either SCSI persistent reservation or SCSI reservation. If you want to share tape drives, note that the drive itself must support one of these types of reservation. To enable SCSI reservation in Data ONTAP, enter either of the following at the ONTAP command line on the filer:

```
options tape.reservations scsi
options tape.reservations persistent
```

To enable the SCSI reservation in the NetBackup web UI, open **Hosts > Host properties**. Select the media server and click **Edit media server**. Then click **Media servers**. Ensure that you select the same type of SCSI reservation as you set on the filer.

- In ONTAP 8.0, both ONTAP 7 mode and ONTAP 10 mode are combined into a single release. You cannot run both modes on the same filer concurrently.
- The NDMP service is controlled by means of the Data ONTAP administrative interface or the following commands:


```
ndmpd on (Starts the NDMP service)
ndmpd off (Stops the NDMP service)
ndmpd status (Displays the status of the NDMP service including any active NDMP sessions)
ndmpd probe session-number (Displays details about the specified session)
```
- By default, the NDMP service is not started at boot time. To start it, add the following line to the end of the `/etc/rc` file on the NetApp system:


```
ndmpd on
```
- To determine the number of objects in a volume, enter the following command:


```
maxfiles
```

Tips for robotic devices

- To display the robot device file, sign on to the NetApp host and enter the following command:


```
sysconfig -m
```

The device names in the output are in the format `mcN`, where `N` is 0 or higher.

Example `sysconfig` output:

```
Medium changer (6a.4) HP C6280-7000
mc0 - medium changer device
```

Tips for tape drives

- To display tape device files, sign on to the NetApp host and enter the following command:


```
sysconfig -t
```

Always use the drive names that begin with `nr` (such as `nrst0a`) because these are the non-rewinding devices.

Example `sysconfig` output:

```
Tape drive (6a.5) Quantum DLT7000
rst0l - rewind device,          format is: 81633 bpi 40 GB (w/comp)
nrst0l - no rewind device,      format is: 81633 bpi 40 GB (w/comp)
urst0l - unload/reload device,  format is: 81633 bpi 40 GB (w/comp)
rst0m - rewind device,          format is: 85937 bpi 35 GB
nrst0m - no rewind device,      format is: 85937 bpi 35 GB
urst0m - unload/reload device,  format is: 85937 bpi 35 GB
rst0h - rewind device,          format is: 85937 bpi 50 GB (w/comp)
```

```
nrst0h - no rewind device,          format is: 85937 bpi 50 GB (w/comp)
urst0h - unload/reload device,      format is: 85937 bpi 50 GB (w/comp)
rst0a - rewind device,              format is: 85937 bpi 70 GB (w/comp)
nrst0a - no rewind device,          format is: 85937 bpi 70 GB (w/comp)
urst0a - unload/reload device,      format is: 85937 bpi 70 GB (w/comp)
```

NetBackup configuration

- If you eject a tape from a NetApp-attached drive and then try to open the device, it will reload the tape. This happens when the device is still UP and the NetBackup automatic-volume-recognition daemon (`avrd`) polls it.
- Image Backup (formerly referred to as SnapMirror to Tape or SMTape) is a Data ONTAP 8.0 feature that backs up an entire volume as a single file. Before Data ONTAP 8.0, the feature was called SMTape, and it required the customer to obtain a Product Variance Request from NetApp.

Note: Because Image Backup backs up an entire volume as though it is a single file, only the entire volume can be restored, not individual files within that volume.

To enable Image Backup, enter the following environment variables in the **Backup Selections** tab (file list) of the NetBackup policy:

```
SET type = smtape
SET SMTAPE_DELETE_SNAPSHOT = Y
/volume_to_back_up
```

Explanation of the variables

- `SET type = smtape`
Specifies the image backup feature.
- `SET SMTAPE_DELETE_SNAPSHOT = Y`
Deletes the snapshot after the backup completes. A snapshot is taken of the volume before the backup is written to tape. Deleting the snapshot saves storage space.
- `/volume_to_back_up`
Specifies the volume you want to back up, such as `/vol/vol1`.

Note: This feature is not currently supported with NDMP backups from NetApp storage configured with NetBackup Replication Director.

Troubleshooting

The logs on the NetApp filer must be viewed through an NFS or CIFS mount point. On the NetApp filer, general messages appear in `/etc/messages`.

Other

Known restrictions

- The user name used with the `tpconfig` command must be defined as `root` for each data mover.
- When restoring files, if the NetApp filer does not use Direct Access Recovery (DAR), the destination path that you specify for the restore must end with the original folder and file name. If the original backup path was `/vol/vol1/mydir/myfile`, the destination path for the restore must end with `/mydir/myfile`. Otherwise, NetBackup appends `/mydir/myfile` to the end of the destination path.

For more details on DAR, and to determine whether DAR has been disabled in NetBackup, refer to the *NetBackup for NDMP Administrator's Guide*.

Using NetBackup with NetApp's Data ONTAP 8.2 cluster mode

In the Clustered Data ONTAP (cDOT) 8.2 release, NetApp released an NDMP extension called Cluster Aware Backup (CAB). This extension allows backing up a Vserver (virtual server) or storage virtual machine (SVM) as an NDMP host (client) in a NetBackup policy. It is the default in new installs of ONTAP 8.2 and beyond. In environments where a cluster is upgraded from an older version of ONTAP, or in environments running multiple ONTAP versions, the behavior is to use node names as the NDMP host name. This is configurable using the following ONTAP command:

```
system services ndmp node-scope-mode [on|off]
```

Where `on` is Node-scope NDMP mode, and `off` is Vserver Aware NDMP mode.

ONTAP 8.2 C-mode (cluster mode) allows volumes to be moved from one node to another node in a cluster. Movement of volumes is performed to provide non-disruptive operations, high availability (failover), and resource balancing. NetApp automatically moves volumes during a failover. However, moving volumes to another node for maintenance or to provide load balancing will be performed by the NetApp storage admin.

NetBackup supports the CAB extension. There are important considerations when configuring NetBackup to protect a NetApp Clustered Data ONTAP environment running in either node-scope NDMP mode or Vserver Aware NDMP mode.

In NetBackup, data is tracked by client name, which is the NDMP hostname that is used to access the data. In cDOT, data is associated with a Vserver and hosted on a physical node. These things need to be considered when configuring NetBackup.

Another consideration is the availability of resources from the cluster, as illustrated in [Table 22-1](#):

Table 22-1 Availability of resources from the cluster

Interface Type	Volume Visibility			Tape Drive Visibility		
	Node-scope-mode	Vserver-mode		Node-scope-mode	Vserver-mode	
		Non-CAB Aware NetBackup	CAB-Aware NetBackup		Non-CAB Aware NetBackup	CAB-Aware NetBackup
Cluster Management	N/A	All volumes on the same node as LIF	All volumes in cluster	N/A	N/A	All tape drives in cluster
Intercluster	N/A	All volumes on the same node as LIF	All volumes in cluster	N/A	N/A	All tape drives in cluster
Vserver	N/A	Volumes in Vserver and hosted on same node as LIF	All volumes in Vserver	N/A	N/A	N/A
Node Name	All volumes on node	N/A	N/A	All tape drives on node	N/A	N/A

An Intercluster LIF is very similar to a Cluster Management LIF, but needs to be configured on each node of a cluster.

Using a node name as the NDMP client name in all versions of NetBackup

With this method, node-scope-mode is enabled on the cluster, and the name of each node is provided as a client name in an NDMP policy in NetBackup.

Pros

- Volumes can be backed up to locally attached tape drives, instead of over a network connection (3-way).

- Using NetBackup 7.6 and higher, with the introduction of the `ALL_FILESYSTEMS` file list directive, it is not necessary to modify the NetBackup policy if a volume moves to another node.

Cons

- When a volume moves to another node, the moved volume and data is now tracked by that other node name in NetBackup. When performing a restore, NetBackup will display all backups from the current selected node. However, to restore data from backups taken when the volume was under a different node, you will need either to maintain a list of those prior nodes or search the other nodes in the cluster for that specific nodename and volume combination.
- Once a volume has been moved, three-way backups from the current node to the original node may occur depending on policy and storage unit configuration.
- If using a version of NetBackup before 7.6, or not using the `ALL_FILESYSTEMS` file list directive, and if a volume moves to a different node in the cluster, the NDMP host name in the NetBackup policy must be modified to that of the node now hosting the volume, using the following command:

```
/usr/opensv/netbackup/bin/admincmd/bpplclients policy_name -rename  
old_host_name new_host_name
```

Using a data Vserver LIF as the NDMP client name in non-CAB-aware versions of NetBackup

With this method, node-scope-mode is disabled on the cluster and a data Vserver name is configured as the client name in an NDMP policy in NetBackup

Pros

- Because the NetBackup catalog tracks backups by client name, it is easier to track down data from that Vserver when restoring.
- Assuming that care has been taken if a volume moves to another node, all volumes backed up using a Vserver name will be displayed when performing a restore.

Cons

- The cDOT filer will only send data from volumes that are associated with a Vserver and are hosted on the same node as the data Vserver LIF. Therefore, if only a volume moves to another node without its corresponding LIF, it will not be backed up. This is not an error. Users will need to very carefully monitor backups to ensure that all of their data was backed up.
- If a volume moves to another node, the vsver LIF must be moved to the other node as well using the following command:

```
net int migrate -vsver <vsver_name> -lif <vsver-LIF>  
-dest-node <dest-node> -dest-port <dest-port>
```

- Because a data Vserver cannot see any tape drives attached to the cluster, every backup of a data Vserver will be a three-way backup.

Using a cluster_mgmt vsriver LIF as the NDMP client name in non-CAB-aware versions of NetBackup

With this method, node-scope-mode is disabled on the cluster and a cluster_mgmt Vserver LIF name(s) are configured in an NDMP policy in NetBackup.

Pros

- A cluster_mgmt LIF can see every volume on the node, so with the proper configuration it is not necessary to change either the NetBackup policy or move a LIF if a volume moves to another node.
- This method is most like previous versions of ONTAP and non-CAB-aware versions of NetBackup.

Cons

- There is no visibility to tape drives in this configuration, so backups will still be 3-way.
- A cluster_mgmt LIF has to be created on each node for the admin vsriver.
- It is still necessary to carefully track data when requesting a restore to ensure the desired data is restored.

Using a cluster_mgmt Vserver LIF as the NDMP client name in CAB-aware versions of NetBackup

With this method, node-scope-mode is disabled on the cluster and a cluster_mgmt Vserver name is configured in an NDMP policy in NetBackup.

All volumes and all attached tape drives will be available through a single interface on the cluster.

Regardless of where a volume is located on the cluster a backup of that volume will be directed to a tape drive (if available) attached to the same node, resulting in a performance improvement

Nexenta

General Information

This information is provided to help you use NetBackup for NDMP with a NexentaStor system.

For further details, contact Nexenta.

Access Configuration

To enable NDMP

- 1 Log in to **Nexenta Management View**.
- 2 Go to **Settings > Misc. Services > NDMP Server**.
- 3 Enable NDMP service.

Troubleshooting

- To enable debugging, edit the `/lib/svc/method/svc-ndmp` file in line 43 to add `-d`. For example:

```
/usr/lib/ndmp/ndmpd -d 2>&1 &
```
- The NDMP logs can be found at `/var/log/ndmp` on each node.
- To monitor system status from the **Nexenta Management View**, go to **Alerts** to view alert activity.

Other

To enable NDMP DAR

- 1 Connect to the Nexenta host with a secure shell (SSH).
- 2 Run `!bash`.
- 3 Enter the following command:

```
ndmpadm set dar-support=yes
```
- 4 To verify the NDMP properties, run the following command:

```
ndmpadm get
```

In 4.0, DAR can also be enabled from the NDMP Configuration page of the user interface.

To enable NDMP DAR from the Nexenta Management View (4.0 only)

- 1 Log in to **Nexenta Management View**.
- 2 Go to **Settings > Misc. Services > NDMP Server > Configure**.
- 3 Select the DAR support option.

Nexsan

General Information

This information is provided to help you use NetBackup for NDMP with a Nexsan system.

For further details, contact Imation.

Access Configuration

To enable NDMP and set the user name and password

- 1 Log in to the Nexsan administration console.

Note: The Nexsan administration console requires Adobe Flash.

- 2 Click the system name on the top of the window, then expand the options pane by clicking the arrows in the bottom right of the window.
- 3 In the options pane, click on the **NDMP** tab
- 4 Click **Enable NDMP**.
- 5 Set the user name and password
- 6 Click **Apply**.

Troubleshooting

To enable debugging for NDMP

- 1 Log in to the Nexsan system.
- 2 Run the following command:

```
nstndmp set debug-enable=true
```

Logs can be viewed in the Nexsan administration console, under **System Events > View**.

Oracle Axiom Series

General Information

This information is provided to help you use NetBackup for NDMP with the Axiom system.

For further details on the Axiom storage system, consult the Oracle website:

<http://www.oracle.com/us/corporate/Acquisitions/pillardata/index.html>

Access Configuration

NDMP settings should be checked on the Oracle Axiom system and verified to correspond to the NetBackup defined configuration. From the Axiom Storage Manager GUI, navigate to **Data Protection > NDMP Backup Settings > Modify**

NDMP Backup Settings. This location also turns the NDMP port on and off on the Axiom system.

Device configuration

To manually detect locally attached tape devices on the Axiom system, use the Axiom Storage Manager GUI to navigate to **Data Protection > Tape Devices**, and then select **Check for Tape Devices** from the **Actions** pull-down menu in the middle of the window.

Troubleshooting

To access Axiom system logs

- 1 Log on to the Axiom Storage Manager GUI.
- 2 Click the **Support** button and then select **Collect System Information** under the **Tools** menu on the left.
- 3 From the **Actions** pull-down menu in the bottom-middle of the window, select **Collect System Information**.
- 4 Select the desired items from the **Collection Scope** list. Collecting system information can take several minutes depending on the items selected.
- 5 After collection of system information completes, select **Download Information to Client** from the **Actions** pull-down menu and select the location to save the System Information. (It will be a tar file.)

Oracle Solaris Server

General Information

This information is provided to help you use NetBackup for NDMP with an Oracle Solaris NDMP server.

For further details, contact Oracle.

Pre-requisites

To install the server software on any Solaris 11 server:

```
pkg install service/storage/ndmp
```

To allow control of locally-attached tape libraries, configure the sgen driver:

```
update_drv -a -I "scsiclass,08" sgen
```

```
reboot
```

Note: It is not recommended to run NetBackup on the same machine. NetBackup Tape Library control cannot run in conjunction with Solaris Tape Library control.

Service Configuration

To display the existing configuration:

```
ndmpadm get
```

The username and password for NDMP access is set when the service is enabled:

```
ndmpadm enable -a cram-md5 -u <username>
```

Enable DAR:

```
ndmpadm set dar-support=yes
```

Enable BSD-style drive access:

```
ndmpadm set drive-type=bsd
```

Configure volumes for NDMP export:

```
ndmpadm set fs-export=/<path1>,[/<path2>, etc.]
```

To start/stop the NDMP service:

```
svcadm [enable|disable] ndmpd
```

Troubleshooting

To display where the service logs:

```
# svcs -l ndmpd
fmri          svc:/system/ndmpd:default
name          NDMP Service
enabled       true
state         online
next_state    none
state_time    May 13, 2015 12:54:07 PM CDT
logfile       /var/svc/log/system-ndmpd:default.log
restarter     svc:/system/svc/restarter:default
contract_id   123
manifest      /lib/svc/manifest/system/ndmp.xml
dependency    require_all/error svc:/milestone/self-assembly-complete
              (online)
```

Other

If you specify an incorrect path name in the NetBackup policy's Backup Selections list (file list), the entire backup fails with status code 99, "NDMP backup failure."

Stratus V Series

General Information

This information is provided to help you use NetBackup for NDMP with a Stratus V Series system.

For further details on the Stratus V Series system, contact Stratus Technologies.

NetBackup configuration

The following directives must be placed at the start of the NetBackup policy's **Backup Selections** tab (file list):

```
SET TYPE=save
SET SAVE_OPTIONS='-backup'
```

To learn more about the additional directives available for the Stratus V Series, see the Stratus V Series documentation.

Troubleshooting

The Stratus VOS Enterprise Backup Agent creates the following log files. All are stored on the Stratus system in the `>system>ndmpd>log` directory:

- `ndmpd_log.YY_MM_DD.out`
- `save.YY_MM_DD.hh_mm_ss.process_id`
- `macro.YY_MM_DD.hh_mm_ss.process_id`

For more information about these files, refer to "Log Files" in the VOS Enterprise Backup Agent online documentation.

Other

Known restrictions

- The Stratus V Series does not support directory names greater than 32 characters. The limitation for files names is 255 characters.
- The Stratus V Series does not support CIFS or NFS for file system access. To access the Stratus file system, you must use SAMBA. To learn more about file system access to the Stratus V Series, refer to your Stratus documentation.
- The Stratus V Series uses its own operating system called VOS. To access the VOS operating system directly, you must use a terminal emulator such as TTWIN 3.
- The Stratus V Series supports the NDMP version 3 protocol only.

Backup and restore procedures

This chapter includes the following topics:

- [Performing a manual backup with an NDMP policy](#)
- [Perform an NDMP restore](#)

Performing a manual backup with an NDMP policy

The following procedures explain how the NetBackup administrator can perform the backup manually. Only a NetBackup administrator can initiate an NDMP backup. In the NetBackup Web UI, a user must have the **Administrator** role.

Perform a manual backup of NDMP (NetBackup Web UI)

To perform a manual backup of NDMP with the Web UI

- 1 On the primary server, open the NetBackup Web UI.
- 2 On the right, open **Protection > Policies**.
- 3 Locate and select the NDMP policy. Then click **Manual Backup**.
- 4 Select a schedule and then select the clients that you want to back up.

If you do not select any schedules, NetBackup uses the schedule with the highest retention level. If you do not select any clients, NetBackup backs up all configured NDMP hosts.

- 5 Click **OK** to start the backup.

Perform an NDMP restore

An administrator can perform a restore as follows:

- With the NetBackup Web UI, from the primary server.
- With the **Backup, Archive, and Restore** interface from a NetBackup primary server or media server.

NetBackup administrators can restore files to the original NDMP host or to a different NDMP host.

Note: User-directed restores of files are not allowed, because no NetBackup client software is installed on an NDMP host.

Restore NDMP (NetBackup Web UI)

To restore NDMP with the NetBackup Web UI

- 1 On the primary server, open the NetBackup Web UI.
- 2 On the left, select **Recovery**.
- 3 On the **Regular recovery** card, click **Start recovery**.
- 4 Select the following information and click **Next**.

Policy type	NDMP
Source client	Select the appropriate NDMP (NAS) host.
Destination client	Select the appropriate NDMP (NAS) host. The destination host must be an NDMP host that is compatible with the data format of the source. (The source and destination must be of the same NAS vendor type.)

- 5 NetBackup automatically displays the most recent backup. To select a different date range, click **Edit**.
- 6 Select the files or folders that you want to restore. Then click **Next**.
- 7 Select the recovery options that you want for the restore. Then click **Next**.

Warning: An NDMP restore always overwrites existing files.

Restore NDMP (Backup, Archive, and Restore interface)

To restore NDMP with the BAR interface

- 1 In the **Backup, Archive, and Restore** interface on a NetBackup server, click **Actions > Specify NetBackup Machines and Policy Type**.

- 2 For the server, select the NetBackup primary server.

If your configuration has multiple primary servers, specify the primary server that has the policy for the NDMP host that you plan to restore. If the server name is not in the pull-down list, use **Edit Server List** to add it.

- 3 For the source clients and destination clients, select the appropriate NDMP (NAS) hosts.

The destination host must be an NDMP host compatible with the data format of the source. (The source and destination must be of the same NAS vendor type.)

Warning: An NDMP restore always overwrites existing files.

If the hosts that you want are not available in the pull-down menu, use **Edit Client List** to add the client.

- 4 In the policy type field, select **NDMP**.

Troubleshooting

This chapter includes the following topics:

- [About NetBackup for NDMP logs](#)
- [General NetBackup for NDMP operating notes and restrictions](#)
- [NetBackup for NDMP troubleshooting suggestions](#)
- [About robot tests](#)

About NetBackup for NDMP logs

NetBackup uses two types of logging, unified logging and legacy logging. Both logging types are described in the "Using Logs and Reports" topic in the [NetBackup Troubleshooting Guide](#).

Note the following:

- All unified logs are written to `/usr/opensv/logs` (UNIX) or `install_path\logs` (Windows). Unlike legacy logging, you do not need to create logging directories.
- Use the `vxlogview` command to examine unified logs:
See "[Viewing NetBackup for NDMP logs](#)" on page 232.
On UNIX: `/usr/opensv/netbackup/bin/vxlogview`
On Windows: `install_path\NetBackup\bin\vxlogview`
Refer to the [NetBackup Troubleshooting Guide](#) for assistance in using the `vxlogview` command.
See also the `vxlogview` man page or the [NetBackup Commands Guide](#).

Viewing NetBackup for NDMP logs

The following procedure describes how to view NetBackup logs.

Note: The legacy and unified logging files can consume a lot of disk space. Delete the log files when you are finished and set logging to a lower level of detail.

To view the NetBackup logs

- 1** In the **NetBackup web UI**, select **Host > Host properties**.
- 2** Select **Logging** and set the **Global logging level** to 5.
- 3** Click **Apply** and then **OK**.
- 4** View the unified logging information in `/usr/opensv/logs` (UNIX) or `install_path\logs` (Windows) for the following processes:
 - `ndmpagent` (originator ID 134)
 - `ndmp` (originator ID 151)
 - `nbpem` (originator ID 116)
 - `nbjm` (originator ID 117)
 - `nbrb` (originator ID 118)
- 5** For `ndmpagent` logs, try the `vxlogview` command as follows:


```
/usr/opensv/netbackup/bin/vxlogview -I ndmpagent -d T,s,x,p
```
- 6** For `ndmp` logs, try the `vxlogview` command as follows:


```
/usr/opensv/netbackup/bin/vxlogview -I ndmp -d T,s,x,p
```
- 7** On the NetBackup for NDMP server, create `bptm`, and `bpbrm` legacy debug log folders in the `/usr/opensv/netbackup/logs` directory (UNIX) or `install_path\NetBackup\logs` folder (Windows):
 - `bpbrm`
 - `bpfis`
 - `bpmount`
 - `bptm`
 - `bppfi`

NetBackup writes legacy log files in these directories, if the directories exist.

NDMP backup levels

At the start of a debug log, you may see an entry titled `LEVEL`. This entry refers to an environment variable that NetBackup set based on the type of backup. Here is an example from a `bptm` log:

```
08:48:38.816 [22923] <2> write_data_ndmp: backup environment
values:
08:48:38.816 [22923] <2> write_data_ndmp: Environment 1:
TYPE=dump
08:48:38.816 [22923] <2> write_data_ndmp: Environment 2:
FILESYSTEM=/vol/vol0/2million
08:48:38.817 [22923] <2> write_data_ndmp: Environment 3:
PREFIX=/vol/vol0/2million
08:48:38.817 [22923] <2> write_data_ndmp: Environment 4: LEVEL=0
```

The NDMP backup level is modeled after UNIX dump levels. The backup level is a number in the range of 0 to 9.

An NDMP backup level of 0 is a full backup. A backup level greater than 0 is an incremental backup of all objects that were modified since the last backup of a lower level. For example, level 1 is a backup of all objects that were modified since the full backup (level 0). Level 3 is a backup of all objects that were modified since the last level 2 incremental.

Table 24-1 NetBackup backup types and corresponding NDMP backup levels

NetBackup backup types	NDMP backup levels
NetBackup Full	NDMP level 0
NetBackup Cumulative Incremental	NDMP level 1
NetBackup Differential Incremental	NDMP level (last level + 1, up to 9) See “About NAS appliances support” on page 198. for valid level values for your device. Some vendors support level values that are greater than 9..

More information is available on environment variables.

See [“About environment variables in the backup selections list”](#) on page 165.

General NetBackup for NDMP operating notes and restrictions

Before you try to troubleshoot a suspected problem, review the following operating notes:

- A tape that was created on an NDMP storage unit is in backup format. It cannot be restored from a non-NDMP storage unit. If you duplicate an NDMP backup image, the new copy is still in backup format. It cannot be used for restores on a non-NDMP storage unit.
- In the backup selections list for an NDMP policy, you can include only directory paths. Individual file names are not allowed. Wildcard characters are allowed in backup selections, though some limitations apply to some filers. More information about wildcards in NDMP backup selections is available:
See [“Wildcard characters in backup selections for an NDMP policy”](#) on page 160.
- In a NetBackup NDMP policy, you cannot include a path in the file list that is more than 1024 characters long. This limitation may be further restricted for certain vendors. See [“About NAS appliances support”](#) on page 198. for path name length information for specific filers.
- Observe the following restrictions to the use of the `ALL_FILESYSTEM` directive and the `VOLUME_EXCLUDE_LIST` directive:
 - A `VOLUME_EXCLUDE_LIST` statement may include a maximum of 256 characters. Create multiple `VOLUME_EXCLUDE_LIST` statements if necessary to avoid exceeding the limit of 256 characters. If you specify more than 256 characters, the volume list is truncated. A truncated statement may result in a backup job failure, and the error message `Invalid command parameter(20)` is displayed. `VOLUME_EXCLUDE_LIST` applies only to `ALL_FILESYSTEMS`. It does not apply to explicit backup selections or wildcard-based backup selections.
 - With NetBackup Replication Director, if the backup selection includes read-only volumes or full volumes, an NDMP backup job fails with the status code 20 (`Invalid command parameter(20)`). If you encounter a similar NDMP backup job error, review the `ostfi` logs to identify the volumes for which the failure occurred. You can use `VOLUME_EXCLUDE_LIST` statements with the `ALL_FILESYSTEMS` statement to exclude the read-only volumes and the volumes with insufficient space.

Note: This restriction applies only to NetBackup Replication Director environments.

More information about these directives is available:

See [“ALL_FILESYSTEMS and VOLUME_EXCLUDE_LIST directives”](#) on page 163.

- The NDMP protocol uses port 10000 for communication.
- On UNIX systems, the NetBackup `avrd` process uses Internet Control Message Protocol (ICMP) to ping NDMP hosts to verify network connectivity. This protocol is required for the NetBackup for NDMP product.
- If backup jobs or restore jobs are running slowly, verify that the network interface cards (NIC) are set to full duplex. Half-duplex often causes poor performance. For assistance viewing and resetting duplex mode for a particular NAS host, consult the documentation that the manufacturer provides. You may be able to use the `ifconfig` (or `ipconfig`) command, as explained in the [NetBackup Troubleshooting Guide](#).
- Do not perform incremental backups of the same NDMP data from two different policies. Incremental backups performed by one of the policies may be incomplete, because NDMP filers perform level-based incremental backups instead of time-based incremental backups. Consider the following example:

Policy A performs a full backup of `/vol/vol1` (level 0).

Policy B then performs a full backup of `/vol/vol1` (level 0). The filer now considers the policy B backup to be the last full (level 0) backup of `/vol/vol1`.

Policy A performs an incremental backup of `/vol/vol1` (level 1). The policy A incremental backup captures only the data that changed since the full backup that was done by policy B. The incremental backup misses any changes that occurred between the policy A full backup and the policy B full backup.

- NDMP restore jobs may complete successfully even though no data (0 KB) has been restored. This situation can occur when a target volume does not have enough space for an image you are trying to restore.
 - Workaround: Check the restore job details for entries similar to the following messages:

```
mm/dd/yyyy hh:mm:ss PM - Info ndmpagent(pid=11071) fas2050c1: RESTORE: We recommend that 19
inodes and 907620 kbytes of disk space be available on the target volume order to restore
this dump. You have 466260 inodes and 5316 kbytes of disk space on volume /vol/abc_15gb
mm/dd/yyyy hh:mm:ss PM - Info ndmpagent(pid=11071) fas2050c1: RESTORE: This restore will
proceed, but may fail when it runs out of inodes and/or disk space on this volume.
```

Confirm that the target volume does not have enough space for the restore image. If it does not, either free up enough space on the volume to complete the restore job successfully or specify a different restore volume.

NetBackup for NDMP troubleshooting suggestions

Try the following troubleshooting suggestions:

- Check the NetBackup All Log Entries report for information about the failed job.
- To verify that the appropriate services are running, use one of the following: the NetBackup Activity Monitor, the Windows control panel (on Windows systems), or the `bpps` command (UNIX systems).
- If NDMP host backups terminate with a status code of 154 (storage unit characteristics mismatch requests), the problem may be one of the following:
 - Verify that the NetBackup configuration is correct.
 - There may be a conflict between the policy type and storage unit type. (For example, if the policy type is Standard and the storage unit is of type NDMP.)
- If your NDMP backup fails with a status code of 99 (NDMP backup failure), no paths in your NDMP policy backup selections list were backed up. Check the NetBackup All Log Entries report for more information. A possible cause of this status is that none of the backup paths exist on the NDMP host.
 For more information about status code 99 and NDMP backup failures, refer to the following tech note:
<https://support.cohesity.com/s/article/article-100023868>
- NetBackup does not support client-side deduplication of NDMP hosts. The backup jobs fail if you try to use client-side deduplication for NDMP hosts.

Troubleshooting NDMP media and devices on Windows

To troubleshoot media and devices on Windows, try the following:

- For legacy logging, enable debug logging by creating `reqlib` and `daemon` directories in the `install_path\Volmgr\debug` directory on the NetBackup for NDMP server.
- Check the Windows Event Viewer Application log for troubleshooting clues.
 For more information on the **Event Viewer logging** option, refer to the [NetBackup Troubleshooting Guide](#).
- Use the **Activity Monitor** utility or the Windows control panel to verify that the **Media and Device Management** utilities are running.

- Drives can be unexpectedly set to the DOWN state.
 This action is due to communication problems between `avrd` on the NetBackup for NDMP server and the NDMP server application on the NDMP host. Some possible causes for the communication problems are:
 - Network cable on the NDMP host was unplugged.
 - NIS (Network Information System) problems on the NetBackup for NDMP server (NDMP client).
 - The NDMP host was halted for too long.

Note: Whatever the cause, if the `avrd` connection to the NDMP host fails, the drive is set to DOWN. It is not automatically set to UP when the communication problem is corrected.

Troubleshooting NDMP media and devices on UNIX

To troubleshoot media and devices on UNIX, try the following:

- Ensure that the `syslogd` logs debug messages relating to `ltid` and other device processes.
 For more information on `syslogd`, refer to the [NetBackup Troubleshooting Guide](#).
- Start `ltid` with the `-v` option. Check the system's syslog for troubleshooting clues.
- Use `vmops` to make sure that the appropriate daemons are running.
- Drives can be unexpectedly set to the DOWN state. This action is due to communication problems between `avrd` on the NetBackup for NDMP server and the NDMP server application on the NDMP host.
 Further details are available.
 See "[Troubleshooting NDMP media and devices on Windows](#)" on page 237.

Troubleshooting NDMP DirectCopy

When NetBackup enables NDMP DirectCopy for a backup image duplication, the NetBackup progress log includes the message "NDMP DirectCopy should be used." If NDMP DirectCopy was not enabled for the duplication, no specific messages about NDMP DirectCopy are listed in the progress log. For detailed messages (such as why NDMP DirectCopy was not used), consult the legacy debug logs for the admin log or the `bptm` log.

Refer to the [NetBackup Troubleshooting Guide](#) for information on legacy NetBackup logs.

Troubleshooting Direct Access Recovery (DAR) with NetBackup for NDMP

Note the following points when using Direct Access Recovery (DAR):

- DAR can be used when restoring NetBackup 4.5 or later backups. Starting with NetBackup 4.5, NetBackup stores the required DAR offset information on each backup.
- Backups must have been performed with the NetBackup catalog set to binary mode. If backups were made with the catalog set to ASCII mode, restores cannot use DAR. ASCII mode did not store the required DAR offset information on each backup. Note that all backups that were made before NetBackup 4.5 used ASCII catalog mode.

Note: Starting with NetBackup 6.0, all backups are in binary mode.

- To use DAR with NetBackup, the NDMP host you want to restore must support DAR. Some NDMP host vendors do not currently support DAR.

The following table lists the messages that may appear in the unified logs for `ndmpagent` (originator ID 134) on the NetBackup media server. These messages are also written to the progress log.

Table 24-2 DAR log messages

Message	Explanation
Data host does not support DAR recovery	The current NDMP host does not support DAR.
DAR disabled—continuing restore without DAR	DAR information is not available for the file.
DAR disabled—backup was performed before NB 4.5	The DAR feature can be used to restore the backups that NetBackup 4.5GA or later made. Starting with NetBackup 4.5GA, NetBackup stores the required DAR offset information on each backup. For pre-4.5GA NetBackup backups, restores cannot use DAR because the pre-4.5 versions did not store DAR offset information.
DAR disabled—NDMP host did not provide DAR info during backup	The backup was performed with an NDMP host version that does not support DAR. Ask the NDMP host vendor if a later NAS software version is available that supports DAR.

Table 24-2 DAR log messages (*continued*)

Message	Explanation
DAR disabled—Exceeded optimal DAR parameters for this image size	NetBackup determined that the restore would take longer with DAR than without it.
DAR disabled—Directory DAR not supported	DAR is automatically disabled when a restore job specifies a directory to restore. DAR can be used to restore files, but not to restoring directories.
DAR disabled by host parameters	DAR was disabled on the Primary or Media Server Properties dialog box. See “About enabling or disabling DAR” on page 167.

About robot tests

Depending on the type of robot, use the tests in the following table to exercise the robot.

Table 24-3 Robot types and tests

Robot type	Test
TLD	tldtest
ACS	acstest

TLD robot test example for UNIX

To exercise drive 1 in the TLD robot `c2t3l0` the NDMP host `stripes` controls, use the following commands on UNIX:

```
/usr/opensv/volmgr/bin/tldtest -r stripes:c2t3l0 -d1 stripes:/dev/RMT/Ocbl
```

At the prompt, enter `?` for help information.

`inquiry` (Displays the Vendor ID and Product ID. If you get a UNIT ATTENTION message, try the `mode` command and then continue your testing.)

`s s` (Checks slot status.)

`s d` (Checks drive status.)

`m s3 d1` (Moves a tape from slot 3 to drive 1.)

`unload d1` (Unloads the tape.)

m d1 s3 (Moves the tape back to slot 3.)

Using NetBackup for NDMP scripts

This chapter includes the following topics:

- [About the NetBackup for NDMP scripts](#)
- [ndmp_start_notify script \(UNIX\)](#)
- [ndmp_start_notify.cmd script \(Microsoft Windows\)](#)
- [ndmp_end_notify script \(UNIX\)](#)
- [ndmp_end_notify.cmd script \(Microsoft Windows\)](#)
- [ndmp_start_path_notify script \(UNIX\)](#)
- [ndmp_start_path_notify.cmd script \(Microsoft Windows\)](#)
- [ndmp_end_path_notify script \(UNIX\)](#)
- [ndmp_end_path_notify.cmd script \(Microsoft Windows\)](#)
- [ndmp_moving_path_notify script \(UNIX\)](#)
- [ndmp_moving_path_notify.cmd script \(Microsoft Windows\)](#)

About the NetBackup for NDMP scripts

This topic provides information that you can use to customize the NDMP-specific notification scripts.

NetBackup for NDMP provides the following scripts (commands on Windows) for collecting information and providing notification of events.

Table 25-1 Scripts to run on the NetBackup for NDMP server

Scripts for UNIX	Scripts for Windows
ndmp_start_notify	ndmp_start_notify.cmd
ndmp_end_notify	ndmp_end_notify.cmd
ndmp_start_path_notify	ndmp_start_path_notify.cmd
ndmp_end_path_notify	ndmp_end_path_notify.cmd
ndmp_moving_path_notify	ndmp_moving_path_notify.cmd

The scripts are similar to those already included in your NetBackup server installation. To create the scripts on UNIX, copy the `bpstart_notify` and `bpend_notify` scripts from

```
/usr/opensv/netbackup/bin/goodies (UNIX)
```

to

```
/usr/opensv/netbackup/bin
```

on the NetBackup for NDMP server. Then rename the copied scripts and modify as needed.

On Windows, you must create the scripts from scratch.

ndmp_start_notify script (UNIX)

The UNIX scripts are provided as examples only. You must customize the scripts before using them. For example, the `-ne` value in the first `if` statement must be modified to reflect the number of passed parameters. For the `ndmp_start_notify` script, the `-ne` value must be set to 7.

On the UNIX media server, NetBackup calls the `ndmp_start_notify` script each time the client starts a backup operation. To use this script, create a script similar to

```
/usr/opensv/netbackup/bin/goodies/bpstart_notify
```

on the server, and copy it to

```
/usr/opensv/netbackup/bin/ndmp_start_notify
```

on the UNIX NetBackup for NDMP server. Then, modify the script and ensure that you have permission to run it.

Note: Before you use this script, make sure that you can run it by using `other` on the media server. Run `chmod 755 script_name`, where *script_name* is the name of the script.

The `ndmp_start_notify` script runs each time a backup starts and after the tape has been positioned. This script must exit with a status of 0 for the calling program to continue and for the backup to proceed. A nonzero status causes the client backup to exit with a status of `ndmp_start_notify failed`.

If the `/usr/opensv/netbackup/bin/ndmp_start_notify` script exists, it runs in the foreground. The `bptm` process that is on the NetBackup for NDMP server waits for it to complete before continuing. Any commands in the script that do not end with an `&` character run serially.

The server expects the client to respond with a `continue` message within the period of time that the NetBackup `CLIENT_READ_TIMEOUT` option on the server specifies.

The default for `CLIENT_READ_TIMEOUT` is 300. If the script needs more time than 300 seconds, increase the value to allow more time.

NetBackup passes the following parameters to the script:

Table 25-2 Script parameters for `ndmp_start_notify` (UNIX)

Parameter	Description
\$1	Specifies the name of the NDMP host.
\$2	Specifies the policy name from the NetBackup catalog.
\$3	Specifies the schedule name from the NetBackup catalog.
\$4	Specifies one of the following: FULL INCR (differential incremental) CINC (cumulative incremental)
\$5	Specifies the NetBackup status code for the operation.

For example:

```
ndmp_start_notify freddie cd4000s fulls FULL 0
ndmp_start_notify danr cd4000s incrementals INCR 0
ndmp_start_notify hare cd4000s fulls FULL 0
```

To create an `ndmp_start_notify` script for a specific policy or policy and schedule combination, create script files with a `.polycname` or `.polycname.schedulename` suffix. In the following two examples of script names, the policy is named `production` and the schedule is named `fulls`:

```
/usr/opensv/netbackup/bin/ndmp_start_notify.production  
/usr/opensv/netbackup/bin/ndmp_start_notify.production.fulls
```

The first script affects all scheduled backups in the policy that is named `production`. The second script affects scheduled backups in the policy that is named `production` only when the schedule is named `fulls`.

Note: For a given backup, NetBackup uses only one `ndmp_start_notify` script and that is the one with the most specific name. For example, if there are both `ndmp_start_notify.production` and `ndmp_start_notify.production.fulls` scripts, NetBackup uses only `ndmp_start_notify.production.fulls`.

The `ndmp_start_notify` script can use the following environment variables:

```
BACKUPID  
UNIXBACKUPTIME  
BACKUPTIME
```

The NetBackup `bptm` process creates these variables. The following are examples of the strings that are available to the script for use in recording information about a backup:

```
BACKUPID=freddie_0857340526  
UNIXBACKUPTIME=0857340526  
BACKUPTIME=Sun Mar 2 16:08:46 1997
```

ndmp_start_notify.cmd script (Microsoft Windows)

When you use Windows NetBackup for NDMP media servers, you can create the batch scripts that provide notification whenever the client starts a backup. These scripts must reside on the media server in the following directory:

```
install_path\NetBackup\bin
```

where *install_path* is the directory where NetBackup is installed.

You can create `ndmp_start_notify` scripts that provide notification for all backups or only for backups of a specific policy or schedule. The `ndmp_start_notify` script runs each time a backup starts and after the tape is positioned.

To create a script that applies to all backups, name the script:

```
install_path\netbackup\bin\ndmp_start_notify.cmd
```

To create an `ndmp_start_notify` script that applies only to a specific policy or policy and schedule combination, add a `.polycname` or `.polycname.schedulename` suffix to the script name. The following are two examples:

- The following script applies only to a policy named `days`:

```
install_path\netbackup\bin\ndmp_start_notify.days.cmd
```

- The following script applies only to a schedule that is named `fulls`, which is in a policy named `days`:

```
install_path\netbackup\bin\ndmp_start_notify.days.fulls.cmd
```

The first script affects the scheduled backups in the policy named `days`. The second script affects the scheduled backups in the policy named `days` only when the schedule is named `fulls`.

For a given backup, NetBackup calls only one `ndmp_start_notify` script and checks for them in the following order:

```
ndmp_start_notify.policy.schedule.cmd
ndmp_start_notify.policy.cmd
ndmp_start_notify.cmd
```

For example, if there are both `ndmp_start_notify.policy.cmd` and `ndmp_start_notify.policy.schedule.cmd` scripts, NetBackup uses only the `ndmp_start_notify.policy.schedule.cmd` script.

Note: If you also use `ndmp_end_notify` scripts, they can provide a different level of notification than the `ndmp_start_notify` scripts. For example, if you had one of each, they could be `ndmp_start_notify.policy.cmd` and `ndmp_end_notify.policy.schedule.cmd`.

When the backup starts, NetBackup passes the following parameters to the script:

Table 25-3 Script parameters for `ndmp_start_notify.cmd` (Microsoft Windows)

Parameter	Description
%1	Specifies the name of the client from the NetBackup catalog.

Table 25-3 Script parameters for ndmp_start_notify.cmd (Microsoft Windows)
(continued)

Parameter	Description
%2	Specifies the policy name from the NetBackup catalog.
%3	Specifies the schedule name from the NetBackup catalog.
%4	Specifies one of the following: FULL INCR CINC
%5	Specifies the status of the operation is always 0 for bpstart_notify.
%6	Specifies the results file that NetBackup checks for a return code from the script. NetBackup uses %6 to pass the file name and then expects the script to create the file in the same directory as the script. If the script applies to a specific policy and schedule, the results file must be named <pre>install_path\netbackup\bin\NDMP_START_NOTIFY_RES.policy.schedule</pre> If the script applies to a specific policy, the results file must be named <pre>install_path\NetBackup\bin\NDMP_START_NOTIFY_RES.policy</pre> If the script applies to all backups, the results file must be named <pre>install_path\NetBackup\bin\NDMP_START_NOTIFY_RES</pre> An echo 0> %6 statement is one way for the script to create the file. NetBackup deletes the existing results file before it calls the script. After the script runs, NetBackup checks the new results file for the status. The status must be 0 for the script to be considered successful. If the results file does not exist, NetBackup assumes that the script was successful.

The server expects the client to respond with a `continue` message within the period of time that the NetBackup `CLIENT_READ_TIMEOUT` option on the server specifies. The default is 300 seconds. If the script needs more than 300 seconds, increase the value to allow more time.

ndmp_end_notify script (UNIX)

The `ndmp_end_notify` script is run at the end of the backup. The backup does not wait for the script to complete.

Note: Before you use this script, make sure you can run it by using `other` on the media server. Run `chmod 755 script_name`, where *script_name* is the name of the script.

The UNIX scripts are provided as examples only. You must customize the scripts before using them. For example, the `-ne` value in the first if statement must be modified to reflect the number of passed parameters. For the `ndmp_end_notify` script, the `-ne` value must be set to 7.

For a UNIX media server, if you need notification whenever the NDMP host completes a backup, copy

```
/usr/opensv/netbackup/bin/goodies/bpend_notify
```

from the server, to

```
/usr/opensv/netbackup/bin/ndmp_end_notify
```

on the UNIX NetBackup for NDMP host. Then, modify the script and ensure that you have permission to run it.

The `ndmp_end_notify` script runs each time a backup completes.

NetBackup passes the following parameters to the `ndmp_end_notify` script:

Table 25-4 Script parameters for `ndmp_end_notify` (UNIX)

Parameter	Description
\$1	Specifies the name of the client from the NetBackup catalog.
\$2	Specifies the policy name from the NetBackup catalog.
\$3	Specifies the schedule name from the NetBackup catalog.
\$4	Specifies one of the following: FULL INCR (differential incremental) CINC (cumulative incremental)
\$5	Specifies the exit code from <code>bptm</code> .

For example:

```
ndmp_end_notify freddie cd4000s fulls FULL 0  
ndmp_end_notify danr cd4000s incrementals INCR 73
```

To create an `ndmp_end_notify` script for a specific policy or policy and schedule combination, create script files with a `.policyname` or `.policyname.schedulename` suffix. In the following two examples of script names, the policy is named `production` and the schedule is named `fulls`:

```
/usr/opensv/netbackup/bin/ndmp_end_notify.production  
/usr/opensv/netbackup/bin/ndmp_end_notify.production.fulls
```

The first script affects all scheduled backups in the policy that is named `production`. The second script affects scheduled backups in the policy that is named `production` only when the schedule is named `fulls`.

Note: For a given backup, NetBackup uses only one `ndmp_end_notify` script and that is the one with the most specific name. For example, if there are both `ndmp_end_notify.production` and `ndmp_end_notify.production.fulls` scripts, NetBackup uses only `ndmp_end_notify.production.fulls`.

The `ndmp_end_notify` script can use the following environment variables:

```
BACKUPID  
UNIXBACKUPTIME  
BACKUPTIME
```

The NetBackup `bptm` process creates these variables. The following are examples of the strings that are available to the script for use in recording information about a backup:

```
BACKUPID=freddie_0857340526  
UNIXBACKUPTIME=0857340526  
BACKUPTIME=Sun Mar 2 16:08:46 1997
```

ndmp_end_notify.cmd script (Microsoft Windows)

For Windows media servers, you can create the batch scripts that provide notification whenever the client completes a backup. These scripts must reside on the media server in the same directory as the NetBackup binaries:

```
install_path\NetBackup\bin
```

where *install_path* is the directory where NetBackup is installed.

You can create `ndmp_end_notify` scripts that provide notification for all backups or only for backups of a specific policy or schedule.

To create an `ndmp_end_notify` script that applies to all backups, name the script:

```
install_path\netbackup\bin\ndmp_end_notify.cmd
```

To create a script that applies only to a specific policy or policy and schedule combination, add a `.policyname` or `.policyname.schedulename` suffix to the script name. The following are two examples:

- The following script applies only to a policy named `days`:

```
install_path\netbackup\bin\ndmp_end_notify.days.cmd
```

- The following script applies only to a schedule that is named `fulls`, which is in a policy named `days`:

```
install_path\netbackup\bin\ndmp_end_notify.days.fulls.cmd
```

The first script affects all scheduled backups in the policy named `days`. The second script affects scheduled backups in the policy named `days` only when the schedule is named `fulls`.

For a given backup, NetBackup calls only one `ndmp_end_notify` script and checks for them in the following order:

```
ndmp_end_notify.policy.schedule.cmd  
ndmp_end_notify.policy.cmd  
ndmp_end_notify.cmd
```

For example, if there are both `ndmp_end_notify.policy.cmd` and `ndmp_end_notify.policy.schedule.cmd` scripts, NetBackup uses only `ndmp_end_notify.policy.schedule.cmd`.

Note: If you also use `ndmp_start_notify` scripts, they can provide a different level of notification than the `ndmp_end_notify` scripts. For example, if you had one of each, they could be `ndmp_start_notify.policy.cmd` and `ndmp_end_notify.policy.schedule.cmd`.

When the backup completes, NetBackup passes the following parameters to the script:

Table 25-5 Script parameters for `ndmp_end_notify.cmd` (Microsoft Windows)

Parameter	Description
%1	Specifies the name of the client from the NetBackup catalog.
%2	Specifies the policy name from the NetBackup catalog.

Table 25-5 Script parameters for ndmp_end_notify.cmd (Microsoft Windows)
(continued)

Parameter	Description
%3	Specifies the schedule name from the NetBackup catalog.
%4	Specifies one of the following: FULL INCR CINC
%5	Specifies the status of the operation. It is the same as the status sent to the NetBackup server. This status is 0 for successful backups and 1 for partially successful backups. If an error occurs, the status is the value associated with that error.
%6	Note: The following file is not checked at the end of a backup. Specifies the results file that NetBackup checks for a return code from the script. NetBackup uses %6 to pass the file name and then expects the script to create the file in the same directory as the script. If the script applies to a specific policy and schedule, the results file must be named <i>install_path\NetBackup\bin\NDMP_END_NOTIFY_RES.policy.schedule</i> If the script applies to a specific policy, the results file must be named <i>install_path\netbackup\bin\NDMP_END_NOTIFY_RES.policy</i> If the script applies to all backups, the results file must be named <i>install_path\NetBackup\bin\NDMP_END_NOTIFY_RES</i> An <code>echo 0> %6</code> statement is one way for the script to create the file. NetBackup deletes the existing results file before it calls the script. After the script runs, NetBackup checks the new results file for the status. The status must be 0 for the script to be considered successful. If the results file does not exist, NetBackup assumes that the script was successful.

ndmp_start_path_notify script (UNIX)

The UNIX scripts are provided as examples only. You must customize the scripts before using them. For example, the `-ne` value in the first `if` statement must be modified to reflect the number of passed parameters. For the `ndmp_start_path_notify` script, the `-ne` value must be set to 7.

To use this script, create a script similar to

```
/usr/opensv/netbackup/bin/goodies/bpstart_notify
```

on the server, and copy it to

```
/usr/opensv/netbackup/bin/ndmp_start_path_notify
```

on the UNIX NetBackup for NDMP server. Then, modify the script and ensure that you have permission to run it.

On the UNIX media server, the `ndmp_start_path_notify` script runs before the backup process is issued to the NAS machine. This script must exit with a status of 0 for the calling program to continue and for the backup to proceed. A nonzero status causes the client backup to exit with a status of 99 (NDMP backup failure).

Note: Before you use this script, make sure you can run it by using `other` on the media server. Run `chmod 755 script_name`, where *script_name* is the name of the script.

If the `/usr/opensv/netbackup/bin/ndmp_start_path_notify` script exists, it runs in the foreground. The `bptm` process on the NetBackup for NDMP server waits for it to complete before continuing. Any commands in the script that do not end with an `&` character run serially.

The server expects the client to respond with a `continue` message within the period of time that the NetBackup `CLIENT_READ_TIMEOUT` option on the server specifies.

The default for `CLIENT_READ_TIMEOUT` is 300. If the script needs more time than 300 seconds, increase the value to allow more time.

NetBackup passes the following parameters to the script:

Table 25-6 Script parameters for `ndmp_start_path_notify` (UNIX)

Parameter	Description
\$1	Specifies the name of the NDMP host.
\$2	Specifies the policy name from the NetBackup catalog.
\$3	Specifies the schedule name from the NetBackup catalog.
\$4	Specifies one of the following: FULL INCR (differential incremental) CINC (cumulative incremental)
\$5	Specifies the NetBackup status code for the operation.

Table 25-6 Script parameters for ndmp_start_path_notify (UNIX) (*continued*)

Parameter	Description
\$6	Not used.
\$7	Specifies the path being backed up.

For example:

```
ndmp_start_path_notify freddie cd4000s fulls FULL
ndmp_start_path_notify danr cd4000s incrementals INCR
ndmp_start_path_notify hare cd4000s fulls FULL
```

To create an `ndmp_start_path_notify` script for a specific policy or policy and schedule combination, create script files with a `.polycname` or `.polycname.schedulename` suffix. In the following two examples of script names, the policy is named `production` and the schedule is named `fulls`:

```
/usr/opensv/netbackup/bin/ndmp_start_path_notify.production
/usr/opensv/netbackup/bin/ndmp_start_path_notify.production.fulls
```

The first script affects all scheduled backups in the policy that is named `production`. The second script affects scheduled backups in the policy that is named `production` only when the schedule is named `fulls`.

Note: For a given backup, NetBackup uses only one `ndmp_start_path_notify` script and that is the one with the most specific name. For example, if there are both `ndmp_start_path_notify.production` and `ndmp_start_path_notify.production.fulls` scripts, NetBackup uses only `ndmp_start_path_notify.production.fulls`.

The `ndmp_start_path_notify` script can use the following environment variables:

```
BACKUPID
UNIXBACKUPTIME
BACKUPTIME
```

The NetBackup `bptm` process creates these variables. The following are examples of the strings that are available to the script for use in recording information about a backup:

```
BACKUPID=freddie_0857340526
UNIXBACKUPTIME=0857340526
BACKUPTIME=Sun Mar 2 16:08:46 1997
```

ndmp_start_path_notify.cmd script (Microsoft Windows)

For Windows media servers, you can create the batch scripts that provide notification before the backup process is issued to the NAS machine. These scripts must reside on the media server in the same directory as the NetBackup binaries:

```
install_path\NetBackup\bin
```

where *install_path* is the directory where NetBackup is installed.

You can create `ndmp_start_path_notify` scripts that provide notification for all backups or only for backups of a specific policy or schedule.

To create an `ndmp_start_path_notify` script that applies to all backups, name the script:

```
install_path\netbackup\bin\ndmp_start_path_notify.cmd
```

To create a script that applies only to a specific policy or policy and schedule combination, add a `.polycname` or `.polycname.schedulename` suffix to the script name. The following are two examples:

- The following script applies only to a policy named `days`:

```
install_path\netbackup\bin\ndmp_start_path_notify.days.cmd
```

- The following script applies only to a schedule that is named `fulls`, which in a policy named `days`:

```
install_path\netbackup\bin\ndmp_start_path_notify.days.fulls.cmd
```

The first script affects all scheduled backups in the policy named `days`. The second script affects scheduled backups in the policy named `days` only when the schedule is named `fulls`.

For a given backup, NetBackup calls only one `ndmp_start_path_notify` script and checks for them in the following order:

```
ndmp_start_path_notify.policy.schedule.cmd  
ndmp_start_path_notify.policy.cmd  
ndmp_start_path_notify.cmd
```

For example, if there are both `ndmp_start_path_notify.policy.cmd` and `ndmp_start_path_notify.policy.schedule.cmd` scripts, NetBackup uses only `ndmp_start_path_notify.policy.schedule.cmd`.

Note: If you also use `ndmp_start_notify` scripts, they can provide a different level of notification than the `ndmp_start_path_notify` scripts. For example, if you had one of each, they could be `ndmp_start_notify.policy.cmd` and `ndmp_start_path_notify.policy.schedule.cmd`.

When the backup starts, NetBackup passes the following parameters to the script:

Table 25-7 Script parameters for `ndmp_start_path_notify.cmd` (Microsoft Windows)

Parameter	Description
%1	Specifies the name of the client from the NetBackup catalog.
%2	Specifies the policy name from the NetBackup catalog.
%3	Specifies the schedule name from the NetBackup catalog.
%4	Specifies one of the following: FULL INCR CINC
%5	Specifies the status of the operation. It is the same as the status sent to the NetBackup server. This status is 0 for successful backups and 1 for partially successful backups. If an error occurs, the status is the value associated with that error.
%6	Specifies the results file that NetBackup checks for a return code from the script. NetBackup uses %6 to pass the file name and then expects the script to create the file in the same directory as the script. If the script applies to a specific policy and schedule, the results file must be named <code>install_path\netbackup\bin\NDMP_START_PATH_NOTIFY_RES.policy.schedule</code> If the script applies to a specific policy, the results file must be named <code>install_path\NetBackup\bin\NDMP_START_PATH_NOTIFY_RES.policy</code> If the script applies to all backups, the results file must be named <code>install_path\NetBackup\bin\NDMP_START_PATH_NOTIFY_RES</code> An <code>echo 0> %6</code> statement is one way for the script to create the file. NetBackup deletes the existing results file before it calls the script. After the script runs, NetBackup checks the new results file for the status. The status must be 0 for the script to be considered successful. If the results file does not exist, NetBackup assumes that the script was successful.
%7	Pathname being backed up.

ndmp_end_path_notify script (UNIX)

The UNIX scripts are provided as examples only. You must customize the scripts before using them. For example, the `-ne` value in the first `if` statement must be modified to reflect the number of passed parameters. For the `ndmp_end_path_notify` script, the `-ne` value must be set to 7.

Note: Before you use this script, make sure you can run it by using `other` on the media server. Run `chmod 755 script_name`, where *script_name* is the name of the script.

For a UNIX media server, if you need notification whenever the NDMP host completes a backup, copy

```
/usr/opensv/netbackup/bin/goodies/bpend_notify
```

from the server, to

```
/usr/opensv/netbackup/bin/ndmp_end_path_notify
```

on the UNIX NetBackup for NDMP host. Then, modify the script and ensure that you have permission to run it.

The `ndmp_end_path_notify` script runs after the NAS machine has informed NetBackup that it has completed sending data.

NetBackup passes the following parameters to the `ndmp_end_notify` script:

Table 25-8 Script parameters for `ndmp_end_path_notify` (UNIX)

Parameter	Description
\$1	Specifies the name of the client from the NetBackup catalog.
\$2	Specifies the policy name from the NetBackup catalog.
\$3	Specifies the schedule name from the NetBackup catalog.
\$4	Specifies one of the following: FULL INCR (differential incremental) CINC (cumulative incremental)
\$5	Specifies the exit code from <code>bptm</code> .
\$6	Not used.

Table 25-8 Script parameters for ndmp_end_path_notify (UNIX) (*continued*)

Parameter	Description
\$7	Specifies the path being backed up.

For example:

```
ndmp_end_path_notify freddie cd4000s fulls FULL 0
ndmp_end_path_notify danr cd4000s incrementals INCR 73
```

To create an `ndmp_end_path_notify` script for a specific policy or policy and schedule combination, create script files with a `.polycname` or `.polycname.schedulename` suffix. In the following two examples of script names, the policy is named `production` and the schedule is named `fulls`:

```
/usr/opensv/netbackup/bin/ndmp_end_path_notify.production
/usr/opensv/netbackup/bin/ndmp_end_path_notify.production.fulls
```

The first script affects all scheduled backups in the policy that is named `production`. The second script affects scheduled backups in the policy that is named `production` only when the schedule is named `fulls`.

Note: For a given backup, NetBackup uses only one `ndmp_end_path_notify` script and that is the one with the most specific name. For example, if there are both `ndmp_end_path_notify.production` and `ndmp_end_path_notify.production.fulls` scripts, NetBackup uses only `ndmp_end_path_notify.production.fulls`.

The `ndmp_end_path_notify` script can use the following environment variables:

```
BACKUPID
UNIXBACKUPTIME
BACKUPTIME
```

The NetBackup `bptm` process creates these variables. The following are examples of the strings that are available to the script for use in recording information about a backup:

```
BACKUPID=freddie_0857340526
UNIXBACKUPTIME=0857340526
BACKUPTIME=Sun Mar 2 16:08:46 1997
```

ndmp_end_path_notify.cmd script (Microsoft Windows)

For Windows media servers, you can create the batch scripts that provide notification whenever the client is finished writing to tape. These scripts must reside on the media server in the same directory as the NetBackup binaries:

```
install_path\NetBackup\bin
```

where *install_path* is the directory where NetBackup is installed.

You can create `ndmp_end_path_notify` scripts that provide notification for all backups or only for backups of a specific policy or schedule.

To create an `ndmp_end_path_notify` script that applies to all backups, name the script:

```
install_path\netbackup\bin\ndmp_end_path_notify.cmd
```

To create a script that applies only to a specific policy or policy and schedule combination, add a `.polycname` or `.polycname.schedulename` suffix to the script name. The following are two examples:

- The following script applies only to a policy named `days`:

```
install_path\netbackup\bin\ndmp_end_path_notify.days.cmd
```

- The following script applies only to a schedule that is named `fulls`, which is in a policy named `days`:

```
install_path\netbackup\bin\ndmp_end_path_notify.days.fulls.  
cmd
```

The first script affects all scheduled backups in the policy named `days`. The second script affects scheduled backups in the policy named `days` only when the schedule is named `fulls`.

For a given backup, NetBackup calls only one `ndmp_end_path_notify` script and checks for them in the following order:

```
ndmp_end_path_notify.policy.schedule.cmd  
ndmp_end_path_notify.policy.cmd  
ndmp_end_path_notify.cmd
```

For example, if there are both `ndmp_end_path_notify.policy.cmd` and `ndmp_end_path_notify.policy.schedule.cmd` scripts, NetBackup uses only `ndmp_end_path_notify.policy.schedule.cmd`.

Note: If you also use `ndmp_end_notify` scripts, they can provide a different level of notification than the `ndmp_end_path_notify` scripts. For example, if you had one of each, they could be `ndmp_end_notify.policy.cmd` and `ndmp_end_path_notify.policy.schedule.cmd`.

When the backup completes, NetBackup passes the following parameters to the script:

Table 25-9 Script parameters for `ndmp_end_path_notify.cmd` (Microsoft Windows)

Parameter	Description
%1	Specifies the name of the client from the NetBackup catalog.
%2	Specifies the policy name from the NetBackup catalog.
%3	Specifies the schedule name from the NetBackup catalog.
%4	Specifies one of the following: FULL INCR CINC
%5	Specifies the status of the operation. It is the same as the status sent to the NetBackup server. This status is 0 for successful backups and 1 for partially successful backups. If an error occurs, the status is the value associated with that error.

Table 25-9 Script parameters for ndmp_end_path_notify.cmd (Microsoft Windows) (*continued*)

Parameter	Description
%6	Note: The following file is not checked when using <code>ndmp_end_path_notify</code>. Specifies the results file that NetBackup checks for a return code from the script. NetBackup uses %6 to pass the file name and then expects the script to create the file in the same directory as the script. If the script applies to a specific policy and schedule, the results file must be named <i>install_path\NetBackup\bin\NDMP_END_PATH_NOTIFY_RES.policy.schedule</i> If the script applies to a specific policy, the results file must be named <i>install_path\netbackup\bin\NDMP_END_PATH_NOTIFY_RES.policy</i> If the script applies to all backups, the results file must be named <i>install_path\NetBackup\bin\NDMP_END_PATH_NOTIFY_RES</i> An <code>echo 0 > %6</code> statement is one way for the script to create the file. NetBackup deletes the existing results file before it calls the script. After the script runs, NetBackup checks the new results file for the status. The status must be 0 for the script to be considered successful. If the results file does not exist, NetBackup assumes that the script was successful.
%7	Specifies the pathname being backed up.

ndmp_moving_path_notify script (UNIX)

The UNIX scripts are provided as examples only. You must customize the scripts before using them. For example, the `-ne` value in the first if statement must be modified to reflect the number of passed parameters. For the `ndmp_moving_path_notify` script, the `-ne` value must be set to 7.

To use this script, create a script similar to

```
/usr/opensv/netbackup/bin/goodies/bpstart_notify
```

on the server, and copy it to

```
/usr/opensv/netbackup/bin/ndmp_moving_path_notify
```

on the UNIX NetBackup for NDMP server. Then, modify the script and ensure that you have permission to run it.

On UNIX media servers, the `ndmp_moving_path_notify` script runs after the backup process sends data to NetBackup.

Note: Before you use this script, make sure you can run it using other on the media server. Run `chmod 755 script_name`, where *script_name* is the name of the script.

If the `/usr/opensv/netbackup/bin/ndmp_moving_path_notify` script exists, it runs in the foreground. The `bptm` process that is on the NetBackup for NDMP server waits for it to complete before continuing. Any commands in the script that do not end with an `&` character run serially.

The server expects the client to respond with a `continue` message within the period of time that the NetBackup `CLIENT_READ_TIMEOUT` option on the server specifies.

The default for `CLIENT_READ_TIMEOUT` is 300 seconds. If the script needs more than 300 seconds, increase the value to allow more time.

NetBackup passes the following parameters to the script:

Table 25-10 Script parameters for `ndmp_moving_path_notify` (UNIX)

Parameter	Description
\$1	Specifies the name of the NDMP host.
\$2	Specifies the policy name from the NetBackup catalog.
\$3	Specifies the schedule name from the NetBackup catalog.
\$4	Specifies one of the following: FULL INCR (differential incremental) CINC (cumulative incremental)
\$5	Specifies the NetBackup status code for the operation.
\$6	Not used.
\$7	Specifies the path being backed up.

For example:

```
ndmp_moving_path_notify freddie cd4000s fulls FULL
ndmp_moving_path_notify danr cd4000s incrementals INCR
ndmp_moving_path_notify hare cd4000s fulls FULL
```

To create an `ndmp_moving_path_notify` script for a specific policy or policy and schedule combination, create script files with a `.polycname` or `.polycname.schedulename` suffix. In the following two examples of script names, the policy is named `production` and the schedule is named `fulls`:

```
/usr/opensv/netbackup/bin/ndmp_moving_path_notify.production
/usr/opensv/netbackup/bin/ndmp_moving_path_notify.production.fulls
```

The first script affects all scheduled backups in the policy that is named `production`. The second script affects scheduled backups in the policy that is named `production` only when the schedule is named `fulls`.

Note: For a given backup, NetBackup uses only one `ndmp_moving_path_notify` script and that is the one with the most specific name. For example, if there are both `ndmp_moving_path_notify.production` and `ndmp_moving_path_notify.production.fulls` scripts, NetBackup uses only `ndmp_moving_path_notify.production.fulls`.

The `ndmp_moving_path_notify` script can use the following environment variables:

```
BACKUPID
UNIXBACKUPTIME
BACKUPTIME
```

The NetBackup `bptm` process creates these variables. The following are examples of the strings that are available to the script for use in recording information about a backup:

```
BACKUPID=freddie_0857340526
UNIXBACKUPTIME=0857340526
BACKUPTIME=Sun Mar 2 16:08:46 1997
```

ndmp_moving_path_notify.cmd script (Microsoft Windows)

For Windows media servers, you can create the batch scripts that provide notification whenever the NAS machine starts sending data. These scripts must reside on the media server in the same directory as the NetBackup binaries:

```
install_path\NetBackup\bin
```

where *install_path* is the directory where NetBackup is installed.

You can create `ndmp_moving_path_notify` scripts that provide notification for all backups or only for backups of a specific policy or schedule.

To create an `ndmp_moving_path_notify` script that applies to all backups, name the script:

```
install_path\netbackup\bin\ndmp_moving_path_notify.cmd
```

To create a script that applies only to a specific policy or policy and schedule combination, add a `.policyname` or `.policyname.schedulename` suffix to the script name. The following are two examples:

- The following script applies only to a policy named `days`:

```
install_path\netbackup\bin\ndmp_moving_path_notify.days.cmd
```

- The following script applies only to a schedule that is named `fulls`, which is in a policy named `days`:

```
install_path\netbackup\bin\ndmp_moving_path_notify.days.fulls.cmd
```

The first script affects all scheduled backups in the policy named `days`. The second script affects scheduled backups in the policy named `days` only when the schedule is named `fulls`.

For a given backup, NetBackup calls only one `ndmp_moving_path_notify` script and checks for them in the following order:

```
ndmp_moving_path_notify.policy.schedule.cmd  
ndmp_moving_path_notify.policy.cmd  
ndmp_moving_path_notify.cmd
```

For example, if there are both `ndmp_moving_path_notify.policy.cmd` and `ndmp_moving_path_notify.policy.schedule.cmd` scripts, NetBackup uses only `ndmp_moving_path_notify.policy.schedule.cmd`.

Note: If you also use `ndmp_start_notify` scripts, they can provide a different level of notification than the `ndmp_moving_path_notify` scripts. For example, if you had one of each, they could be `ndmp_start_notify.policy.cmd` and `ndmp_moving_path_notify.policy.schedule.cmd`.

When the backup starts, NetBackup passes the following parameters to the script.

Table 25-11 Script parameters for ndmp_moving_path_notify.cmd (Microsoft Windows)

Parameter	Description
%1	Specifies the name of the client from the NetBackup catalog.
%2	Specifies the policy name from the NetBackup catalog.
%3	Specifies the schedule name from the NetBackup catalog.
%4	Specifies one of the following: FULL INCR CINC
%5	Specifies the status of the operation. It is the same as the status sent to the NetBackup server. This status is 0 for successful backups and 1 for partially successful backups. If an error occurs, the status is the value associated with that error.
%6	Note: The following is not checked when using <code>ndmp_moving_path_notify</code>. Specifies the results file that NetBackup checks for a return code from the script. NetBackup uses %6 to pass the file name and then expects the script to create the file in the same directory as the script. If the script applies to a specific policy and schedule, the results file must be named <code>install_path\netbackup\bin\NDMP_END_NOTIFY_RES.policy.schedule</code> If the script applies to a specific policy, the results file must be named <code>install_path\NetBackup\bin\NDMP_END_NOTIFY_RES.policy</code> If the script applies to all backups, the results file must be named <code>install_path\NetBackup\bin\NDMP_END_NOTIFY_RES</code> An <code>echo 0> %6</code> statement is one way for the script to create the file. NetBackup deletes the existing results file before it calls the script. After the script runs, NetBackup checks the new results file for the status. The status must be 0 for the script to be considered successful. If the results file does not exist, NetBackup assumes that the script was successful.
%7	Specifies the pathname being backed up.