

Dell EMC NetWorker

Version 19.1

Snapshot Management for NAS Devices Integration Guide

302-005-699

REV 01

Copyright © 1990-2019 Dell Inc. or its subsidiaries. All rights reserved.

Published May 2019

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Tables		7
Figures		9
Preface		11
Chapter 1	Overview of NAS Snapshot Features	17
	NetWorker Snapshot Management for NAS feature description.....	18
	NAS snapshot operations.....	18
	Types of NAS snapshot-based backups.....	19
	Types of NAS snapshot and clone recoveries.....	20
	NetWorker clone support.....	20
	Backup configuration methods.....	20
	Recovery interfaces.....	21
	Monitoring and reporting NAS snapshot operations.....	21
	Internationalization support.....	21
	Multiple rollover sessions for NAS.....	21
	Creating a multiple rollover session for NAS.....	22
	Components of the NAS snapshot network.....	22
	NAS devices.....	22
	NAS snapshot naming convention.....	24
	NDMP environment.....	25
	NetWorker server.....	25
	NetWorker client.....	25
	NetWorker storage node.....	25
	Backup storage media.....	25
	Licensing requirements.....	26
	Example NAS snapshot environments.....	26
	Example of a snapshot and rollover to storage media.....	26
	Example of a snapshot from replication.....	27
	Examples of snapshot recovery.....	28
Chapter 2	Data Protection Policies	31
	Introduction.....	32
	Default data protection policies in NMC's NetWorker Administration window	32
	Strategies for NAS snapshot backups.....	33
	Overview of configuring a new data protection policy.....	33
	Creating a policy.....	34
	Create a workflow for a new policy in NetWorker Administration.....	36
	Protection groups for NAS snapshot devices.....	38
	Creating a basic client group.....	39
	Creating a dynamic client group.....	40
	Creating a NAS device group.....	40
	Creating a save set group.....	41
	Creating a query group.....	42

	Supported actions in snapshot workflows.....	44
	Action in snapshot workflows.....	45
	Creating a probe action.....	46
	Creating a check connectivity action.....	50
	Creating a snapshot backup action.....	53
	Creating a discover action.....	57
	Creating a generate index action.....	58
	Creating a clone action.....	59
	Visual representation of NAS snapshot workflows.....	64
Chapter 3	Software Configuration	67
	Backup group resource migration.....	68
	Road map for NAS snapshot configuration.....	70
	NetWorker with NAS configuration prerequisites.....	70
	NAS device prerequisites.....	71
	Replication prerequisites.....	71
	NetApp prerequisites.....	71
	Storage node prerequisite.....	72
	Configuring NDMP.....	72
	Configuring a Client resource with the wizard.....	73
	Performing post Client Configuration wizard steps.....	76
	Configuring a NAS device client resource manually.....	77
	Configuring NSM with NetApp NAS snapshots.....	78
	Configuring NSM with NetApp ONTAP 7-mode	78
	Configuring NSM with NetApp ONTAP Cluster Mode	78
	Configuring the Application Information variables.....	79
	Configuring preprocessing and postprocessing scripts.....	79
Chapter 4	Data Management and Recovery	81
	Snapshot management.....	82
	Snapshot reporting.....	82
	Save set IDs and expiration policies.....	82
	Browsing snapshot and clone save sets.....	82
	Browsing and recovery on Windows hosts.....	83
	NAS snapshot recovery support and limitations.....	83
	Recovering data from a snapshot with the Recovery Wizard.....	84
Appendix A	Application Information Variables	89
	Using Application Information variables.....	90
	Common Application Information variables.....	90
Appendix B	Command-Line Operations for NAS Recovery	93
	Using CLI commands for NetWorker operations.....	94
	NAS snapshot commands.....	94
	Using nsrsnapadmin operations.....	94
	Example nsrsnapadmin operations.....	95
	Querying with the mminfo command.....	97
	Using nsrnassnap_recover operations.....	98
Appendix C	Snapshot limitations	99
	NAS snapshot support and limitation.....	100
	Device-specific limitations for snapshot recovery.....	100

Glossary

105

CONTENTS

TABLES

1	Revision history.....	11
2	Style conventions.....	13
3	NAS device support and limitations.....	23
4	Save set criteria.....	42
5	Schedule icons.....	47
6	Schedule icons.....	51
7	Backup type icons.....	54
8	Snapshot discovery icons.....	57
9	Snapshot generate icons.....	59
10	Schedule icons.....	60
11	Migration of Group attributes.....	68
12	Common Application Information variables.....	90
13	Commands and options that are supported in nrsnapadmin interactive mode.....	94
14	Device limitations and support for snapshot management.....	100
15	Device support for recovery of save sets and individual files.....	102
16	NetApp support for types of local replication and platforms.....	103
17	NetApp C-Mode local replication and platforms.....	103
18	NetApp support for types of remote replication and platforms.....	104

TABLES

FIGURES

1	NAS flow diagram.....	18
2	Remote replication with snapshot and clone to storage.....	27
3	In-place recovery restores data to the original location on the source device.....	28
4	Out-of-place recovery restores data to a different location on the source device.....	29
5	Directed recovery restores data to the NetWorker client that runs the recovery.....	29
6	Platinum policy configuration.....	32
7	Data protection policy example.....	34
8	All possible workflow actions for a snapshot backup.....	45
9	Workflow path from a snapshot backup action.....	46
10	Workflow path from a discover action.....	46
11	Sample NAS snapshot workflow.....	64

FIGURES

Preface

As part of an effort to improve product lines, periodic revisions of software and hardware are released. Therefore, all versions of the software or hardware currently in use might not support some functions that are described in this document. The product release notes provide the most up-to-date information on product features.

If a product does not function correctly or does not function as described in this document, contact a technical support professional.

Note

This document was accurate at publication time. To ensure that you are using the latest version of this document, go to the Support website <https://www.dell.com/support>.

Purpose

This document provides planning, practices, and configuration information for the use of the NetWorker Snapshot Management (NSM) for network-attached storage (NAS) devices features within an NetWorker backup and storage management environment.

Audience

This document is intended for system administrators. Readers of this document must be familiar with the following tasks:

- Identifying the different hardware and software components that make up the NetWorker datazone.
- Following procedures to configure storage management operations.
- Following guidelines to locate problems and implement solutions.

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
01	May 20, 2019	First release of the document for NetWorker 19.1.

Related documentation

The NetWorker documentation set includes the following publications, available on the Support website:

- *NetWorker E-LAB Navigator*
Provides compatibility information, including specific software and hardware configurations that NetWorker supports. To access E-LAB Navigator, go to <https://elabnavigator.emc.com/eln/elhome>.
- *NetWorker Administration Guide*
Describes how to configure and maintain the NetWorker software.
- *NetWorker Network Data Management Protocol (NDMP) User Guide*
Describes how to use the NetWorker software to provide data protection for NDMP filers.

- *NetWorker Cluster Integration Guide*
Contains information related to configuring NetWorker software on cluster servers and clients.
- *NetWorker Installation Guide*
Provides information on how to install, uninstall, and update the NetWorker software for clients, storage nodes, and servers on all supported operating systems.
- *NetWorker Updating from a Previous Release Guide*
Describes how to update the NetWorker software from a previously installed release.
- *NetWorker Release Notes*
Contains information on new features and changes, fixed problems, known limitations, environment and system requirements for the latest NetWorker software release.
- *NetWorker Command Reference Guide*
Provides reference information for NetWorker commands and options.
- *NetWorker Data Domain Boost Integration Guide*
Provides planning and configuration information on the use of Data Domain devices for data deduplication backup and storage in a NetWorker environment.
- *NetWorker Performance Optimization Planning Guide*
Contains basic performance tuning information for NetWorker.
- *NetWorker Server Disaster Recovery and Availability Best Practices Guide*
Describes how to design, plan for, and perform a step-by-step NetWorker disaster recovery.
- *NetWorker Snapshot Management Integration Guide*
Describes the ability to catalog and manage snapshot copies of production data that are created by using mirror technologies on storage arrays.
- *NetWorker Snapshot Management for NAS Devices Integration Guide*
Describes how to catalog and manage snapshot copies of production data that are created by using replication technologies on NAS devices.
- *NetWorker Security Configuration Guide*
Provides an overview of security configuration settings available in NetWorker, secure deployment, and physical security controls needed to ensure the secure operation of the product.
- *NetWorker VMware Integration Guide*
Provides planning and configuration information on the use of VMware in a NetWorker environment.
- *NetWorker Error Message Guide*
Provides information on common NetWorker error messages.
- *NetWorker Licensing Guide*
Provides information about licensing NetWorker products and features.
- *NetWorker REST API Getting Started Guide*
Describes how to configure and use the NetWorker REST API to create programmatic interfaces to the NetWorker server.
- *NetWorker REST API Reference Guide*
Provides the NetWorker REST API specification used to create programmatic interfaces to the NetWorker server.
- *NetWorker 19.1 with CloudBoost 19.1 Integration Guide*
Describes the integration of NetWorker with CloudBoost.

- *NetWorker 19.1 with CloudBoost 19.1 Security Configuration Guide*
Provides an overview of security configuration settings available in NetWorker and Cloud Boost, secure deployment, and physical security controls needed to ensure the secure operation of the product.
- NetWorker Management Console Online Help
Describes the day-to-day administration tasks performed in the NetWorker Management Console and the NetWorker Administration window. To view the online help, click **Help** in the main menu.
- NetWorker User Online Help
Describes how to use the NetWorker User program, which is the Windows client interface, to connect to a NetWorker server to back up, recover, archive, and retrieve files over a network.

Special notice conventions that are used in this document

The following conventions are used for special notices:

NOTICE

Identifies content that warns of potential business or data loss.

Note

Contains information that is incidental, but not essential, to the topic.

Typographical conventions

The following type style conventions are used in this document:

Table 2 Style conventions

Bold	Used for interface elements that a user specifically selects or clicks, for example, names of buttons, fields, tab names, and menu paths. Also used for the name of a dialog box, page, pane, screen area with title, table label, and window.
<i>Italic</i>	Used for full titles of publications that are referenced in text.
Monospace	Used for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, file names, file name extensions, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Used for variables.
Monospace bold	Used for user input.
[]	Square brackets enclose optional values.
	Vertical line indicates alternate selections. The vertical line means or for the alternate selections.
{ }	Braces enclose content that the user must specify, such as x, y, or z.
...	Ellipses indicate non-essential information that is omitted from the example.

You can use the following resources to find more information about this product, obtain support, and provide feedback.

Where to find product documentation

- <https://www.dell.com/support>
- <https://community.emc.com>

Where to get support

The Support website <https://www.dell.com/support> provides access to product licensing, documentation, advisories, downloads, and how-to and troubleshooting information. The information can enable you to resolve a product issue before you contact Support.

To access a product-specific page:

1. Go to <https://www.dell.com/support>.
2. In the search box, type a product name, and then from the list that appears, select the product.

Knowledgebase

The Knowledgebase contains applicable solutions that you can search for either by solution number (for example, KB000xxxxxx) or by keyword.

To search the Knowledgebase:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Knowledge Base**.
3. In the search box, type either the solution number or keywords. Optionally, you can limit the search to specific products by typing a product name in the search box, and then selecting the product from the list that appears.

Live chat

To participate in a live interactive chat with a support agent:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Contact Support**.
3. On the **Contact Information** page, click the relevant support, and then proceed.

Service requests

To obtain in-depth help from Licensing, submit a service request. To submit a service request:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Service Requests**.

Note

To create a service request, you must have a valid support agreement. For details about either an account or obtaining a valid support agreement, contact a sales representative. To get the details of a service request, in the *Service Request Number* field, type the service request number, and then click the right arrow.

To review an open service request:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Service Requests**.
3. On the **Service Requests** page, under **Manage Your Service Requests**, click **View All Dell Service Requests**.

Online communities

For peer contacts, conversations, and content on product support and solutions, go to the Community Network <https://community.emc.com>. Interactively engage with customers, partners, and certified professionals online.

How to provide feedback

Feedback helps to improve the accuracy, organization, and overall quality of publications. You can send feedback to DPAD.Doc.Feedback@emc.com.

CHAPTER 1

Overview of NAS Snapshot Features

This chapter includes the following topics:

- [NetWorker Snapshot Management for NAS feature description](#)..... 18
- [Components of the NAS snapshot network](#)..... 22
- [Licensing requirements](#)..... 26
- [Example NAS snapshot environments](#)..... 26

NetWorker Snapshot Management for NAS feature description

The NetWorker 9.0 and later server software includes optional features to enable NetWorker Snapshot Management for network-attached storage (NAS), also called NSM for NAS. Supported NAS devices are Isilon, VNX, VNX2, VNXe 3200, NetApp, and Unity Storage.


A NAS device may also be called a Network Data Management Protocol (NDMP) data server, data mover, NDMP client, or NAS filer.

Before you plan, configure, and administer the snapshot environment, become familiar with the concepts in this chapter. You need an advanced working knowledge of the NAS devices that you use with NetWorker Snapshot Management.

The *NetWorker Hardware Compatibility Guide* provides details on NAS versions that NetWorker supports.

The *NetWorker Release Notes* provides the details on NetWorker features, support, fixed problems, and known problems with workarounds for each release.

Figure 1 NAS flow diagram



NAS snapshot operations

A network-attached storage (NAS) device is a dedicated file server that provides primary data storage for application servers in a heterogeneous network environment. A NAS device is configured as a client of the NetWorker server (for example, by using the Client Configuration wizard) but you do not install NetWorker client software on the NAS device.

A NAS device presents stored data as file systems. Its snapshot and replication features are covered by NAS array licenses. NDMP enables the data transfer that is required between the primary NAS device storage and secondary storage on tape or VTL devices.

The amount of data that is stored on NAS devices can often be large, and the backup times for tape and VTL devices can exceed the available backup window. By scheduling local point-in-time snapshot copies of the NAS data, NetWorker can use whatever time is required to clone the data to backup media independently of the normal backup windows.

NetWorker works with NAS devices to perform the following snapshot and replication operations:

- Create local snapshot save sets of the specified data on the NAS device.
- Replicate the specified snapshot data on a source NAS device to a different location on the same device or to a location on a different NAS device.
- Perform an immediate or a delayed clone backup of specified snapshot save sets to secondary storage by using NDMP.
- Apply retention policies to manage the lifecycles of the snapshot, replication, and clone save sets.
- Recover the specified data from snapshots and clones.

Discovery of non-NetWorker snapshots

NetWorker can discover snapshots on the NAS device that NetWorker did not create.

NetWorker can manage these snapshots as follows:

- NetWorker catalogs all discovered snapshots as snapshot save sets.
- NetWorker can clone, recover, and report details of discovered snapshots but cannot delete them.
- You can schedule discovery or perform discovery manually.

Types of NAS snapshot-based backups

The type of network-attached storage (NAS) snapshot-based backup that you configure depends on where you intend to create and store the snapshot data, as follows:

- Local snapshot backup—The NetWorker server creates a snapshot of the data on the NAS device and retains the snapshot on the NAS device. The NetWorker server catalogs the snapshot as a backup in the media database and can perform a recovery from the snapshot.
- Replicated snapshot backup:
 - Isilon SyncIQ—The NetWorker server starts the data replication on the Isilon device. The device replicates the data to a local or remote Isilon device. NetWorker creates a snapshot of the replicated data and catalogs it as a backup. The NetWorker server can clone to conventional storage from that snapshot.
 - NetApp SnapVault and SnapMirror—The NetWorker server starts the data replication on the NetApp device. The source device creates a snapshot on the same device and then sends information only about new, changed, or deleted data to the remote mirror NetApp device. For SnapMirror, NetWorker catalogs the remote mirror as a backup and can clone to conventional storage from the mirror. NetWorker clones the data from the source volume.

[Example NAS snapshot environments](#) on page 26 provides more details.

Types of NAS snapshot and clone recoveries

The types of recovery that you can perform for snapshot-based save sets depends on the location of the snapshot backup, and where you want to recover the data:

- Snapshot recovery—NetWorker mounts the snapshot volume, where you can view the save sets and select the directories or the individual files to recover, from the snapshot volume. You can perform one of the following types of recoveries from the snapshot:
 - In-place—Recovers files and directories to the original locations on the NAS device.
 - Out-of-place—Recovers files and directories to a new location on the NAS device that you specify.
 - Directed recovery—Recovers files and directories to the NetWorker host that runs the recovery. You can configure the destination directory as an NFS or CIFS file share, which enables you to recover to a remote location. For example, another client or a NAS device.
 - In NetWorker 19.1, Isilon supports snapshot recovery to different cluster.
- Recovery from a clone—The same as recovery from a backup. You perform a conventional NetWorker recovery from the backup storage media.

Note

You can perform a clone directed restore to a similar NAS box only.

NetWorker clone support

NetWorker uses cloning to copy snapshots to the following types of media:

- Data Domain Boost, Advanced File Type Device (AFTD), Tape—You can clone any types of snapshots to these types of conventional media.

NetWorker cloning supports full and incremental cloning. Cloning can leverage traditional NetWorker restore methods.

Cloning is also supported to Data Domain CloudTier and CloudBoost.

NetWorker catalogs snapshots and clone copies in the media database as follows:

- File system backups—NetWorker records the contents of the snapshots in the client file index (CFI) only during a clone operation to conventional media.

The *NetWorker Administration Guide* provides details.

Note

NetWorker 8.2.x and earlier versions do not support the cloning of snapshot save sets.

Backup configuration methods

You can configure snapshot backups by using the Client Backup Configuration wizard or the Client Properties window in the NetWorker Administration interface:

- Client Backup Configuration wizard—It is recommended that you use the wizard to create and modify NAS snapshot configurations, because the wizard provides the correct sequence of steps and verifications. You cannot use the wizard to configure some special use directives and variables. After you use the wizard to

create the client, you can modify the client resource to specify directives and variables.

Recovery interfaces

You can recover snapshot-based data by using the following interfaces:

- **NetWorker Recovery Wizard**—The recommended interface to recover data from snapshots and conventional storage media.
- `nsrsnapadmin` command—Provides an interactive Command Line Interface (CLI) session for various snapshot-related operations, that you can use to recover data from snapshots and conventional storage media.
- `nsrnassnap_recover` command—Provides another CLI method that you can use to recover data from snapshots.

The *NetWorker Command Reference Guide* provides details of the NetWorker commands.

Monitoring and reporting NAS snapshot operations

NetWorker reporting features enable you to monitor NAS snapshot operations for each NetWorker client or network-attached storage (NAS) device, including snapshots that you did not create with NetWorker. NetWorker catalogs all snapshots as backup save sets. You can monitor the progress of the snapshot creation, mounting, deletion, and clone operations.

The NetWorker `nwnassnap.raw` log provides details of snapshot operations. Also, policy logs may provide more detailed information about the actions and workflows that have been run.

Internationalization support

NetWorker provides NAS snapshot operations with the standard NetWorker client support for non-ASCII international character sets.

Multiple rollover sessions for NAS

Multiple rollover is the process of creating copies of multiple NAS snapshots simultaneously. When data is moved from a NAS device to a NetWorker device, the clone operations happen simultaneously and is a much faster process. For example, if there are three snapshot savesets, the data movement happens simultaneously for the savesets based on the parallelism settings.

Clone action depends on the following parallelism values:

- If the number of savesets and the parallelism value is same, the clone action runs simultaneously. For example, if the number of savesets and the parallelism value is five, the clone action for the five savesets happen simultaneously.
- If the parallelism value is lesser than the number of savesets, the clone action runs simultaneously based on the value set. For example, if the parallelism value is two and the number of savesets are five, the clone action happens two at a time. It starts with the first two savesets while the other three are queued. Once it is complete, the clone action for the next two savesets start followed by the remaining one saveset.
- If the parallelism value is greater than the number of savesets, the clone action happens simultaneously. For example, if the parallelism value is six and the number of savesets are five, the clone action for the five savesets happen simultaneously.

Note

- Parallelism primarily depends on the NAS capabilities and must be set to a value appropriate for NAS being backed up.
 - If rollover by query is tried for multiple snapshots of savesets specified in the client with parallelism value more than the number of savesets, during clone operation the parallelism value is reduced to the number of saveset specified in the client.
 - During the multiple rollover session in NMC, the clone session displayed in the session tab of the monitoring window represents the entire clone process.
-

Creating a multiple rollover session for NAS

The following instructions describe how to create a multiple rollover session for NAS:

Procedure

1. Configure a NAS snapshot client using the NMC Client Configuration wizard by specifying the **NAS device name**, **Management credentials**, and **NDMP credentials**.
2. Add five NFS or CIFS NAS savesets for different numbers of snapshots.
3. Create a group and add the client to the group.
4. Create the policy, add the group created, and create workflows which has the snapshot and the clone action.
5. Select **modify client** properties for the client. **Under Global 1 of 2** tab, set the parallelism value and click **OK**.
6. Start the workflow of the policy.
7. Verify whether the backup is successful by checking the monitoring logs.
8. Verify that the clone operation has started after backup completion.
9. Validate whether the multiple rollover is successful by verifying the following:
 - a. In NMC Log section, check whether clones have started simultaneously according to the parallelism.
 - b. Verify whether the `nsrndmp_save` processes are started simultaneously for the savesets according to the parallelism value.
 - c. Verify that the number of clones generated are equal to the number of savesets using the `mminfo -avot` command.

Components of the NAS snapshot network

You can deploy various required and optional devices, hosts, and connectivity for network-attached storage (NAS) snapshot operations in a NetWorker datazone.

NAS devices

A network-attached storage (NAS) device is a dedicated file server that provides primary data storage that you can use for Network Data Management Protocol (NDMP) data operations. A NAS device may also be called an NDMP data server, data mover, NDMP client, or NAS filer. You must configure a NAS device as a client of the

NetWorker server, but you do not install NetWorker client software on the NAS device.

All the supported NAS devices can support a directed recovery to the host that runs the recovery. You can recover to an NFS file share directory on that host, which enables recovery to a remote host or remote device.

Table 3 NAS device support and limitations

NAS device	Characteristics
Isilon	<p>This device supports:</p> <ul style="list-style-type: none"> • Snapshots with SnapshotIQ™, which requires OneFS® 7.02 or later. • Replications with SyncIQ®, which requires OneFS 7.1 or later. • Snapshots and replication at the directory level. • In-place and out-of-place recovery of snapshot directories. • Directed recovery. <p>This device does not support in-place or out-of-place recovery from a remotely replicated snapshot.</p>
VNX/VNX2	<p>This device supports:</p> <ul style="list-style-type: none"> • Snapshots at the file system level. • The Client resource must specify the data mover name. • Supports directed recovery. <p>This device does not support:</p> <ul style="list-style-type: none"> • In-place or out-of-place snapshot recovery. • Replication.
VNXe	<p>This device requires:</p> <ul style="list-style-type: none"> • Unisphere Management client (uemcli) on the NetWorker server. • The Client resource must specify the NAS server name. <p>This device supports:</p> <ul style="list-style-type: none"> • Snapshots at the file system level. • Directed recovery. <p>This device does not support:</p> <ul style="list-style-type: none"> • In-place or out-of-place snapshot recovery. • Replication.
Unity	<p>This device requires:</p> <ul style="list-style-type: none"> • Unisphere Management client (uemcli) on the NetWorker server. • The Client resource must specify the NAS server name. <p>This device supports:</p> <ul style="list-style-type: none"> • Snapshots at the file system level. • Directed recovery. <p>This device does not support:</p> <ul style="list-style-type: none"> • In-place or out-of-place snapshot recovery.

Table 3 NAS device support and limitations (continued)

NAS device	Characteristics
	<ul style="list-style-type: none"> Replication.
NetApp	<p>This device requires operating system libraries from the NetApp Manageability SDK, available from NetApp. This device supports:</p> <ul style="list-style-type: none"> Snapshots using 7-Mode and Cluster Mode, and requires NetAppData ONTAP 8.0 or later. <hr/> <p>Note</p> <p>The <i>NetWorker E-LAB Navigator</i> provides details about Data ONTAP versions with NetApp Cluster-Mode support.</p> <hr/> <ul style="list-style-type: none"> SnapMirror for Local and Remote Replication of snapshots in C-mode and 7-Mode. SnapVault for Local and Remote Replication of snapshots in 7-Mode. SnapVault for Local and Remote Replication only of snapshots in C-Mode. Cluster Aware Backup (CAB) for affinity awareness of volumes and tape devices. In-place and out-of-place recovery of snapshot files but not directories or save sets (except if the save set is a single file), and only to the same volume as the source. Directed recovery. <hr/> <p>Note</p> <p>Earlier when you used NetWorker Snapshot Management (NSM) with a NetApp appliance to perform snapshot backups, you were restricted to connect to the NAS device using only port number 80. However, with NetWorker 19.1 and later you can connect to the NAS device using a user defined port.</p> <hr/> <p>This device does not support:</p> <ul style="list-style-type: none"> NetApp on Linux systems does not support a temporary mount point for snapshot browse and recover, but must use an existing, manually configured NFS mount point. SnapMirror in-place or out-of-place recovery, does not support recovery for local and remotely replicated snapshots.

NAS snapshot naming convention

In NetWorker 19.1, the network-attached storage (NAS) snapshot name format that is generated by NSM, has changed from "EMCDP-<long client id>-<index>-<UTC time stamp>" to "NSMNAS-<snapshot time [yyyymmddhhmmss]>-<session id>-<index>."

For example:

- Old format: EMCDP-eabd217f-00000004-56b5453f-56b5453e-00255000-7f76b156-0-1454724889

- New format: NSMNAS-20160916114959-1474051799-0

You must add the attribute in the **Application Information** area when you create the client. You also now have an option to change the prefix of the snapshot name using the attribute NSR_SNAPSHOT_NAME_PREFIX.

For example: NSR_SNAPSHOT_NAME_PREFIX=<Hostx-myfs1> The prefix can be a maximum length of 40 characters.

Note

The names must be alphanumeric with acceptable special characters, such as -, _.

The replicated snapshot names will not reflect these prefixes. The prefix for replicated snapshot depends on the NAS replication configuration settings.

NDMP environment

NetWorker software uses the Network Data Management Protocol (NDMP) functionality to enable access to storage in a heterogeneous network environment. NDMP uses TCP/IP to control the movement of the data, and specifies various device drivers to store the data on devices.

NetWorker server

The NetWorker server manages the NAS devices and the configuration settings that are required to create the snapshots and perform the clone and the recovery operations. It is the Data Management Agent (DMA).

NetWorker client

The network-attached storage (NAS) devices do not run NetWorker client software. However, for remote replication, if you want to perform backup, clone, or indexing actions, you must have a client resource for the remote target NAS device, and you must create a NetWorker Client resource for the NAS device. The NetWorker Client resource specifies the snapshot replication target and provides the credentials to the NetWorker server for replication, cloning, and recovery operations.

NetWorker storage node

The NetWorker storage node hosts and manages the devices that are used for cloning to secondary storage media, such as tape or VTL devices. The storage node is the Network Data Management Protocol (NDMP) tape server.

If you plan to use network-attached storage (NAS) devices only to create and recover snapshots without cloning to secondary storage, then the use of a storage node is optional.

Backup storage media

NetWorker uses Network Data Management Protocol (NDMP) services to clone network-attached storage (NAS) snapshot save sets to tape or VTL storage media, DD, cloud tier, cloud boost, and AFTD. During cloning, NetWorker catalogs the clone save set files in its client file index.

Licensing requirements

Network-attached storage (NAS) snapshot and replication features are covered by NAS array licenses.

The *NetWorker Licensing Guide* and a sales representative can provide details about licensing for NetWorker Snapshot Management (NSM).

Example NAS snapshot environments

Plan the network-attached storage (NAS) environment to manage replication and snapshot operations efficiently as illustrated by the following examples.

Example of a snapshot and rollover to storage media

This example illustrates the data flow in a typical environment for a local snapshot backup of network-attached storage (NAS) client data with an optional rollover to a Network Data Management Protocol (NDMP) storage device.

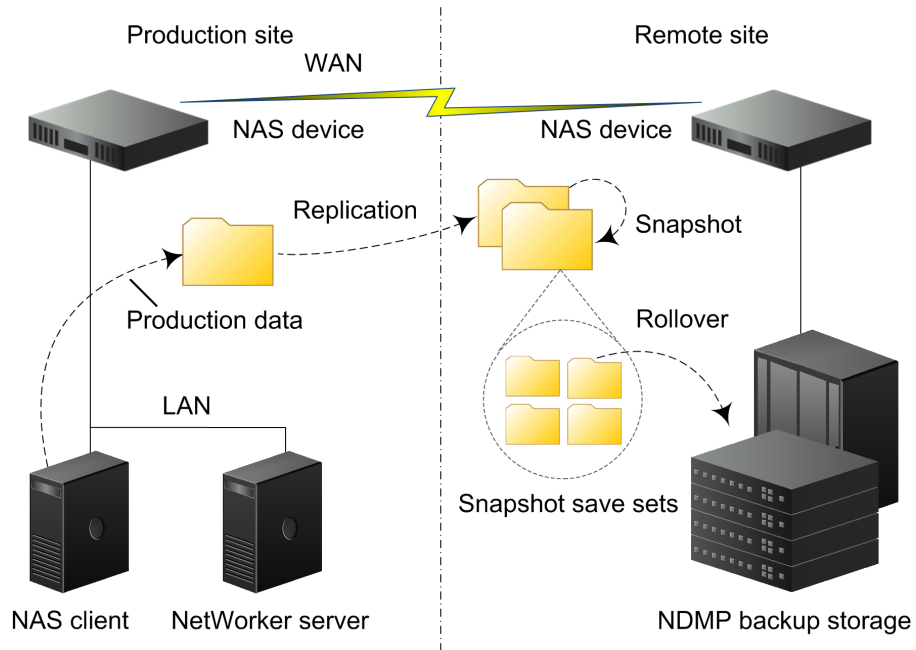
The process flow is as follows:

1. The NAS client writes its production data to primary storage on the NAS device.
2. The NetWorker server, at a specified time, instructs the NAS device to use its proprietary operations to create a snapshot (replication) of production file systems, as specified in the NetWorker Client resource for the device. The NetWorker server catalogs the snapshot as a backup save set from which recoveries can be made.
3. The NetWorker workflow may also include a clone action. In that case, the NetWorker server instructs the NAS device to clone the snapshot save set to secondary storage such as an NDMP tape or a VTL device.
4. The NetWorker server manages the snapshot save set and any secondary backup save set according to the options that are configured in the NetWorker Data Protection Policy.
5. The NAS device retains the snapshot until the snapshot expires or until a NetWorker operator deletes it.

Example of a snapshot from replication

This example illustrates the data flow for a snapshot of network-attached storage (NAS) client data that is remotely replicated, with an optional clone to Network Data Management Protocol (NDMP) storage media.

Figure 2 Remote replication with snapshot and clone to storage



The process flow is as follows:

1. The NAS client writes its production data to primary storage on the NAS device.
2. The NetWorker server, at a specified time, instructs the NAS device to create a remote replication/snapshot of production file systems. The NAS device uses its proprietary operations to create the replication on the remote NAS device:
 - For Isilon devices, the production file systems to replicate are specified by the NetWorker Client resource for the device.
 - For NetApp devices, the device creates a local full point-in-time (PIT) snapshot of the production data on the same device, and then updates the Local and Remote NetApp mirror device with any new, changed, or deleted data.
3. The NetWorker server, using the schedule that is specified in the Data Protection Policy for the remote NAS device, instructs the device to use its proprietary operations to create a snapshot/replication of the replicated file systems. The NetWorker server catalogs the remote snapshot as a backup save set from which recoveries can be made.
4. The NetWorker server, if specified in the NetWorker Data Protection policy may also include a clone action, in that case instructs the remote NAS device to clone the snapshot save set to secondary storage such as an NDMP tape or a VTL device, DD, cloud tier, cloud boost, and AFTD.
5. The NetWorker server manages the remote snapshot save set and any secondary backup save set according to the retention options and other options configured in the Data Protection Policy.

6. The NAS device retains the snapshot until the snapshot expires or until a NetWorker operator deletes it.

Examples of snapshot recovery

This example illustrates the data flow for the recovery of directories or files from a snapshot save set. NetWorker supports the save set recovery of network-attached storage (NAS) snapshots where the entire save set is restored. VNX and Isilon support save set recovery. NetApp does not support save set recovery.

The process flow is as follows:

1. You select the snapshot volume that contains the data that you want to recover. NetWorker mounts the snapshot volume.
2. You browse for the files, file systems, or volumes that you want to recover.
3. You specify where to recover the data on the NAS device host or on a NetWorker host:
 - In-place recovery—Recover to the original location on the NAS device.
 - Out-of-place recovery—Recover to a different location on the NAS device.
 - Directed recovery—Recover to a directory on the NetWorker client host that is running the recovery. The recovery directory may be a local directory or a locally mounted CIFS or NFS remote file system. In NetWorker 19.1, Isilon supports snapshot recovery to different Isilon clusters.
4. NetWorker copies the data from the snapshot volume to the specified location.

Figure 3 In-place recovery restores data to the original location on the source device




Figure 4 Out-of-place recovery restores data to a different location on the source device





Figure 5 Directed recovery restores data to the NetWorker client that runs the recovery



CHAPTER 2

Data Protection Policies

This chapter includes the following topics:

- [Introduction](#)..... 32
- [Default data protection policies in NMC's NetWorker Administration window](#)... 32
- [Strategies for NAS snapshot backups](#)..... 33
- [Overview of configuring a new data protection policy](#)..... 33
- [Creating a policy](#)..... 34
- [Create a workflow for a new policy in NetWorker Administration](#)..... 36
- [Protection groups for NAS snapshot devices](#)..... 38
- [Supported actions in snapshot workflows](#)..... 44
- [Action in snapshot workflows](#)..... 45
- [Visual representation of NAS snapshot workflows](#)..... 64

Introduction

You can protect your data on NAS devices in two ways:

- Using the NDMP protocol.
- Using snapshot management technology.

NDMP follows the traditional backup strategy.

For a detailed overview about creating, editing, and deleting groups and policies, refer to the Data Protection Policies chapter in the *NetWorker Administration Guide*.

Default data protection policies in NMC's NetWorker Administration window

The NMC **NetWorker Administration** window provides you with pre-configured data protection policies that you can use immediately to protect the environment, modify to suit the environment, or use as an example to create resources and configurations. To use these pre-configured data protection policies, you must add clients to the appropriate group resource.

Note

NMC also includes a pre-configured Server Protection policy to protect the NetWorker and NMC server databases.

Platinum policy

The Platinum policy provides an example of a data protection policy for an environment that contains supported storage arrays or storage appliances and requires backup data redundancy. The policy contains one workflow with two actions, a snapshot backup action, followed by a clone action.

Figure 6 Platinum policy configuration



Gold policy

The Gold policy provides an example of a data protection policy for an environment that contains virtual machines and requires backup data redundancy.

Silver policy

The Silver policy provides an example of a data protection policy for an environment that contains machines where file systems or applications are running and requires backup data redundancy.

Bronze policy

The Bronze policy provides an example of a data protection policy for an environment that contains machines where file systems or applications are running.

Strategies for NAS snapshot backups

Multiple data protection policy strategies are available to help you optimize the way that snapshots are performed and backed up.

When you protect NAS devices by using snapshot technology, there are four actions that you can perform on the NAS data:

- Create a snapshot backup.
- Discover snapshots on the NAS device that were created independent of NetWorker:
 - Add newly discovered snapshots to the media database on the NetWorker server.
 - Remove previously discovered snapshots from the media database when the snapshots are no longer available on the NAS device.
- Update client file index entries with information about:
 - Snapshots that are created by NetWorker.
 - Snapshots created independent of NetWorker.
- Clone a snapshot.

Overview of configuring a new data protection policy

The following steps are an overview of the tasks to complete, to create and configure a data protection policy.

Procedure

1. Create a policy resource.

When you create a policy, you specify the name and notification settings for the policy.

2. Within the policy, create a workflow resource for each data type.

For example, create one workflow to protect file system data and one workflow to protect application data. When you create a workflow, you specify the name of the workflow, the time to start the workflow, notification settings for the workflow, and the protection group to which the workflow applies.

3. Create a protection group resource.

The type of group that you create depends on the types of clients and data that you want to protect. The actions that appear for a group depend on the group type.

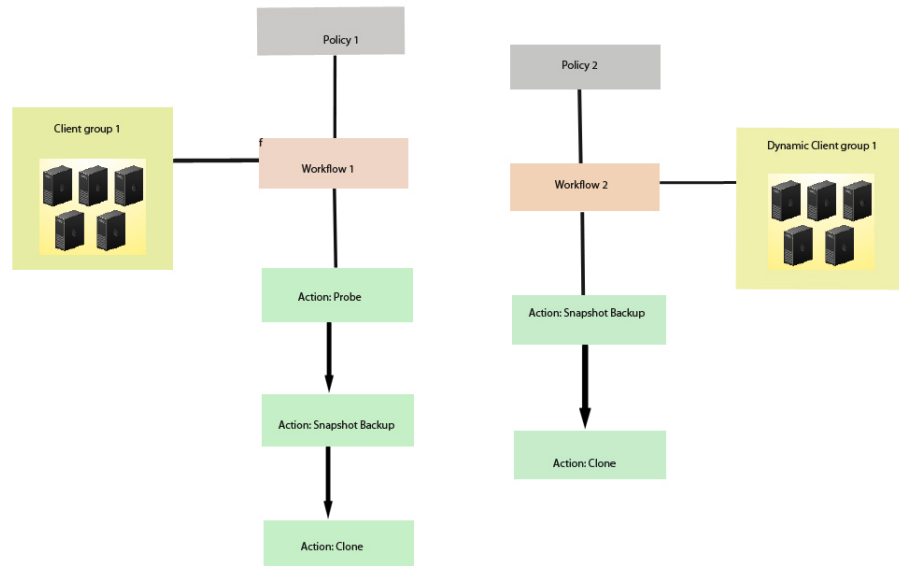
4. Create one or more action resources for the workflow resource.
5. Configure client resources, to define the backup data that you want to protect, and then assign the client resources to a protection group.

Example 1 Example of a data protection policy with 2 workflows

The following figure illustrates a policy with two different workflows. Workflow 1 performs a probe action, then a backup of the client resources in Client group 1, and then a clone of the save sets from the backups. Workflow 2 performs a backup of the

Example 1 Example of a data protection policy with 2 workflows (continued)

client resources in Dynamic client group 1, and then a clone of the save sets from the backup.

Figure 7 Data protection policy example**Note**

For more information on configuring a new data protection policy using the NetWorker Management Web UI, see the *NetWorker Administration Guide*.

Creating a policy

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Policies**, and then select **New**.
The **Create Policy** dialog box appears.
3. On the **General** tab, in the **Name** field, type a name for the policy.
The maximum number of characters for the policy name is 64.
 - Legal Characters: _ : - + = # , . % @
 - Illegal Characters: / \ * ? [] () \$! ^ ; ' " ` ~ > < & | { }

Note

After you create a policy, the **Name** attribute is read-only.

4. In the **Comment** field, type a description for the policy.
5. From the **Send Notifications** list, select whether to send notifications for the policy:
 - To avoid sending notifications, select **Never**.
 - To send notifications with information about each successful and failed workflow and action, after the policy completes all the actions, select **On Completion**.
 - To send a notification with information about each failed workflow and action, after the policy completes all the actions, select **On Failure**.
6. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

To send email messages or the `smtpmail` application on Windows, use the default mailer program on Linux:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:

```
nsrlog -f policy_notifications.log
```

- On Linux, to send an email notification, type the following command:

```
mail -s subject recipient
```

- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

```
/usr/sbin/sendmail -v recipient_email "subject_text"
```

- On Windows, to send a notification email, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.
- `recipient1@mailserver`—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

7. To specify the Restricted Data Zone (RDZ) for the policy, select the **Restricted Data Zones** tab, and then select the RDZ from the list.
8. Click **OK**.

After you finish

Create the workflows and actions for the policy.

Create a workflow for a new policy in NetWorker Administration

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the left pane, expand **Policies**, and then select the policy that you created.
3. In the right pane, select **Create a new workflow**.
4. In the **Name** field, type the name of the workflow.

The maximum number of allowed characters for the **Name** field is 64.

- Legal Characters: `_ : - + = # , . % @`
- Illegal Characters: `/*?[]()$!^;'"`~><&|{}`

5. In the **Comment** box, type a description for the workflow.

The maximum number of allowed characters for the **Comment** field is 128.

6. From the **Send Notifications** list, select how to send notifications for the workflow:
 - To use the notification configuration that is defined in the policy resource to specify when to send a notification, select **Set at policy level**.
 - To send notifications with information about each successful and failed workflow and action, after the workflow completes all the actions, select **On Completion**.
 - To send notifications with information about each failed workflow and action, after the workflow completes all the actions, select **On Failure**.
7. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages, or use the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:


```
nsrlog -f policy_notifications.log
```
- On Linux, to send an email notification, type the following command:


```
mail -s subject recipient
```
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

```
/usr/sbin/sendmail -v recipient_email "subject_text"
```

- On Windows, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- **-s *subject***—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- **-h *mailserver***—Specifies the hostname of the mail server to use to relay the SMTP email message.
- ***recipient1@mailserver***—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

8. In the **Running** section, perform the following steps to specify when and how often the workflow runs:
 - a. To ensure that the actions that are contained in the workflow run when the policy or workflow starts, in the **Enabled** box, leave the option selected. To prevent the actions in the workflow from running when the policy or workflow that contains the action starts, clear this option.
 - b. To start the workflow at the time that is specified in the **Start time** attribute, on the days that are defined in the action resource, in the **AutoStart Enabled** box, leave the option selected. To prevent the workflow from starting at the time that is specified in the **Start time** attribute, clear this option.
 - c. To specify the time to start the actions in the workflow, in the **Start Time** attribute, use the spin boxes.
The default value is 9:00 PM.
 - d. To specify how frequently to run the actions that are defined in the workflow over a 24-hour period, use the **Interval** attribute spin boxes. If you are performing transaction log backup as part of application-consistent protection, you must specify a value for this attribute in order for incremental transaction log backup of SQL databases to occur.
The default **Interval** attribute value is 24 hours, or once a day. When you select a value that is less than 24 hours, the **Interval End** attribute appears. To specify the last start time in a defined interval period, use the spin boxes.
 - e. To specify the duration of time in which NetWorker can manually or automatically restart a failed or canceled workflow, in the **Restart Window** attribute, use the spin boxes.
If the restart window has elapsed, NetWorker considers the restart as a new run of the workflow. NetWorker calculates the restart window from the start of the last incomplete workflow. The default value is 24 hours.
For example, if the **Start Time** is 7:00 PM, the **Interval** is 1 hour, and the **Interval End** is 11:00 PM., then the workflow automatically starts every hour beginning at 7:00 PM. and the last start time is 11:00 PM.

9. To create the workflow, click **OK**.

After you finish

Create the actions that will occur in the workflow, and then assign a group to the workflow. If a workflow does not contain a group, a policy does not perform any actions.

Protection groups for NAS snapshot devices

Create protection groups for NAS devices that support snapshot technology.

The type of protection group that you create to protect snapshot data depends on the data operation that you want to perform:

- **Basic client group**—A static list of client resources to back up.
- **Dynamic client group**—A dynamic list of client resources to back up. A dynamic client group automatically generates a list of the client resources that use a client tag which matches the client tag that is specified for the group.
- **NAS device group**—Use this group to define a static list of NAS devices on which you want to perform discovery operations, to detect NAS snapshots and backups that were created or deleted independent of NetWorker. The discovery operation will create entries in the media database for the new backups and remove media database entries for deleted backups.
- **Save set ID List**—Use this group to generate a static list of NAS snapshot backup save sets in the media database. Use this list to perform one of the following operations:
 - Create client file index entries for the save sets.
 - Clone the save sets.
- **Save Set Query**—Use this group to generate a dynamic list of NAS snapshot backup save sets in the media database, based on matching criteria. Criteria that you can define in a query group includes:
 - Date and time range for the save set.
 - Backup level of the save set.
 - Client resource that is associated with a save set.
 - Policy, workflow, group, or action that is used to generate the save set.

Note

A NAS environment does not support the use of the Query by Pool or the Number of clone copies criteria to generate a dynamic list of NAS snapshot backup save sets.

Use a query group to generate a save set list to perform one of the following operations:

- Create client file index entries for the save sets.
- Clone the save sets.

Note

For query groups that include NAS devices you cannot specify a volume. If you use this query to do a clone, the clone will fail.

Create multiple protection groups to perform different types of snapshot backups or discovery for different devices, or to perform snapshot backups or discovery on different schedules. For example:

- Create one basic client group for a workflow that performs a snapshot backup of NAS devices.
- Create one NAS device group for a workflow that performs discovery, and then a query group to generate a list of save sets for a workflow that generates index data for the snapshot save sets at a different time.

Creating a basic client group

Use basic client groups to specify a static list of client resources for a traditional backup, a check connectivity action, or a probe action.

Before you begin

Create the policy and workflow resources in which to add the protection group to.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Groups** and select **New** from the drop-down, or right-click an existing group and select **Edit** from the drop-down.

The **Create Group** or **Edit Group** dialog box appears, with the **General** tab selected.

3. In the **Name** attribute, type a name for the group.

The maximum number of characters for the group name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: / \ * ? [] () \$! ^ ; ' " ` ~ > < & | { }

Note

After you create a group, the **Name** attribute is read-only.

4. From the **Group Type** list, leave the default selection of **Clients**.
5. In the **Comment** field, type a description of the group.
6. From the **Policy-Workflow** list, select the workflow that you want to assign the group to.

Note

You can also assign the group to a workflow when you create or edit a workflow.

7. (Optional) To specify the Restricted Datazone (RDZ) for the group, on the **Restricted Datazones** tab, select the RDZ from the list.
8. Click **OK**.

After you finish

Create Client resources. Assign clients to a protection group, by using the Client Configuration wizard or the **General** tab on the **Client Properties** page.

Creating a dynamic client group

Dynamic client groups automatically include group settings when you add client resources to the NetWorker datazone. You can configure a dynamic group to include all the clients on the NetWorker server or you can configure the dynamic client group to perform a query that generates a list of clients that is based on a matching tag value.

A tag is a string attribute that you define in a Client resource. When an action starts in a workflow that is a member of a tagged dynamic protection group, the policy engine dynamically generates a list of client resources that match the tag value.

Use dynamic client groups to specify a dynamic list of Client resources for a traditional backup, a probe action, a check connectivity action, or a server backup action.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Groups** and select **New** from the drop-down, or right-click an existing group and select **Edit** from the drop-down.

The **Create Group** or **Edit Group** dialog box appears, with the **General** tab selected.

3. In the **Name** attribute, type a name for the group.

The maximum number of characters for the group name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[](\$!^;"'~><&|{}

Note

After you create a group, the **Name** attribute is read-only.

4. From the **Group Type** list, select **Dynamic Clients**. For steps 5 to 8, follow the instructions given in [Creating a client group](#).

Creating a NAS device group

A NAS device group defines a list of NAS devices for discovery.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Groups** and select **New** from the drop-down, or right-click an existing group and select **Edit** from the drop-down.

The **Create Group** or **Edit Group** dialog box appears, with the **General** tab selected.

3. In the **Name** attribute, type a name for the group.

The maximum number of characters for the group name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[](\$!^;"'~><&|{}

Note

After you create a group, the **Name** attribute is read-only.

4. From the **Group Type** list, select **NAS Device**.
 5. In the **Comment** field, type a description of the group.
 6. From the **Policy-Workflow** list, select the workflow that you want to assign the group to.
-

Note

You can also assign the group to a workflow when you create or edit a workflow.

7. Add and remove clients in the NAS Device section, under the Add/Remove selected Clients to/from the following Groups area:
 - a. To add a client, select the checkbox next to the device, and then click **OK**.
If you want to select all the clients that are listed, click **Highlight All**.
 - b. To remove a client, select the checkbox next to the device, and then click **OK**.
8. (Optional) To specify the Restricted Datazone (RDZ) for the group, on the **Restricted Datazones** tab, select the RDZ from the list.
9. Click **OK**.

Creating a save set group

A save set group defines a static list of save sets for cloning or for snapshot index generation.

Before you begin

Determine the save set ID or clone ID (ssid/clonid) of the save sets for the group by using the **Administration > Media** user interface or the `mminfo` command.

Procedure

1. In the **Administration** window, click **Protection**.
 2. In the expanded left pane, right-click **Groups**, and then select **New**.
The **Create Group** dialog box appears, starting with the **General** tab.
 3. In the **Name** field, type a name for the group.
The maximum number of characters for the group name is 64.
 - Legal Characters: `_ : - + = # , . % @`
 - Illegal Characters: `/*?[]$!^;'"`~><&|{}`
-

Note

After you create a group, the **Name** attribute is read-only.

4. From the **Group Type** list, select **Save Set ID List**.
5. In the **Comment** field, type a description of the group.

6. (Optional) To associate the group with a workflow, from the **Workflow (Policy)** list, select the workflow.

You can also assign the group to a workflow when you create or edit a workflow.

7. In the **Clone specific save sets (save set ID/clone ID)** field, type the save set ID/clone ID (ssid/clonid) identifiers.

To specify multiple entries, type each value on a separate line.

8. To specify the Restricted Data Zone (RDZ) for the group, select the **Restricted Data Zones** tab, and then select the RDZ from the list.
9. Click **OK**.

Creating a query group

A query group defines a list of save sets for cloning or snapshot index generation, based on a list of save set criteria.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Groups**, and then select **New**.
The **Create Group** dialog box appears, starting with the **General** tab.
3. In the **Name** field, type a name for the group.

The maximum number of characters for the group name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[]()\$!^;'"`~><&|{ }

Note

After you create a group, the **Name** attribute is read-only.

4. From the **Group Type** list, select **Save Set Query**.
5. In the **Comment** field, type a description of the group.
6. (Optional) To associate the group with a workflow, from the **Workflow (Policy)** list, select the workflow.
You can also assign the group to a workflow when you create or edit a workflow.
7. Specify one or more of the save set criteria in the following table.

Note

When you specify more than one save set criteria, the list of save sets only includes save sets that match all the specified criteria.

Table 4 Save set criteria

Criteria	Description
Date and time range	Specify the start date and time range for the save sets.

Table 4 Save set criteria (continued)

Criteria	Description
	<p>To specify the current date and time as the end date for the range, select Up to now.</p> <p>To specify a time period, select Up to.</p>
Backup level	<p>In the Filter save sets by level section, next to the backup level for the save set, select the full checkbox.</p> <hr/> <p>Note</p> <p>Only the full backup level is applicable for network-attached storage (NAS) devices.</p> <hr/>
Limit the number of clones	<p>Specify the number for the limit in the Limit number of clones list. The clone limit is the maximum number of clone instances that can be created for the save set. By default, the value is set to 1, and cannot be changed for NAS or Block.</p> <hr/> <p>Note</p> <p>When this criteria is set to 1, which is the default value, you may experience volume outage issues with Data Domain and advanced file type devices.</p> <hr/>
Client	Next to one or more client resources that are associated with the save set in the Client list, select the checkbox.
Policy	Next to the policy used to generate the save set in the Policy list, select the checkbox.
Workflow	Next to the workflow used to generate the save set in the Workflow list, select the checkbox.
Action	Next to the action used to generate the save set in the Action list, select the checkbox.
Group	Next to the group associated with the save set in the Group list, select the checkbox.
Pools	<p>Next to the media pool on which the save set is stored in the Pools list, select the checkbox.</p> <hr/> <p>Note</p> <p>You cannot select Pools for NAS.</p> <hr/>
Name	<p>In the Filter save sets by name field, specify the name of the save set.</p> <hr/> <p>Note</p> <p>You cannot use wildcards to specify the save set name.</p> <hr/>

If you specify multiple criteria, the save set must match all the criteria to belong to the group.

8. To specify the Restricted Data Zone (RDZ) for the group, select the **Restricted Data Zones** tab, and then select the RDZ from the list.
9. Click **OK**.

Supported actions in snapshot workflows

Snapshot workflow supports the following actions types.

Probe

A probe action runs a user-defined script on a NetWorker client before the start of a backup. A user-defined script is any program that passes a return code. If the return code is 0 (zero), then a client backup is required. If the return code is 1, then a client backup is not required.

Only a backup action can follow a probe action.

Check connectivity

A check connectivity action tests the connectivity between the clients and the NetWorker server before the start of a probe or backup action occurs. If the connectivity test fails, then the probe action and backup action does not start for the client.

Snapshot backup

A snapshot backup action performs a snapshot on a supported NAS device, backs up the snapshot, and then generates a save set entry for the snapshot backup in the NetWorker media database.

Note

A snapshot backup does not generate an entry for the snapshot backup in the client file index. To index a snapshot backup, use the generate index action after the snapshot backup action.

Discover

A discover action discovers snapshots on a NAS device that the current NetWorker server did not create. NetWorker then creates save set entries for the discovered snapshots on a NAS client in the media database.

The discover action also compares the save sets in the media database for the NAS client against the actual snapshots that are stored on the client. If the media database entry does not have a corresponding snapshot on the NAS client, then the media database entry is removed.

Note

A discover action does not generate entries for the discovered snapshots in the client file indexes. To index a discovered snapshot, use the generate index action after the discover action.

Generate index

A generate index action creates entries in the client file indexes for snapshot backups and discovered snapshots.

You can configure a generate index action to occur automatically in a workflow after a snapshot backup or discover action. Alternatively, you can configure a generate index action to occur separately for a group of save sets.

Clone

A clone action creates a copy of one or more save sets. Cloning enables secure offsite storage, the transfer of data from one location to another, and the verification of backups.

You can configure a clone action to occur after a backup in a single workflow, or concurrently with a backup action in a single workflow. You can use save set and query groups to define a specific list of save sets to clone, in a separate workflow.

Note

The clone action clones the scheduled backup save sets only, and it does not clone the manual backup save sets. Some NetWorker module backups might appear to be scheduled backups that are initiated by a policy backup action, but they are manual backups because they are initiated or converted by a database or application. The *NetWorker Module for Databases and Applications Administration Guide* and the *NetWorker Module for SAP Administration Guide* provides more details.

Action in snapshot workflows

Workflows enable you to chain together multiple actions and run them sequentially or concurrently.

You can perform a snapshot or you can use the discover action to discover snapshots on the NAS device that the current NetWorker server did not create, and create save set entries in the NetWorker media database for the snapshots.

Neither the snapshot backup action nor the discover action create entries for snapshots in the client file indexes. Use a generate index action to create snapshot entries in the client file indexes. You can configure the generate index action to occur automatically after the snapshot backup action or the discover action as part of a single workflow. To index the snapshot at a different time than the snapshot backup or discovery, create a separate workflow with the generate index action.


The *NetWorker Snapshot Management for NAS Devices Integration Guide* provides details on backing up NAS devices with NetWorker.

The following supported actions can follow the lead action and other actions in a workflow.

All possible workflow actions for a snapshot backup


You can perform a check connectivity and probe action before a snapshot backup action, and a clone action after the snapshot action.

Figure 8 All possible workflow actions for a snapshot backup




Workflow path from a snapshot backup action

You can perform a generate index action (to generate an index of the snapshot) or a clone action after a snapshot backup action.

Figure 9 Workflow path from a snapshot backup action**Workflow path from a snapshot discover action**

When you perform snapshot discovery, the discover action must be the first action in the workflow. You can optionally include a generate index or clone action concurrently with or after the discovery.

Figure 10 Workflow path from a discover action

Creating a probe action

A probe action runs a user-defined script on a NetWorker client before the start of a backup. A user-defined script is any program that passes a return code. If the return code is 0 (zero), then a client backup is required. If the return code is 1, then a client backup is not required.

Before you begin

- Create the probe resource script on the NetWorker clients that use the probe. Create a client probe resource on the NetWorker server. Associate the client probe resource with the client resource on the NetWorker server.
- Create the policy and workflow that contain the action.
- Optional. Create a check connectivity action to precede the probe action in the workflow. A check connectivity action is the only supported action that can precede a probe action in a workflow.

Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
 - Illegal Characters: \ ^ * ? [] () \$! ^ ; ' " ` ~ ~ > < & | { }
3. In the **Comment** field, type a description for the action.
 4. To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.



Note

When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

5. From the **Action Type** list, select **Probe**.
6. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
7. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
8. Specify a weekly or monthly schedule for the action:
 - To specify a schedule for each day of the week, select **Weekly by day**.
 - To specify a schedule for each day of the month, select **Monthly by day**.
9. Specify the days to probe the client:
 - To perform a probe action on a specific day, click the **Execute** icon on the day.
 - To skip a probe action, click the **Skip** icon on the day.
 - To perform a probe action every day, select **Execute** from the list, and then click **Make All**.

The following table provides details on the icons.

Table 5 Schedule icons

Icon	Label	Description
	Execute	Perform the probe on this day.
	Skip	Do not perform a probe on this day.

10. Click **Next**.
The **Specify the Probe Options** page appears.
11. Specify when to start the subsequent backup action:
 - To start the backup only if all the probes associated with client resources in the assigned group succeed, select the **Start backup only after all probes succeed** checkbox.

- To start the backup if any of the probes are associated with a client resource in the assigned group succeed, clear the **Start backup only after all probes succeed** checkbox.

12. Click **Next**.

The **Specify the Advanced Options** page appears.

13. In the **Retries** field, specify the number of times that NetWorker should retry a failed probe or backup action, before NetWorker considers the action as failed. When the **Retries** value is 0, NetWorker does not retry a failed probe or backup action.

Note

The **Retries** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option for other actions, NetWorker ignores the values.

14. In the **Retry Delay** field, specify a delay in seconds to wait before retrying a failed probe or backup action. When the **Retry Delay** value is 0, NetWorker retries the failed probe or backup action immediately.

Note

The **Retry Delay** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. When you specify a value for this option in other actions, NetWorker ignores the values.

15. In the **Inactivity Timeout** field, specify the maximum number of minutes that a job run by an action can try to respond to the server.

If the job does not respond within the specified time, the server considers the job a failure and NetWorker retries the job immediately to ensure that no time is lost due to failures.

Increase the timeout value if a backup consistently stops due to inactivity. Inactivity might occur for backups of large save sets, backups of save sets with large sparse files, and incremental backups of many small static files.

Note

The **Inactivity Timeout** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option in other actions, NetWorker ignores the value.

16. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.

17. From the **Failure Impact** list, specify what to do when a job fails:

- To continue the workflow when there are job failures, select **Continue**.
- To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.

Note

The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.

- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.
-

Note

If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

18. Do not change the default selections for the Notification group box. NetWorker does not support notifications for probe actions and ignores and specified values.
19. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
20. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
21. (Optional) In **Start Time** specify the time to start the action.
Use the spin boxes to set the hour and minute values, and select one of the following options from the drop-down list:
 - **Disabled**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
22. (Optional) Configure overrides for the task that is scheduled on a specific day.
To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:
 - Select the day in the calendar, which changes the action task for the specific day.
 - Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.
 - To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

Note

- You can edit or add the rules in the **Override** field.
 - To remove an override, delete the entry from the **Override** field.
-

23. Click **Next**.

The **Action Configuration Summary** page appears.

24. Review the settings that you specified for the action, and then click **Configure**.

Creating a check connectivity action

A check connectivity action tests the connectivity between the clients and the NetWorker server, usually before another action such as a backup occurs.

Before you begin

Create the policy and the workflow that contain the action. The check connectivity action should be the first action in the workflow.

Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: / \ * ? [] () \$! ^ ; ' " ` ~ > < & | { }

3. In the **Comment** field, type a description for the action.
4. To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.

Note



When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

5. From the **Action Type** list, select **Check Connectivity**.
6. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
7. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.

8. Specify a weekly or monthly schedule for the action:
 - To specify a schedule for each day of the week, select **Weekly by day**.
 - To specify a schedule for each day of the month, select **Monthly by day**.
9. Specify the days to check connectivity with the client:
 - To check connectivity on a specific day, click the **Execute** icon on the day.
 - To skip a connectivity check on a specific day, click the **Skip** icon on the day.
 - To check connectivity every day, select **Execute** from the list, and then click **Make All**.

The following table provides details about the icons.

Table 6 Schedule icons

Icon	Label	Description
	Execute	Check connectivity on this day.
	Skip	Do not check connectivity on this day.

10. Click **Next**.
The **Specify the Connectivity Options** page appears.
11. Select the success criteria for the action:
 - To specify that the connectivity check is successful only if the connectivity test is successful for all clients in the assigned group, select the **Succeed only after all clients succeed** checkbox.
 - To specify that the connectivity check is successful if the connectivity test is successful for one or more clients in the assigned group, clear the checkbox.
12. Click **Next**.
The **Specify the Advanced Options** page appears.
13. (Optional) Configure advanced options and schedule overrides.

Note

Although the **Retries**, **Retry Delay**, **Inactivity Timeout**, or the **Send Notification** options appear, the Check Connectivity action does not support these options and ignores the values.

14. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.
15. From the **Failure Impact** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.

Note

The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.

- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

Note

If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

16. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.

17. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.

18. (Optional) In **Start Time** specify the time to start the action.

Use the spin boxes to set the hour and minute values, and select one of the following options from the drop-down list:

- **Disabled**—Do not enforce an action start time. The action will start at the time defined by the workflow.
- **Absolute**—Start the action at the time specified by the values in the spin boxes.
- **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.

19. (Optional) Configure overrides for the task that is scheduled on a specific day.

To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:

- Select the day in the calendar, which changes the action task for the specific day.
- Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.
 - To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

Note

- You can edit or add the rules in the **Override** field.
- To remove an override, delete the entry from the **Override** field.

20. Click **Next**.

The **Action Configuration Summary** page appears.

- Review the settings that you specified for the action, and then click **Configure**.

After you finish

(Optional) Create one of the following actions to automatically occur after the check connectivity action:

- Probe
- Traditional backup

Note

This option is not available for NAS snapshot backups.

- Snapshot backup

Creating a snapshot backup action

A snapshot backup action performs a snapshot on a supported storage device, and then generates a save set entry for the snapshot-based backup in the NetWorker media database.

Before you begin

- Create the policy and workflow that contain the action.
- (Optional) Create actions to precede the snapshot backup action. Supported actions that can precede a snapshot backup include:
 - Probe
 - Check connectivity

Procedure

- In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

- From the **Action Type** list, select **Backup**.
- From the secondary action list, select **Snapshot**.
- If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
- Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
- Specify a weekly or monthly schedule for the action:
 - To specify a schedule for each day of the week, select **Weekly by day**.
 - To specify a schedule for each day of the month, select **Monthly by day**.
- Specify the type of backup to perform on each day:



- To specify a level on a specific day, click the backup type icon on the day.
- To specify the same type of backup on each day, select the backup type from the list, and then click **Make All**.

Note

The schedule for a snapshot backup or discovery defines the days of the week or month on which to perform the snapshot backup or discovery. For a snapshot backup action, the schedule also defines the level of backup to perform on each day. This level also applies to the clone action, if created.

The following table provides details on the backup type that each icon represents.

Table 7 Backup type icons

Icon	Label	Description
	Full	Perform a full backup on this day. Full backups include all files, regardless of whether the files changed.
	Skip	Do not perform a backup on this day.

8. Click **Next**.

The **Snapshot Options** page appears.

9. From the **Destination Storage Node** list box, select the storage node with the devices on which to store the backup data.
10. From the **Destination Pool** list box, select the media pool in which to store the backup data.
11. From the **Retention** list box, specify the amount of time to retain the backup data.

After the retention period expires, the save set is removed from the media database and the snapshot is deleted.

12. From the **Minimum Retention Time** list box, specify the minimum amount of time to retain the backup data.

After the specified amount of time, an in-progress snapshot action can remove the snapshot from the storage device to ensure that there sufficient disk space is available for the new snapshot.

13. Click **Next**.

The **Specify the Advanced Options** page appears.

14. In the **Retries** field, specify the number of times that NetWorker should retry a failed probe or backup action, before NetWorker considers the action as failed. When the **Retries** value is 0, NetWorker does not retry a failed probe or backup action.

Note

The **Retries** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option for other actions, NetWorker ignores the values.

15. In the **Retry Delay** field, specify a delay in seconds to wait before retrying a failed probe or backup action. When the **Retry Delay** value is 0, NetWorker retries the failed probe or backup action immediately.
-

Note

The **Retry Delay** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. When you specify a value for this option in other actions, NetWorker ignores the values.

16. In the **Inactivity Timeout** field, specify the maximum number of minutes that a job run by an action can try to respond to the server.

If the job does not respond within the specified time, the server considers the job a failure and NetWorker retries the job immediately to ensure that no time is lost due to failures.

Increase the timeout value if a backup consistently stops due to inactivity. Inactivity might occur for backups of large save sets, backups of save sets with large sparse files, and incremental backups of many small static files.

Note

The **Inactivity Timeout** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option in other actions, NetWorker ignores the value.

17. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.
18. From the **Failure Impact** list, specify what to do when a job fails:

- To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.
-

Note

The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.

- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.
-

Note

If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

19. From the **Send Notifications** list box, select whether to send notifications for the action:
 - To use the notification configuration that is defined in the Policy resource to send the notification, select **Set at policy level**.
 - To send a notification on completion of the action, select **On Completion**.
 - To send a notification only if the action fails to complete, select **On Failure**.
20. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:


```
nsrlog -f policy_notifications.log
```
- On Linux, to send an email notification, type the following command:


```
mail -s subject recipient
```
- On Window, to send a notification email, type the following command:


```
smtpmail -s subject -h mailserver recipient1@mailserver recipient2@mailserver...
```

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
 - `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.
 - `recipient1@mailserver`—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.
21. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
 22. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
 23. Click **Next**.
The **Action Configuration Summary** page appears.
 24. Review the settings that you specified for the action, and then click **Configure**.

Creating a discover action

A discover action discovers snapshots on a NAS device that the current NetWorker server did not create. NetWorker then creates save set entries for the discovered snapshots on a NAS client in the media database. The discover action also compares the save sets in the media database for the NAS client against the actual snapshots stored on the client. If the media database entry does not have a corresponding snapshot on the NAS client, then the media database entry is removed.

Before you begin

Create the policy and workflow that contain the action. Ensure that the group assigned to the workflow contains only NAS devices. The discover action should be the first action in the workflow.

Procedure



1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. From the **Action Type** list, select **Discover**.
3. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
4. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
5. Specify a weekly or monthly schedule for the action:
 - To specify a schedule for each day of the week, select **Weekly by day**.
 - To specify a schedule for each day of the month, select **Monthly by day**.
6. To specify whether to perform snapshot discovery, click the icon on each day.

The following table provides details on the icons.

Table 8 Snapshot discovery icons

Icon	Label	Description
	Execute	Perform snapshot discovery on this day.
	Skip	Do not perform snapshot discovery on this day.

To perform snapshot discovery every day, select **Execute** from the list and click **Make All**.

7. Click **Next**.

The **Discover Options** page appears.

8. From the **Discover Type** list, select the NAS snapshot data to discover.
9. Click **Next**.

The **Specify the Advanced Options** page appears.

After you finish

(Optional) Create a generate index action to create entries in the client file indexes for the snapshot save sets, or a clone action to automatically clone the save sets after the discover action.

Creating a generate index action

A generate index action creates client file indexes of existing NAS snapshots.

Before you begin

- Create the policy and workflow that contains the action.
- (Optional) Create actions to precede the generate index action in the workflow. Supported actions that can precede a generate index action include:
 - Discover
 - Snapshot backup

Note

This point is only optional if you are using a saveset ID group. If you are using a NAS Device group, or Client group, one of these two actions must be used for the index action.



Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.
2. From the **Action Type** list, select **Generate Index**.
3. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
4. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
5. Specify a weekly or monthly schedule for the action:
 - To specify a schedule for each day of the week, select **Weekly by day**.
 - To specify a schedule for each day of the month, select **Monthly by day**.
6. To specify whether to perform snapshot index generation, click the icon on each day.

The following table provides details on the icons.

Table 9 Snapshot generate icons

Icon	Label	Description
	Execute	Perform snapshot index generation on this day.
	Skip	Do not perform snapshot index generation on this day.

To perform snapshot index generation every day, select **Execute** from the list and click **Make All**.

- Click **Next**.

The **Index Options** page appears.

- From the **Type of saveset to be indexed** list, select the NAS snapshot to index.
- Click **Next**.

The **Specify the Advanced Options** page appears.

After you finish

(Optional) Create a clone action to automatically clone the data after the generate index action. A clone action is the only supported action after a generate index action in a workflow.

Creating a clone action

A clone action creates a copy of one or more save sets. Cloning allows for secure offsite storage, the transfer of data from one location to another, and the verification of backups.

Procedure

- In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:
 - If the action is the first action in the workflow, select **Create a new action**.
 - If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

- In the **Name** field, type the name of the action.

The maximum number of characters for the action name is 64.

- Legal Characters: _ : - + = # , . % @
- Illegal Characters: /*?[]()\$!^;"'` ~><&|{ }

- In the **Comment** field, type a description for the action.
- To ensure that the action runs when the policy or workflow that contains the action is started, in the **Enabled** box, select the option. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.



Note

When you clear the **Enabled** option, actions that occurs after a disabled action do not start, even if the subsequent options are enabled.

5. From the **Action Type** list, select **Clone**.
6. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
7. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
8. Specify a weekly or monthly schedule for the action:
 - To specify a schedule for each day of the week, select **Weekly by day**.
 - To specify a schedule for each day of the month, select **Monthly by day**.
9. Specify the days to perform cloning:
 - To clone on a specific day, click the **Execute** icon on the day.
 - To skip a clone on a specific day, click the **Skip** icon on the day.
 - To check connectivity every day, select **Execute** from the list, and then click **Make All**.

The following table provides details on the icons.

Table 10 Schedule icons

Icon	Label	Description
	Execute	Perform cloning on this day.
	Skip	Do not perform cloning on this day.

10. Click **Next**.
The **Specify the Clone Options** page appears.
11. In the **Data Movement** section, define the volumes and devices to which NetWorker sends the cloned data:
 - a. From the **Destination Storage Node** list, select the storage node with the devices on which to store the cloned save sets.
 - b. In the **Delete source save sets after clone completes** box, select the option to instruct NetWorker to move the data from the source volume to the destination volume after clone operation completes. This is equivalent to staging the save sets.
 - c. From the **Destination Pool** list, select the target media pool for the cloned save sets.
 - d. From the **Retention** list, specify the amount of time to retain the cloned save sets.

After the retention period expires, the save sets are marked as recyclable during an expiration server maintenance task.

12. In the **Filters** section, define the criteria that NetWorker uses to create the list of eligible save sets to clone. The eligible save sets must match the requirements that are defined in each filter. NetWorker provides the following filter options:
 - a. **Time filter**—In the **Time** section, specify the time range in which NetWorker searches for eligible save sets to clone in the media database. Use the spin boxes to specify the start time and the end time. The **Time** filter list includes the following options to define how NetWorker determines save set eligibility, based on the time criteria:
 - **Do Not Filter**—NetWorker inspects the save sets in the media database to create a clone save set list that meets the time filter criteria.
 - **Accept**—The clone save set list includes save sets that are saved within the time range and meet all the other defined filter criteria.
 - **Reject**—The clone save set list does not include save sets that are saved within the time range and meet all the other defined filter criteria.
 - b. **Save Set filter**—In the **Save Set** section, specify whether to include or exclude ProtectPoint and Snapshot save sets, when NetWorker searches for eligible save sets to clone in the media database. The **Save Set** filter list includes the following options to define how NetWorker determines save set eligibility, based on the save set filter criteria:
 - **Do Not Filter**—NetWorker inspects the save sets in the media database to create a clone save set list that meets the save set filter criteria.
 - **Accept**—The clone save set list includes eligible ProtectPoint save sets or Snapshot save sets, when you also enable the ProtectPoint checkbox or Snapshot checkbox.
 - **Reject**—The clone save set list does not include eligible ProtectPoint save sets and Snapshot save sets when you also enable the ProtectPoint checkbox or Snapshot checkbox.

Note

For NAS device, only Snapshot save set is applicable.

- c. **Clients filter**—In the **Client** section, specify a list of clients to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The **Client** filter list includes the following options, which define how NetWorker determines save set eligibility, based on the client filter criteria:
 - **Do Not Filter**—NetWorker inspects the save sets that are associated with the clients in the media database, to create a clone save set list that meets the client filter criteria.
 - **Accept**—The clone save set list includes eligible save sets for the selected clients.
 - **Reject**—The clone save set list does not include eligible save sets for the selected clients.
- d. **Levels filter**—In the **Levels** section, specify a list of backup levels to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The **Levels** filter list includes the following options to define how NetWorker determines save set eligibility, based on the level filter criteria:

- **Do Not Filter**—NetWorker inspects the save sets regardless of the level in the media database, to create a clone save set list that meets all the level filter criteria.
- **Accept**—The clone save set list includes eligible save sets with the selected backup levels.
- **Reject**—The clone save set list does not include eligible save sets with the selected backup levels.

Note

For NAS device, only full backup level is applicable.

13. Click **Next**.

The **Specify the Advanced Options** page appears.

14. Configure advanced options, including notifications and schedule overrides.
-

Note

Although the **Retries**, **Retry Delay**, or the **Inactivity Timeout** options appear, the clone action does not support these options and ignores the values.

15. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.
16. From the **Failure Impact** list, specify what to do when a job fails:
- To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.

Note

The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.

- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.
-

Note

If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

17. From the **Send Notifications** list box, select whether to send notifications for the action:
- To use the notification configuration that is defined in the Policy resource to send the notification, select **Set at policy level**.
 - To send a notification on completion of the action, select **On Completion**.
 - To send a notification only if the action fails to complete, select **On Failure**.
18. In the **Send notification** attribute, when you select the **On Completion** option or **On failure** option, the **Command** box appears. Use this box to configure how

NetWorker sends the notifications. You can use the `nsrlog` command to send the notifications to a log file or you can send an email notification.

The default notification action is to send the information to the `policy_notifications.log` file. By default, the `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and in the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows.

Use the default mailer program on Linux to send email messages or the `smtpmail` application on Windows:

- To send notifications to a file, type the following command, where `policy_notifications.log` is the name of the file:

```
nsrlog -f policy_notifications.log
```

- On Linux, to send an email notification, type the following command:

```
mail -s subject recipient
```

- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

```
/usr/sbin/sendmail -v recipient_email "subject_text"
```

- On Windows, to send a notification email, type the following command:

```
smtpmail -s subject -h mailserver recipient1@mailserver  
recipient2@mailserver...
```

where:

- `-s subject`—Includes a standard email header with the message and specifies the subject text for that header. Without this option, the `smtpmail` program assumes that the message contains a correctly formatted email header and nothing is added.
- `-h mailserver`—Specifies the hostname of the mail server to use to relay the SMTP email message.
- `recipient1@mailserver`—Is the email address of the recipient of the notification. Multiple email recipients are separated by a space.

19. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
20. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
21. (Optional) In the **Start Time** option, specify the time to start the action.

Use the spin boxes to set the hour and minute values, and select one of the following options from the list box:

- **Disabled**—Do not enforce an action start time. The action will start at the time defined by the workflow.
- **Absolute**—Start the action at the time specified by the values in the spin boxes.
- **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.

22. (Optional) Configure overrides for the task that is scheduled on a specific day.

To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:

- Select the day in the calendar, which changes the action task for the specific day.
- Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.
 - To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

Note

- You can edit or add the rules in the **Override** field.
 - To remove an override, delete the entry from the **Override** field.
-

23. Click **Next**.

The **Action Configuration Summary** page appears.

24. Review the settings that you specified for the action, and then click **Configure**.

After you finish

(Optional) Create a clone action to automatically clone the save sets again after this clone action. Another clone action is the only supported action after a clone action in a workflow.

Visual representation of NAS snapshot workflows

After you create actions for a workflow, in the Administration interface, you can see a map provides a visual representation of the actions on the right side of the **Protection** window.

Figure 11 Sample NAS snapshot workflow



The oval icon specifies the group to which the workflow applies. The rounded rectangle icons identify actions. The parallelogram icons identify the destination pool for the action.

You can work directly in the visual representation of a workflow to perform the following tasks:

- You can adjust the display of the visual representation by right-clicking and selecting one of the following options:
 - **Zoom In**—Increase the size of the visual representation.

- **Zoom Out**—Decrease the size of the visual representation.
 - **Zoom Area**—Limit the display to a single section of the visual representation.
 - **Fit Content**—Fit the visual representation to the window area.
 - **Reset**—Reset the visual representation to the default settings.
 - **Overview**—View a separate dialog box with a high-level view of the visual representation and a legend of the icons.
- You can view and edit the properties for the group, action, or destination pool by right-clicking the icon for the item, and then select **Properties**.
 - You can create a group, action, or destination pool by right-clicking the icon for the item, and then select **New**.

CHAPTER 3

Software Configuration

This chapter includes the following topics:

- [Backup group resource migration](#).....68
- [Road map for NAS snapshot configuration](#).....70
- [NetWorker with NAS configuration prerequisites](#).....70
- [Configuring a Client resource with the wizard](#).....73
- [Configuring a NAS device client resource manually](#).....77
- [Configuring NSM with NetApp NAS snapshots](#).....78
- [Configuring the Application Information variables](#).....79
- [Configuring preprocessing and postprocessing scripts](#).....79

Backup group resource migration

During the migration process, NetWorker creates resources to replace each Group resource, and then migrates the Group configuration attributes from the 8.2.x and earlier resources to the new NetWorker 19.1 resources.

Resource migration for group resources when Snapshot is enabled

The following table summarizes the Group attribute values that migrate to NetWorker 19.1 resources attributes, when the group is Snapshot enabled.

Table 11 Migration of Group attributes

19.1 Resource type	19.1 Resource name	Migration process overview	Attribute values that are migrated from Group resource
Policy	Backup	One policy resource that is called Backup appears and contains all migrated information for all NetWorker group resources that back up file system.	Not applicable
Protection Group	Snapshot	One protection group resource appears for all Snapshot policies.	
Protection Group	Name of the Group resource	One Protection Group resource appears for each migrated Group resource. Each Protection Group contains the same client resources that were associated with the pre-19.1 group resource.	Comment
Workflow	Name of the Group resource	One Workflow resource appears for each migrated Group resource. Each Workflow resource is associated with the Protection Group resource that was created for the migrated Group resource.	<ul style="list-style-type: none"> • Autostart • Start Time • Next Start • Interval • Restart Window • End Time attribute value is set to Start Time+ (Interval*(n-1)) <p>If the Probe backup group attribute was enabled, the following values are migrated:</p> <ul style="list-style-type: none"> • Probe Interval—To the Interval attribute • Probe Start Time—To the Start Time attribute

Table 11 Migration of Group attributes (continued)

19.1 Resource type	19.1 Resource name	Migration process overview	Attribute values that are migrated from Group resource
			<ul style="list-style-type: none"> • Probe End Time—To the End Time attribute
Probe	Probe	The Probe action resource appears when the Probe based group attribute was enabled in the pre-19.1 migrated group.	Not applicable
Action— Snapshot backup	Backup	The Snapshot Backup action appears for a Group resource that has the Snapshot attribute enabled.	<ul style="list-style-type: none"> • Parallelism • Retries • Retry delay • Success Threshold • Option attributes: <ul style="list-style-type: none"> ▪ No save, Verbose, Estimate, Verify Synthetic Full, Revert to full when Synthetic Full fails • Schedule • Schedule Time • Retention policy • Inactivity Timeout • Soft Runtime Limit—To Soft Limit • Hard Runtime Limit—To Hard Limit • File Inactivity Threshold —To Inactivity Threshold • File Inactivity Alert Threshold—To Inactivity Alert Threshold • Min expiration = (1440/ (backups per day/retain count))-10 • If Retain snapshot=0, then Backup snapshots attribute is set to ALL
Action— Clone	Clone	The Clone action resource appears when the Clone attribute was enabled in the Group resource.	Clone Pool—To the Destination Pool attribute

Note

The *NetWorker Updating from a Previous Release Guide* provides details about resources that are migrated during the update process.

Considerations about migration

1. A clone action is created for every backup action. The clone action moves snapshot data to media.
2. A NetWorker 8.2.x snapshot policy resource that was configured for **nth snapshot** is not migrated to NetWorker 19.1, and a clone action is not created. For example, for a NetWorker 8.2.2 snapshot policy of 1-1-day-All resolves to take 1 snapshot, retain 1, expire snapshots every day, and rollover the data to the rollover device. After the migration a clone action is created to backup every snapshot. This means the rollover action in 8.2.x translates to a clone action in 9.2.x and later.
3. Each backup action within snapshot policy is followed by a clone action. Clone actions have filtering options. Check whether filtering meets the backup requirements of NetWorker Snapshot Management.

Road map for NAS snapshot configuration

Plan the network-attached storage (NAS) snapshot configuration by using a high-level road map that outlines the sequence of configuration tasks that you must perform.

Procedure

1. Verify the configuration prerequisites.
[NetWorker with NAS configuration prerequisites](#) on page 70 provides details.
2. Configure the snapshot and the clone operations by using the Data Protection Action Wizard or the CLI.
The following topics provide details:
 - [Configuring a Client resource with the wizard](#) on page 73
 - [Configuring a NAS device client resource manually](#) on page 77
3. Configure any necessary Application Information variables.
[Configuring the Application Information variables](#) on page 79 provides details.
4. Configure any necessary preprocessing or postprocessing scripts.
[Configuring preprocessing and postprocessing scripts](#) on page 79 provides details.
5. Test the configuration.

NetWorker with NAS configuration prerequisites

Verify the basic compatibility of all systems that are used for NetWorker operations. [Components of the NAS snapshot network](#) on page 22 provides details.

The following sections describe the prerequisites for the hosts that are involved in NetWorker operations.

NAS device prerequisites

Ensure that the network-attached storage (NAS) device is installed with the prerequisite software for snapshot and clone operations:

- Isilon—None required.
- VNX or VNX2—None required.
- VNXe—VNXe Unisphere CLI, available at <https://support.emc.com>.
- Unity—Unisphere CLI, available at <https://support.emc.com>.
- NetApp—NetApp Manageability SDK, available from NetApp.
 - On Microsoft Windows systems, which are installed in a directory in the system-level PATH variable.
For example, `C:\Program Files\EMC Networker\nsr\bin`
 - On Linux systems, which are installed in the `/usr/lib64` directory.
 - On NetApp, ensure the `options.httpd.admin.enable` setting is set to ON.
 - (Optional) for HTTPS access on NetApp, the `options.httpd.admin.ssl.enable` setting is set to ON.
 - You can launch the NetApp GUI from the NetApp Systems Manager. You can access the NetApp Systems Manager from here: <http://mysupport.netapp.com/NOW/cgi-bin/software>

Ensure that you have synchronized the system times on the NAS devices with the NetWorker server. If the clocks are not synchronized within 10 seconds, the snapshot and replication operations can cause problems or fail. Use a common Network Time Protocol (NTP) server to obtain system time, to prevent time drift.

The *NetWorker E-LAB Navigator* for the NAS device that you use, available on the Support site, provides requirement details on systems and software.

Replication prerequisites

For replication operations, ensure that the following prerequisites are met:

- For remotely replicated directories, ensure the remote system has a Client resource set up on the NetWorker server.
- Ensure you have correctly configured the replication policy on the NAS device before you create the Client resource for the NAS device.
- Ensure you have configured the replication policy to create a snapshot when the replication operation is completed.

The NAS device documentation provides details.

NetApp prerequisites

For SnapVault and SnapMirror replication operations, ensure that the following prerequisites are met:

- The source and destination volumes or qtrees must exist before you configure a replication. The volumes may be on the same or on separate NetApp devices on the same or on separate servers that may be NetApp virtual machines or Vserver. Both volumes must be in an online state.
- The source volume must be read-writable. The destination volume must be a data protection type.

- Verify that the SnapVault and SnapMirror replication policies are configured on the NetApp devices.
- Validate the replication configuration by performing an initial replication operation. This operation makes the replication policies available for selection in the NetWorker client wizard.
- In NMC under the Server area, ensure that you assign all privileges under **User Groups**.
- The NetWorker server must be a part of the NetApp export policy for the share to be accessible.
- The NAS shares must be mounted to the NetWorker server before performing a backup.

The NetApp documentation provides details.

Note

Qtree to qtree Remote Replication is not supported.

Storage node prerequisite

If you plan to clone the snapshot save sets to conventional storage media, then a NetWorker storage node is required for recovery operations.

Ensure the following prerequisites are met:

- NetWorker storage node software is installed.
- Network Data Management Protocol (NDMP) backup storage devices are configured on the storage node for the clone operations.

Configuring NDMP

If you plan to clone the snapshot save sets to conventional storage media, complete the NDMP configuration on the NAS client system. The *NetWorker Network Data Management Protocol (NDMP) User Guide* provides information about NDMP configuration.

Note

For the client creation to be successful, you must preconfigure the NDMP username and password for the NAS devices.

Note

The *Isilon OneFS Administration Guide* provides updates and details about Isilon commands. You can also perform this NDMP configuration through the Isilon OneFS Storage Administration user interface.

The following steps are for Isilon version 8. If you are using an earlier version of Isilon, <https://support.emc.com/kb/471904> provides more information.

Procedure

1. To connect to a node in the cluster, use `ssh`.
2. To create the NDMP username and password, use the `isi` command.

```
isi ndmp users create <name> --password <string>
```


For example, the following command creates an NDMP user account with username `ndmp_user` and password `1234`:

```
isi ndmp users create ndmp_user --password=1234
```

3. To enable NDMP, use the `isi` command.

For example: `isi ndmp settings global modify --service=yes`

4. To configure an NDMP backup, type the NDMP settings `isi` command.

For example, the following command configures OneFS to interact with NetWorker:

```
isi ndmp settings global modify --dma=emc
```

Configuring a Client resource with the wizard

The NetWorker Management Console (NMC) configures network-attached storage (NAS) devices as NetWorker clients. Network Data Management Protocol (NDMP) enables NAS devices to mimic NetWorker clients for configuration purposes without NetWorker client software installed. The wizard is the recommended configuration method to create or modify the NAS device configuration.

Before you begin

Verify that the necessary prerequisites are met. [NetWorker with NAS configuration prerequisites](#) on page 70 provides details.

Procedure

1. Run NMC.
2. In the **Enterprise** view, browse tree, right-click the NetWorker server that will manage the snapshots, and then select **Launch Application**.
3. Start the client wizard to create a NAS device from one of the following areas:
 - The **Protection** view
 - The **Devices** view
4. In the **Protection** view, launch the client wizard:
 - To create a NAS device configuration, right-click **Client** in the navigation tree, and then select **New Client Wizard**.
 - To modify a configuration, right-click an existing NAS device in the Client table, and then select **Modify Client Wizard** or **Modify Client Properties**.
5. In the wizard, complete the **Specify the Client Information** page:
 - a. In the **Client Name** field, type an associated hostname.

Note

For Isilon, use a non-aggregated (nonteam) network connection to one of the Isilon cluster nodes that has a fixed IP address. Do not use the SmartConnect hostname.

- b. In the **Client Type** field, select **NAS-Device**.
 - c. Click **Next**.
6. Complete the **Specify Credentials for the NAS device** page:

- a. For the user who is authorized to access data on this client instance of the NAS device, for example, to create, delete, or share snapshots, in the **Management Credentials** fields, type a **User Name** and **Password**. Each instance can have a different user. You may need to provide a management name or address such as a VNX Control Station name if it differs from the NAS device name. For NetApp, you must provide the management name of the NetApp storage virtual machine (SVM), and not the NetApp filer management name.
 - b. For the user who is authorized to perform NDMP operations on the NAS device, such as browsing the file system during configuration, in the **NDMP Credentials** fields, type a **User Name** and **Password**.
 - c. If a Microsoft Windows client will access the snapshot data, complete the **File Server Credentials**. These fields are ignored by other operating system types. Type a **User Name** and **Password** for the user who is authorized to mount and access the snapshot data on a Windows recovery host for this instance of the NAS device. Each instance can have a different user.
 - d. Click **Next**.
7. Complete the **Verify the NAS Device Information** page:
- a. Review the **Vendor Information** details.
 - b. To configure snapshot clones to NDMP devices, complete the **NDMP Clone Information** details. The available values are those supported by the client's NDMP configuration. The *NetWorker Administration Guide* provides details. Specify the **NDMP backup type**.

NAS device	Supported backup types
VNX	<ul style="list-style-type: none"> • tar—Tape Archive. • dump—The best backup type for this device. An inode-based backup that traverses a file tree in directory first and file-based order. • vbb—Use the volume-based backup type to back up the entire volume at the block level rather than at a file level, which can be more efficient than traditional file-based backups. It does not support DDAR, TBB, and Three Party backups. • ts—Tape silvering.
Isilon	<ul style="list-style-type: none"> • tar—Tape Archive. • dump—An inode-based backup that traverses a file tree in directory first and file-based order.
NetApp	<ul style="list-style-type: none"> • dump—An inode-based backup that traverses a file tree in directory first and file-based order.

In the **Advanced App Info** field, you may specify additional NAS-specific variables that are supported by the client's NDMP configuration, one per line. For example, for an Isilon or NetApp remote replication backup, specify `NDMP_MNT_HOST=remote host`. [Common Application Information variables](#) on page 90 and [Configuring the Application Information variables](#) on page 79 provides more information.

Review the **App Info** options that the wizard has selected for the NAS device. You can disable options, as required.

c. Click **Next**.

8. For NAS devices that do not support file browsing, complete the **Select the File System Objects** page:

a. Type the absolute pathnames of file systems, directories, and individual files to include in the snapshot save set. Add each item on a separate line.

b. Click **Next**.

Note

While configuring the NetWorker client for a NetApp storage system in 7-Mode and C-Mode use the following paths to mount the volume and add the save sets:

Mode	Path for mounting the volume	Path for adding the savesets
7-Mode	/vol/volumename	/vol/volumename
C-Mode	/volumename	/vol/volumename

9. Complete the **Specify Snapshot Location** page:

a. Specify a local location or remote location for the snapshot. The available snapshot locations depend on the replication policies that the storage administrator defined on the NAS device. The available snapshot locations are those capable of replicating all the save set paths that you specified earlier in the wizard:

- **Local snapshot**—The device will not replicate the specified data and NetWorker will create a snapshot of the original data.
- **Locally replicated snapshot**—The device will replicate the specified data to a different location on the same device and NetWorker will create a snapshot from the replicated data.
- **Remotely replicated snapshot**—The device will replicate the specified data to a location on a different NAS device and NetWorker will create a snapshot from the replicated data.

For example, the following replication policies are defined on a NAS Isilon device:

NAS policy name	Local system	Local path	Remote system	Replicated path
Local1	System_A	/ifs/home/folders	Not applicable	/ifs/localrep
Remote1	System_A	/ifs/home/folders	System_B	/ifs/remoterep/1
Remote2	System_A	/ifs/home	System_B	/ifs/remoterep/2

Then the specified snapshot location paths will use the corresponding NAS replication policies.

Snapshot save set location	NAS policies
/ifs	none
/ifs/foo	none
/ifs/home	Remote2
/ifs/home/folders	Local1, Remote1, Remote2
/ifs/home/folders/bar	Local1, Remote1, Remote2

Note

For NetApp devices, if the snapshot destination volume is written to, the destination becomes the new source volume and the relationship switches.

- b. Click **Next**.
10. Complete the **Backup Configuration Summary** page:
 - a. Review the attributes and values that are listed in the summary.
To modify a setting, click **Back** or click the link in the steps panel.
 - b. To accept and create the configuration, click **Create**.
11. Complete the **Client Configuration Results** page:
 - a. Ensure that the backup configuration is successfully completed.
 - b. Click **Finish**.

After you finish

The Data Protection Policies chapter describes how to schedule a snapshot backup and clone of the data. [Common Application Information variables](#) on page 90 provides more details. After the Client Configuration wizard creates the NAS client, you can modify the properties of the new NAS client.

Performing post Client Configuration wizard steps

After the Client Configuration wizard creates the NAS client, modify the properties of the new NAS client.

Procedure

1. On the **Administrator** window, in the navigation pane, select **Clients**.
2. Right-click the client that you created, and then select **Modify Client Properties**.
3. On the **Globals (2 of 2)** tab, in the **Remote Access** field, specify the remote access credentials for the NetWorker server as follows:
 - For a Windows NetWorker server—`system@networker-server-name`
 - For a UNIX NetWorker server—`root@networker-server-name`

Where *networker-server-name* is the hostname of the NetWorker server.

4. Click **OK**.

Configuring a NAS device client resource manually

In some situations, you might want to configure a network-attached storage (NAS) device manually instead of using the NMC Client Configuration wizard. For NetWorker purposes, NAS devices are configured as NetWorker clients.

Before you begin

Ensure that the prerequisites are met. [NetWorker with NAS configuration prerequisites](#) on page 70 provides details.

You can manually modify a NAS device configuration by using the NMC property windows. This method enables you to configure special directives or available options that the wizard interface does not support, such as [Application Information Variables](#) on page 89.

The replication policies, which are required for locally replicated and remotely replicated snapshots can be configured in the respective NAS devices.

The *NetWorker Administration Guide* provides additional Client resource details that are not specific to NAS device configuration.

Procedure

1. Run NMC, and in the NMC **Enterprise** view, browse tree, right-click the NetWorker server that manages the snapshots, and then select **Launch Application**.
2. In the NetWorker server's **Protection** view, in the **Clients** table right-click the NAS device, and then select **Modify Client Properties**.
3. On the **General** tab, in the **Save set** field, type or browse, and then select all the file systems, directories, or individual files that you want to include in the snapshot. Add each item on a separate line.
4. On the **Apps & Modules** tab, specify the following settings:
 - a. To clone the snapshot, provide NDMP credentials in the **Remote User** and **Password** fields.
 - b. Select **NAS device**.
 - c. Select **NDMP**.
 - d. In the **Application information** field, type any optional NDMP variables that you want to use in the clone.

Note

NetApp configurations must specify the NSR_MOUNTPOINT_NAME variable to enable data recovery from NetApp replications. NetApp configurations must also specify the NDMP_MNT_HOST variable (for example, NDMP_MNT_HOST=*remote system name*).

Variables must be uppercase. NetWorker does not validate the variables. [Application Information Variables](#) on page 89 and [Configuring the Application Information variables](#) on page 79 provides details.

5. On the **Snapshot Management** tab, specify the following settings:

- a. In the **Management credentials** attribute, type the name and password for the user who is authorized to access data (for example, to create, delete, or share data for this instance of the NAS device). Each instance can have a different user.
 - b. In the **NAS device management name** field, type the fully qualified domain name (FQDN), which is required by some devices that use a special address for data management (for example, the VNX or VNXe Control Station address, which is different than the NAS file server address). For NetApp, you must provide the management name of the NetApp storage virtual machine (SVM), and not the NetAPP filer management name.
 - c. If a Microsoft Windows client accesses the snapshot data, complete the **NAS file access user** fields. These fields are ignored by other operating system types. For the user who is authorized to mount and access the snapshot data on a Windows recovery host for this instance of the NAS device, type a **User Name** and **Password**. Each instance can have a different user.
6. Click **OK**.

Configuring NSM with NetApp NAS snapshots

Use the following workaround procedures to successfully configure NSM with NetApp ONTAP.

Configuring NSM with NetApp ONTAP 7-mode

Configure the following pre-requisites on the NetApp filer.

Procedure

1. Create volumes in the NetApp filer with UNIX and NTFS default permissions.
2. Configure CIFS and NFS protocols on the filer.
3. Mount the volumes on the NetWorker server.
4. Configure NDMP credentials.

Configuring NSM with NetApp ONTAP Cluster Mode

Configure the following pre-requisites on the NetApp filer.

Procedure

1. Create the Vservers with separate credentials. Ensure that the user is assigned the "vsadmin" role.
2. Ensure that the user has ONTAP API SSH and HTTPS access.
3. Configure NDMP credentials.
4. Map the domain controller to the respective Vservers.
5. Ensure that the CIFS and NFS protocol services are up and running.
6. Set up the export policies and name spaces in the Vserver.
7. Create the backup volumes, and then map the backup volumes to the respective export policies.
8. Ensure that the NetWorker server is added to the Vserver export policy.
9. In case of CIFS, create the CIFS share.

10. Mount the volumes on the NetWorker server using the following command:

```
mount <vservername>:/volname /mountpoint
```

11. Complete the following steps from NMC:
 - a. While configuring the client resource, add the Vserver name or IP address.
 - b. Add the Vserver username and password under Management credentials.
 - c. Add the Vserver name in the Management Name field.
 - d. Add the saveset name in the following format:

```
/vol/<volumename>
```

Configuring the Application Information variables

Special Application Information variables provide specific control of snapshot processes. The Client Configuration Wizard cannot configure these variables. Manually configure these variables in the NetWorker Client resource for the NAS device by using the NetWorker property windows. [Application Information Variables](#) on page 89 provides details.

Note

For NetApp recovery operations, configure the NSR_MOUNTPOINT_NAME variable. For SnapMirror and SnapVault the NDMP clone application variables must be set to N as follows:

```
DIRECT=N
```

```
UPDATE=N
```

Configuring preprocessing and postprocessing scripts

You can run user-defined preprocessing and postprocessing scripts from the application client. You can run these scripts only for file system backups.

Use preprocessing scripts and postprocessing scripts for operations such as application quiescing, shutdown, or startup. The scripts can produce output such as log files, but the scripts must return an exit status of 0, which means that the script did not fail and the backup can run. Any other exit code for a preprocessing script causes the backup to fail.

Note

Technical Support does not support user-defined scripts.

Procedure

1. Provide the script files with the following security:
 - On Microsoft Windows systems, provide the script files with security that grants full control only to the local SYSTEM, local Administrators, or Backup Operators groups. Otherwise, the scripts do not run.

To set this security, in Windows Explorer right-click the script file, select **Properties**, and then in the **Properties** window **Security** tab, click **Advanced**.

- On UNIX systems, the root user must own the script files. The scripts can set only owner access permissions, and the scripts must at least have run access. Otherwise, the scripts do not run. The parent directory of the scripts must have at least owner run permissions, and must not have write permissions for the group and world.
2. To add, modify, or run the resident scripts, place the scripts in a directory where a user must have administrator/root privileges. Otherwise, any backups that use the scripts fail.

On Microsoft Windows systems, NetWorker searches for relative pathnames in the `NetWorker_install_path/bin` directory.

3. In the property window of the application Client resource, include the pathnames of user-defined scripts in the **Application Information** attribute by using the following variables:
 - `NSR_PRE_SNAPSHOT_SCRIPT=pre-mirror_split-script_path`
 - `NSR_POST_SNAPSHOT_SCRIPT=post-mirror_split-script_path`
4. After a backup is completed, verify the log files that are generated in the `/nsr/logs` (UNIX) directory on the application client host. The log file name is in the form of `script_name_LOGFILE.txt`. The script output appears in the log file.

CHAPTER 4

Data Management and Recovery

This chapter includes the following topics:

- [Snapshot management](#).....82
- [NAS snapshot recovery support and limitations](#)..... 83
- [Recovering data from a snapshot with the Recovery Wizard](#)..... 84

Snapshot management

The management and the recovery operations for NAS device file system data are described in this chapter.

The NMC Recovery wizard and the NetWorker CLI commands provide features that enable you to browse, delete, change snapshot expiration, and recover snapshot data.

Snapshot reporting

You can create individual reports of NAS snapshot, replicate, and clone activities.

The *NetWorker Administration Guide* provides details.

Procedure

1. In NMC, select the **Reports** view.
2. In the **Legacy Reports** folder, select a summary report or a statement report to view. The **Configure** tab for the selected report type appears in the right panel.
3. On the **Configure** tab, specify the items that you want to include in the report. Select the item parameters. Click the **Remove (<)**, **Remove All (<<)**, or **Add All (>>)** buttons as required. If you do not specify **Save Time** values, the report displays all the available data.

The **View Report** tab shows the resultant report.

Save set IDs and expiration policies

When NetWorker creates a snapshot, it generates a separate save set ID for each snapshot object that is specified in the Client resource.

For example, a single physical snapshot can create save sets for `/home/folders/user1` and `/home/folders/user2` if they both reside on the same NAS device. Each save set has a unique save set ID, even if both save sets belong to the same client and both reside on the same NAS device.

During cloning to conventional storage media, NetWorker assigns a different save set ID to each cloned snapshot object. Each save set has an independent expiration policy, and when one save set expires, you can still use the other save set to perform a recovery.

The retention time of the backup action that created the save set determines the expiration of a snapshot save set. However, you can change the expiration period of a snapshot save set and the change does not affect the browse or expiration periods of an associated clone.

Browsing snapshot and clone save sets

The NetWorker client file index records only the save sets that it clones to conventional storage media. Because NetWorker indexes clones, you can browse the save set files in NMC.

The NetWorker media database contains entries for snapshot save sets. However, unlike clones, NetWorker does not catalog the snapshot save sets in the client file index. To browse snapshot save sets, you must use the NMC Recovery Wizard or the `nsrsnapadmin` command utility. NetWorker mounts the snapshot file system on the mount host, which enables you to browse and select files to recover.

Browsing and recovery on Windows hosts

When performing file-by-file browsing, or performing a directed recovery from the NAS device to a Microsoft Windows host, NetWorker connects to the NAS device by using the File Server credentials in the Client resource for the NAS device.

Microsoft Windows cannot connect to the NAS device by using different credentials than a previously established file service connection. NetWorker does not use the credentials that a previously established file service connection has used.

NAS snapshot recovery support and limitations

The following support and limitations apply to NAS snapshot recovery operations:

- A NAS snapshot recovery is supported in the following user interfaces:
 - NMC Recovery wizard
 - `nsrsnapadmin` command utility
 - `nsrnassnap_recover` command for NAS devices
- You can recover individual files or complete file systems from snapshot save sets.
- You cannot combine individual files from multiple save sets in a single recovery session.
- You can recover data from snapshots that NetWorker has cloned to conventional storage media as you would for any conventional NetWorker backup.
- You can recover data from NAS discovered snapshots that NetWorker has cloned to conventional storage media, as you would for any conventional NetWorker backup.

Note

NetApp only supports file-by-file restore and does not support save set restore. In Netapp, restore to different volume is not supported (limitation from NetApp), customers should choose same parent volume (but the directory path can be different).

To restore from snapshot, choose the following paths:

- For in-place restore, choose the original location.
- For out-of-place restore, add `/vol/<parent-volumename> /<directory to restored>` in the destination path.

To restore from clone, choose the following paths:

- For in-place restore, choose `<vservername>/<parent-volumename>` in the destination path.
- For out-of-place restore, add `/<vservername>/<parent-volumename> /<directory to restored>` in the destination path.

Note

[NAS snapshot support and limitation](#) on page 100 provides more information on NAS file system limitations.

Recovering data from a snapshot with the Recovery Wizard

You can use the NMC Recovery Wizard to recover file system data from a snapshot that is stored on a supported NAS device.

Procedure

1. Run NMC, and in the **Enterprise** view, select the NetWorker server name, and then select **Enterprise > Launch Application**.
2. In the NetWorker server **Recover** view, select **Recover > New Recover** to launch the Recovery Wizard.
3. Complete the **Select the Recovery Hosts** page:
 - a. In the **Source Host** field, specify the NAS device that was the original source for the snapshot files or directories that you want to recover.
 - b. In the **Destination Host** field, specify where you want to recover the snapshot files and directories:
 - **Recover to the same host**—The wizard later prompts you to recover to a location on the source NAS device.
 - **Select a destination host**—The wizard later prompts you to direct the recovery to a location on the NetWorker host that is running the recovery.
 - c. In the **Available Recover Types** field, select **NAS Snapshots** or **Filesystem (NDMP)**.
 - d. Click **Next**.
4. Complete the **Select a NAS Snapshot** page. You can recover the entire snapshot save set or you can recover individual directories and individual files from the snapshot save set. Perform the following steps:
 - a. In the **NAS Snapshots** table, select the snapshot to recover from, based on the snapshot time and save set volumes.
 - b. In **Select save set**, select a single save set volume to mount and recover from.

Note

You can select only one save set for this operation. Each additional save set volume requires a separate pass through the wizard.

- c. In **Recover mode**, select from the following options:
 - **Browse and recover save set**—If you select this option, when you click **Next**, the wizard mounts the snapshot volume for the save set and opens the **Select Data to Recover** page. The mount operation can take some time.
 - **Recover full save set**—If you select this option, when you click **Next**, the wizard opens the **Select the Recovery Options** page.
 - **Create temporary mount point for access**—If you want to use an existing mount point for volume access, clear this checkbox. You can set the mount point by using the NSR_MOUNTPOINT_NAME variable. [Using Application Information variables](#) on page 90 provides details.

Note

Clear this setting for NetApp NAS devices on Linux operating systems, which cannot use a temporary mount point. Use an existing mount point.

- d. Click **Next**. The resultant wizard step depends on the selections.
 5. If you selected the **Browse and recover save set** option, complete the **Select Data to Recover** page. Otherwise, NetWorker skips this step.
 - a. Specify the items to recover in the browse tree or type the full path of a location. Mark which directories or files the wizard will recover in the table.
-

Note

The wizard does not list expired save sets. You can recover expired save sets manually by using the `nsrsnapadmin` command utility with the `R` command option or the `nsrnassnap_recover` command. [Using nsrsnapadmin operations](#) on page 94 and the *NetWorker Command Reference Guide* provide details.

- b. Click **Next**.
6. Complete the **Select the Recovery Options** page:
 - a. In the **File Path for Recovery** field, select, browse, or type a location where the wizard will restore the recovered files:
 - **Original path:**
 - In-place recovery—Restore to the original location on the NAS device.
 - **New destination path:**
 - Out-of-place recovery—Restore to a different location on the NAS device.
 - Directed recovery—Restore to a location on the NetWorker host that is running the recovery. A directed recovery may specify a location that is configured as a CIFS or NFS file share, which directs the recovery to a remote location.

Note

In-place recovery and out-of place recovery on the NAS device is not available for the following devices and conditions:

- VNX, VNX2, VNXe, and Unity devices.
 - Isilon devices for a remotely replicated snapshot.
 - NetApp only supports file-by-file restore and does not support saveset restore. To restore from snapshot, choose the following paths:
 - For in-place restore, choose the original location.
 - For out-of-place restore, add `/vol/<volumename> /<directory to restored>` in the destination path.
-

Note

For out of place restore, the directory to be restored has to be created in advance in the NetApp array.

- To restore from clone, choose the following paths:
 - For in-place restore, choose `<vservername>/volumename` in the destination path.
 - For out-of-place restore, add `/<vservername>/<volumename> /<directory to restored>` in the destination path.
-

[Examples of snapshot recovery](#) on page 28 illustrates the recovery options.

NetWorker 19.1 supports an Isilon cluster recovery to another Isilon in a different cluster or the same cluster. For example, host Isilon A has the original data and you want to recover data to host Isilon B and you have taken a snapshot of Isilon A. On the NMC, run a recovery to Isilon B. Begin by first selecting the NetWorker server as the destination host. Then in the **New destination path** field, specify the Isilon B hostname, followed by a colon, followed by the path to recover to (for example, `IsilonB:/my/path`). You must type the destination path in the correct format, `<host>:<path>`. NetWorker does not validate the format.

You can also perform this recovery from the CLI by typing the following command:

```
nsrcnassnap_recover -c <snapshot_source_Isilon_machine> -s
<NetWorker_server> -S <saveset_ID> -i {N|Y|R} -A
NFS_DESTINATION_CLIENT=<destination_Isilon_machine> -A
"NFS_DESTINATION_PATH=<destination_Isilon_path>" -A
NSR_DIRECTED_RECOVER_HOST=<NetWorker_server>
<source_Isilon_restore_path>
```

b. In the **Duplicate file options** field, specify how the wizard resolves conflicts in file names or directory names:

- **Rename the recovered file**—The wizard recovers the file or directory with a new name that NetWorker automatically generates. This option is available for directed recoveries only.
- **Do not recover the file**—The wizard does not recover the file or directory, and skip it.

- **Overwrite the existing file**—The wizard replaces the file contents or the directory contents.
 - **Fail the recover**—The wizard does not recover any further files or directories. This option is also available for directed recoveries.
-

Note

For NetApp devices, select an option, otherwise the recovery fails.

- c. To specify further options, select **Advanced Options** , and then specify the attributes.
-

Note

- For a Remote Replication recovery in NetApp, specify the `-R` command in the extra recovery options.
 - If recovery has to be performed on the directory that doesn't have parent directories, then use `-P` in recovery option. This ensures the ACLs of the parent directories are intact.
-

- d. Click **Next**.

7. Complete the **Perform the Recover** page:

- a. In the **Recovery Name** field, type a name for the recovery.
- b. In the **Recovery Start Time** field, specify the following attributes:
 - The **Start recovery now** option is the only option that snapshot recovery supports.
 - In the **Specify a hard stop time** option, you can specify a time limit that stops an incomplete recovery process.
- c. In the **Recovery Resource Persistence** field, select the retain or delete option for this recovery resource:
 - **Persist this resource until deleted by user**
 - **Automatically remove this resource based on jobs database retention**
- d. Review the **Summary** of the recovery and make any necessary corrections by going to the previous pages in the wizard.
- e. Click **Run Recover**.

The wizard performs the recovery. The *NetWorker Administration Guide* provides more details on the NMC Recovery Wizard.

APPENDIX A

Application Information Variables

This appendix includes the following topic:

- [Using Application Information variables](#).....90

Using Application Information variables

As part of the manual configuration of a NAS device, some NetWorker configurations require the use of special variables that provide specific control of snapshot processes.

To implement these controls, type the variables and their values in the Application Information attribute of the Client resource for the NAS device.

[Configuring a NAS device client resource manually](#) on page 77 provides the manual configuration procedure that can include Application Information variables.

Common Application Information variables

The following table lists the Application Information variables that are common to the NAS devices that NetWorker supports.

Table 12 Common Application Information variables

Common variable	Definition
NASDeviceType	Specifies the snapshot provider. Usually not required due to device type discovery. Valid values are Isilon, NetApp, VNX, and VNXe. If you do not specify a value, NSM tries each of these values in order.
NSR_DEBUG_LEVEL	Specifies the verbosity level of the error message logs. Valid values are 0 to 9. The default value is 0.
NSR_MOUNTPOINT_NAME	This variable is mandatory to restore from NetApp replications. Specifies the name of the mount point to use when NetWorker needs to mount a snapshot on a client.
NSR_POST_SNAPSHOT_SCRIPT	Specifies the pathname of the postprocessing command script. No default value.
NSR_PRE_SNAPSHOT_SCRIPT	Specifies the pathname of the preprocessing command script. No default value.
NSR_PS_RECOVER_FILE_OPT=(N/R/Y)	Specifies options to handle duplicate file behavior on recover. Options are: <ul style="list-style-type: none"> • N (skip, do not overwrite) • R (rename) • Y (overwrite)
NSR_USE_EXISTING_MOUNTPOINT	Specifies the use of an existing mount point instead of creating a temporary mount point. Not available for NetApp on Linux. The default value is No.

Table 12 Common Application Information variables (continued)

Common variable	Definition
NDMP_MNT_HOST	Specifies the name of the remote host when you configure Remote Replication backups on Isilon or NetApp filers.

APPENDIX B

Command-Line Operations for NAS Recovery

This appendix includes the following topics:

- [Using CLI commands for NetWorker operations](#)..... 94
- [Using nsrsnapadmin operations](#)..... 94
- [Querying with the mminfo command](#)..... 97
- [Using nsrnassnap_recover operations](#)..... 98

Using CLI commands for NetWorker operations

This appendix provides a summary of CLI commands and examples for NetWorker operations.

The *NetWorker Command Reference Guide* and the NetWorker man pages provide details on the commands.

NAS snapshot commands

You can use the following commands for managing NetWorker NAS snapshots:

- `nsrnassnap`
- `nsrnassnap_discover`
- `nsrnassnap_recover`
- `nsrnassnap_save`
- `nsrnassnapck`
- `nsrnassnap_index`

The NetWorker man pages and the *NetWorker Command Reference Guide* provide details.

Using nsrsnapadmin operations

You can run the `nsrsnapadmin` command utility in interactive mode to query, recover, delete, and expire file system snapshot save sets.

Note

The `nsrsnapadmin` interactive commands support only snapshots of file systems. The commands do not support the snapshots of application data, such as NMDA or NMSAP data.

To start interactive mode, type `nsrsnapadmin` at the CLI prompt. When you receive an input prompt, you can type a specific command and its available options to perform the NetWorker options that are listed in the following table.

Table 13 Commands and options that are supported in `nsrsnapadmin` interactive mode

NSM operation	Command and available options
Display snapshot save sets	<code>p [-s <i>nsr_server</i>] [-c <i>client</i>] [-v] [<i>path</i> -S <i>ssid</i>]</code>
Delete a snapshot save set	<code>d [-s <i>nsr_server</i>] [-c <i>client</i>] [-v] [-a] [-y] -S <i>ssid</i> [or -S "<i>ssid ssid ...</i>"]</code>
Perform a save set recovery	<code>R [-s <i>nsr_server</i>] [-c <i>client</i>] [-M <i>mount_host</i>] [-v] -S <i>ssid</i> [-t <i>destination</i>] [-T <i>recover_host</i>] -m <i>path</i> [-A <i>attr=val</i>]</code>
Perform a file-by-file browsing and recovery	<code>r [-s <i>nsr_server</i>] [-c <i>client</i>] [-M <i>mount_host</i>] [-T <i>recover_host</i>] -S <i>ssid</i> [-A <i>attr=val</i>]</code>

Table 13 Commands and options that are supported in nsrsnapadmin interactive mode (continued)

NSM operation	Command and available options
Reset the expiration time for a snapshot save set	<code>e time [-s nsr_server] [-c client] [-v] -S ssid [or -S "ssid ssid ..."]</code>
Exit the program	<code>q</code> or <code>quit</code>

where:

- *nsr_server* is the hostname of the NetWorker server.
- *client* is the name of the NAS device.
- *mount_host* is the hostname of the mount host.
- `-v` is for verbose logging.

The *NetWorker Command Reference Guide* and NetWorker man pages provide details.

Example nsrsnapadmin operations

After you start the `nsrsnapadmin` utility in interactive mode, you can type a specific command and its options at the input prompt to perform an NSM operation. The following sections provide examples of the commands and options that are used for specific NSM operations.

Querying snapshot save sets

When you type the `p` command and its options at the `nsrsnapadmin` prompt, the program queries the NetWorker server for snapshot save sets for the client. The program lists specific properties of the snapshot save sets, such as the creation time and the date of each snapshot. For example:

```
p -s server -c client [-v] path
```

where:

- *server* is the hostname of the NetWorker server.
- *client* is the hostname of the NAS device.
- *path* is the pathname of a particular snapshot save set. Type the pathname to query a single save set only. Otherwise, the output message lists all the save sets.

A message similar to the following appears:

```
nsrsnapadmin> p -c isilon-1.myorg.com
  ssid = 643255862 savetime="4/16/2014 11:05:02 AM"
(1397660702)   expiretime="4/17/2015 9:41:59 AM"
(1429278119)   ssname=/ifs/data
```

File-by-file browsing and recovery

When you type the `r` command and its options at the `nsrsnapadmin` prompt, the program lists the file system as it existed at the time of the snapshot backup. Options enable you to browse, select, and recover the elements of the file system.

For example:

```
r -s server -c client -M mount_host -T recover_host -S ssid
```

Where *client* is the name of the NAS device.

Deleting a snapshot save set

Deleting a snapshot save set is similar to deleting a standard NetWorker save set. The `nsrsnapadmin` command deletes the physical snapshot from the NAS device and deletes all save sets that refer to that physical snapshot from the media database.

For example:

```
nsrsnapadmin -d -s server -c client -S ssid
```

where:

- *server* is the hostname of the NetWorker server.
- *client* is the name of the NAS device.
- *ssid* is the snapshot save set ID.

You can also use the `nsrnassnapck` command to delete a snapshot.

Modifying retention periods of a snapshot save set

To modify the expiration date of a snapshot, type the `e` command at the `nsrsnapadmin` prompt. For example:

```
e time -s server -c client -S ssid
```

where:

- *time* is the date and time when the snapshot save set will expire. Acceptable date and time formats are as follows:
 - mm/dd[/yy]
 - month_name dd[/yy]
 - hh[:mm[:ss]] [meridian] [zone]
 - hhmm [meridian] [zone]
- *server* is the hostname of the NetWorker server.
- *client* is the hostname of the NAS device.
- *ssid* is the ID of the snapshot save set that you want to modify.

A message similar to the following appears:

```
Resetting expire time for ssid : 4090300235
```

The message indicates that you have successfully changed the expiration time.

Note

- If you omit the year, the year defaults to the current year.
- If you omit the meridian, NetWorker uses a 24-hour clock.
- If you omit the time zone (for example, GMT), NetWorker uses the current time zone.
- If you specify a date mm/dd/yy as in `-e 09/04/13`, the time defaults to 00:00:00. NetWorker changes the snapshot save set browse and retention times to 09/04/13 00:00:00.
If you specify a time hh:mm:ss as in `-e 20:00:00`, the date defaults to the system time, for example, 09/03/13. NetWorker changes the snapshot save set browse and retention times to 09/03/13 20:00:00.

Querying with the `mminfo` command

You can query a client's snapshot save sets with the `mminfo` command. The `-q snap` option lists all snapshot save sets for a particular client. You can also query for snapshots by typing one of the following commands:

- `mminfo -qsnap -r "queries specification"` (Sample query specification: `-r "ssid, savetime, name"`) Also, `mminfo -c <client name> -r "<query specification>"`
- The `-ot` command or the `-on` command can also be added to display the save sets in time based order or alphabetical order.

To list the snapshot save sets for a client, type the following at the command prompt:

```
mminfo -s server -S -a -q "client=clientname,snap"
```

where:

- `server` is the hostname of the NetWorker server.
- `clientname` is the hostname of the NAS device.

Example output:

```
mminfo -a -q"client=isilon-1.emc.com,snap" -r "ssid,savetime,name"
ssid      date      name
4259966917 1/29/2014 /ifs
643255862 4/16/2014 /ifs/data
609701430 4/16/2014 /ifs/data/folders
1700220469 1/6/2013 /ifs/data/folders/anne
2736137348 3/4/2014 /ifs/data/linux
576146998 5/7/2012 /ifs/data/amr_jr
1633111605 2/5/2014 /ifs/localtgt/2
777473590 3/28/2014 /ifs/localtgt/2
```

The *NetWorker Command Reference Guide* and the NetWorker man pages provide details on the `mminfo` command.

Using nsrnassnap_recover operations

You can use the `nsrnassnap_recover` command to recover save set files and directories.

When you specify the file or directory to recover, directories must end with a slash character (/).

For example:

Single file:

```
nsrnassnap_recover -S file123 -s nas456.lss.cor.com -c
file789.lss.cor.com -d d:\recover\7777 -A NSR_PS_FILE_RECOVER_OPT=Y -A
NSR_DIRECTED_RECOVER_HOST=file1133 -D 9 /vol/jjr4/test1.xml
```

Multiple files:

```
nsrnassnap_recover -S 3005375043 -s 10.63.29.222 -c 10.63.29.239 -d C:
\recover_test -A NSR_PS_FILE_RECOVER_OPT=Y -A
NSR_DIRECTED_RECOVER_HOST=10.63.29.222 -D 9 /ifs/R3/d1/
t1.txt /ifs/R3/d1/t2.txt/ifs/R3/d1/t3.txt
```

Entire directory tree:

```
nsrnassnap_recover -S file123 -s nas456.lss.cor.com -c
file789.lss.cor.com -d d:\recover\7777 -A NSR_PS_FILE_RECOVER_OPT=Y -A
NSR_DIRECTED_RECOVER_HOST=file1133 -D 9 /vol/jjr4/
```

The recovery options for the `-A NSR_PS_FILE_RECOVER_OPT=(Y/N/R)` command are as follows:

- Y = Overwrite the file or directory
- N = Skip the file or directory
- R = Rename the file or directory

Note

For NetApp devices, a recovery option must be specified, otherwise the recovery fails. For Isilon and NetApp devices, specify the `-R` remote hostname option to recover from a remote replication backup.

APPENDIX C

Snapshot limitations

This appendix includes the following topics:

- [NAS snapshot support and limitation](#).....100
- [Device-specific limitations for snapshot recovery](#).....100

NAS snapshot support and limitation

For backup and recovery operations the following limitations apply:

- LINUX systems must use NFS exports.
- Microsoft Windows systems must use CIFS exports.

Note

On Windows systems, a file share connection cannot exist between the Windows system running the recovery and the NAS device with the stored data. For NetApp remote replication on mirror policy clone, use node-scoped NDMP on the cluster credentials.

Device-specific limitations for snapshot recovery

The following tables list the snapshot recovery limitations and the support specific to each of the supported NAS devices. Restore from clone copies is supported on all supported platforms.

Table 14 Device limitations and support for snapshot management

NAS device	Limitation
Isilon	<ul style="list-style-type: none"> • Remote replication does not support in-place recoveries and out-of-place recoveries. • Renaming the data is supported only for directed recovery. • Directed recoveries overwrite an existing file.
VNX/ VNX2	<ul style="list-style-type: none"> • In-place recoveries and out-of-place recoveries are not supported. • Renaming the data is supported only for directed recovery. • Directed recoveries overwrite an existing file. • Recovery of NFS volumes on Windows systems is not supported. • Recovery of CIFS volumes on Linux systems is not supported.
VNXe	<ul style="list-style-type: none"> • In-place recoveries and out-of-place recoveries are not supported. • Renaming the data is supported only for directed recovery. • Recovery of NFS volumes on Windows systems is not supported. • Recovery of CIFS volumes on Linux systems is not supported. • Directed recoveries overwrite an existing file.
NetApp	<ul style="list-style-type: none"> • In-place recoveries and out-of-place recoveries restore only selected files. NetApp does not restore directory structures. • To enable selected file-by-file (FBF) recovery from snapshot replications, you must specify the NSR_MOUNTPOINT_NAME variable. -A NSR_MOUNTPOINT_NAME=<i>source_volume</i>

Table 14 Device limitations and support for snapshot management (continued)

NAS device	Limitation
	<p>Using Application Information variables on page 90 provides details.</p> <ul style="list-style-type: none"> • If a recovered file exists at the recovery destination then you must specify an overwrite option: <ul style="list-style-type: none"> ▪ For NMC Recovery Wizard recovery, select an overwrite option. ▪ For CLI recovery, when you issue the recover command, you must specify an overwrite option in the <code>nsrsnapadmin</code> prompt, for example: <pre>-A NSR_PS_FILE_RECOVER_OPT=Y</pre> <p>Windows example:</p> <pre>nsrsnapadmin> r -S 3895199998 -T 10.13.999.99 Current working directory is /vol/win_bang/ snaprecover> add NetWorkerEXT.XML.old 1 file(s) marked for recovery snaprecover> recover -A NSR_PS_FILE_RECOVER_OPT=Y</pre> • On Linux operating systems, NetApp devices cannot use a temporary mount point for in-place or out-of-place recovery: <ul style="list-style-type: none"> ▪ Ensure the <code>NSR_MOUNTPOINT_NAME</code> variable is set. ▪ For NMC Recovery Wizard recovery, in the Browse and recover save set step, ensure that Create temporary mount point for access is not selected. ▪ For CLI recovery, when you issue the recover command, in the <code>nsrsnapadmin</code> prompt, you must specify that you will use the existing mount point. <pre>-A NSR_USE_EXISTING_MOUNTPOINT=Yes</pre> <p>Linux example:</p> <pre>nsrsnapadmin> r -S 4046199999 -A NSR_USE_EXISTING_MOUNTPOINT=yes -A NSR_MOUNTPOINT_NAME=/vol/arnab -T 10.13.999.99 Current working directory is /vol/arnab/ snaprecover> cd dir1 snaprecover> add file1 1 file(s) marked for recovery snaprecover> recover -A NSR_PS_FILE_RECOVER_OPT=Y</pre> • On Linux systems, renaming the recovery data is not supported for in-place or out-of-place recoveries. Renaming the data is supported only for directed recovery.

Table 14 Device limitations and support for snapshot management (continued)

NAS device	Limitation
	<p>Note</p> <p>For a successful NDMP recovery, you must add <code>NDMP_MNT_HOST=<remote device></code> to the application variables before a save.</p> <p>For NDMP recovery using a NetApp Cluster Mode system, in the Recovery Configuration wizard, on the Select the Recovery Options window, select the New destination path option. You must specify: <code>/vserver name/volume name</code>. If you are using the command line option using the CLI, specify: <code>-m /vserver name/volume name</code>.</p> <p>For out of place restore, the directory to be restored has to be created in advance in the NetApp array.</p> <p>Only for remotely replicated NDMP recoveries, the <code>-R</code> command is required. For example:</p> <pre>nsrndmp_recover -c <client name>-s <server> -m /<Vserver name>/<volume name> -S <ssid> -R <source vservers></pre> <hr/> <p>Note</p> <p>For Isilon and NetApp remote replication NDMP recoveries, type <code>-R</code> with the source filer specified in the extra recovery options.</p>
Unity	<ul style="list-style-type: none"> • In-place recoveries and out-of-place recoveries are not supported. • Renaming the data is supported only for directed recovery. • Recovery of NFS volumes on Windows systems is not supported. • Recovery of CIFS volumes on Linux systems is not supported. • Directed recoveries overwrite an existing file. • Recovery of NFS volumes on Windows systems is not supported. • Recovery of CIFS volumes on Linux systems is not supported.

Table 15 Device support for recovery of save sets and individual files

Recovery from snapshot	Isilon	Isilon local replication	Isilon remote replication	NetApp	VNX	VNX2	VNXe	Unity
Directed recovery of save set	Yes	Yes	Yes	Not applicable	Yes	Yes	Yes	Yes
In-place recovery of save set	Yes	Yes	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable
Out-of-place recovery of save set	Yes	Yes	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable

Table 15 Device support for recovery of save sets and individual files (continued)

Recovery from snapshot	Isilon	Isilon local replication	Isilon remote replication	NetApp	VNX	VNX2	VNXe	Unity
Directed recovery of individual files	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
In-place recovery of individual files	Yes	Yes	Not applicable	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Out-of-place recovery of individual files	Yes	Yes	Not applicable	Yes	Not applicable	Not applicable	Not applicable	Not applicable

Table 16 NetApp support for types of local replication and platforms

Local replication 7-Mode	SnapVault (Windows)	SnapMirror (Windows)	SnapVault (Linux)	SnapMirror (Linux)
Backup	Yes	Yes	Yes	Yes
PIT restore from CLI				
FBF directed restore	Yes	Yes	Yes	Yes
FBF in-place restore	Yes	No	Yes	No
FBF out-of-place restore	Yes	No	Yes	No
PIT restore from GUI				
FBF directed restore	Yes	Yes	Yes	Yes
FBF in-place restore	Yes	No	Yes	No
FBF out-of-place restore	Yes	No	Yes	No
PIT restore from clone copy				
In-place restore from clone copy	Yes	Yes	Yes	Yes

Note

For C-Mode local replication clones, enable **Node Scope Mode** on the Netapp filer.

Table 17 NetApp C-Mode local replication and platforms

Local replication C-Mode	SnapVault (Windows)	SnapMirror (Windows)	SnapVault (Linux)	SnapMirror (Linux)
Backup	Yes	Yes	Yes	Yes
PIT restore from CLI				
FBF directed restore	Yes	No	Yes	No
FBF in-place restore	Yes	No	Yes	No

Table 17 NetApp C-Mode local replication and platforms (continued)

Local replication C-Mode	SnapVault (Windows)	SnapMirror (Windows)	SnapVault (Linux)	SnapMirror (Linux)
FBF out-of-place restore	Yes	No	Yes	No
PIT restore from GUI				
FBF directed restore	Yes	No	Yes	No
FBF in-place restore	Yes	No	Yes	No
FBF out-of-place restore	Yes	No	Yes	No
PIT restore from clone copy				
In-place restore from clone copy	Yes	Yes	Yes	Yes

Table 18 NetApp support for types of remote replication and platforms

Remote replication NetApp	SnapVault (Windows)	SnapMirror (Windows)	SnapVault (Linux)	SnapMirror (Linux)
Backup	Yes	Yes	Yes	Yes
PIT restore from CLI				
FBF directed restore	Yes	Not applicable	Yes	Not applicable
FBF in-place restore	Yes	Not applicable	Yes	Not applicable
FBF out-of-place restore	Yes	Not applicable	Yes	Not applicable
PIT restore from GUI				
FBF directed restore	Yes	Not applicable	Yes	Not applicable
FBF in-place restore	Yes	Not applicable	Yes	Not applicable
FBF out-of-place restore	Yes	Not applicable	Yes	Not applicable
PIT restore from clone copy				
In-place restore from clone copy	Yes	Not applicable	Yes	Not applicable

GLOSSARY

This glossary contains the definitions of terms used in this guide.

B

- backup** 1. Duplicate of database or application data, or an entire computer system, stored separately from the original, which can be used to recover the original if it is lost or damaged.
2. Operation that saves data to a volume for use as a backup.
- backup volume** A volume used to store backup data. NetWorker backup data cannot be stored on an archive volume or a clone volume.
- browse policy** NetWorker policy that specifies the period of time during which backup entries are retained in the client file index. Backups listed in the index are browsable and readily accessible for recovery.

C

- client** Host on a network, such as a computer, workstation, or application server whose data can be backed up and restored with the backup server software.
- client file index** Database maintained by the NetWorker server that tracks every database object, file, or file system backed up. The NetWorker server maintains a single index file for each client computer. The tracking information is purged from the index after the browse time of each backup expires.
- client-initiated backup** See [manual backup](#).
- Client resource** NetWorker server resource that identifies the save sets to be backed up on a client. The Client resource also specifies information about the backup, such as the schedule, browse policy, and retention policy for the save sets.
- clone** 1. Duplicate copy of backed-up data, which is indexed and tracked by the NetWorker server. Single save sets or entire volumes can be cloned.
2. Type of mirror that is specific to a storage array.
- cluster** Group of linked virtual or physical hosts, each of which is identified as a node, with shared storage that work together and represent themselves as a single host.
- consistent** State of a data set that fully matches an application's active view of the data at any point in time.
- Console server** See [NetWorker Management Console \(NMC\)](#).

D

- datazone** Group of clients, storage devices, and storage nodes that are administered by a NetWorker server.
- device**
1. Storage folder or storage unit that can contain a backup volume. A device can be a tape device, optical drive, autochanger, or disk connected to the server or storage node.
 2. General term that refers to storage hardware.
 3. Access path to the physical drive, when dynamic drive sharing (DDS) is enabled.

F

- file index** See [client file index](#).
- file system**
1. Software interface used to save, retrieve, and manage files on storage media by providing directory structures, data transfer methods, and file association.
 2. Entire set of all files.
 3. Method of storing files.

G

- group** One or more client computers that are configured to perform a backup together, according to a single designated schedule or set of conditions.

H

- host** Computer on a network.
- hostname** Name or address of a physical or virtual host computer that is connected to a network.

L

- live backup** See [rollover-only backup](#).
- logical unit (LUN)** Logical storage space on a storage array that is addressed by SCSI or Fibre Channel protocols. Multiple LUNs can be used to represent a storage volume.
- logical volume manager (LVM)** Software that controls disk resources by mapping data between a logical view of storage space and the actual physical disks.

M

- manual backup** Backup that a user performs from the client, also known as an unscheduled, on-demand, or ad hoc backup.

media	Physical storage, such as a disk file system or magnetic tape, to which backup data is written. See volume .
media database	Database that contains indexed entries of storage volume location and the life cycle status of all data and volumes managed by the NetWorker server.
mirror	<ol style="list-style-type: none"> 1. Exact duplicate of data or another object, such as a disk. 2. To write duplicate data to more than one device.
mount	To make a volume physically available for use, such as the placement of a removable disk volume or tape into a drive for reading or writing.
mount host	Host in a network that is used to mount storage array snapshot volumes to perform snapshot restore and rollover operations.

N

network attached storage (NAS)	Disk array or storage device (NAS filer) that connects directly to the messaging network or LAN interfaces and uses the common communication protocols of TCP/IP or NDMP.
Network Data Management Protocol (NDMP)	Software component that uses TCP/IP standards to specify how heterogeneous network components communicate for the purposes of backup, recovery, and transfer of data between storage systems.
NetWorker Management Console (NMC)	Software program that is used to manage NetWorker servers and clients. The NMC server also provides reporting and monitoring capabilities for all NetWorker processes.
NetWorker server	Computer on a network that runs the NetWorker server software, contains the online indexes, and provides backup and restore services to the clients and storage nodes on the same network.
NetWorker Snapshot Management (NSM)	Technology that provides point-in-time snapshot copies of data. NetWorker software backs up data from the snapshot. This allows applications to continue to write data during the backup operation, and ensures that open files are not omitted.

P

point-in-time copy (PIT copy)	Fully usable copy of a defined collection of data, such as a consistent file system, database, or volume that contains an image of the data as it appeared at a specific point in time. A PIT copy is also called a snapshot or shadow copy.
policy	Set of defined rules for client backups that can be applied to multiple groups. Groups have dataset, schedule, browse, and retention policies.
pool	<ol style="list-style-type: none"> 1. NetWorker sorting feature that assigns specific backup data to be stored on specified media volumes. 2. Collection of NetWorker backup volumes to which specific data has been backed up.

Q

quiesce State in which all writes to a disk are stopped and the file system cache is flushed. Quiescing the database prior to creating the snapshot provides a transactionally consistent image that can be remounted.

R

recover To restore data files from backup storage to a client and apply transaction (redo) logs to the data to make it consistent with a given point-in-time.

resource Software component whose configurable attributes define the operational properties of the NetWorker server or its clients. Clients, devices, schedules, groups, and policies are all NetWorker resources.

restore To retrieve individual data files from backup media and copy the files to a client without applying transaction logs.

retention policy NetWorker setting that determines the minimum period of time that backup data is retained on a storage volume and available for recovery. After this time is exceeded, the data is eligible to be overwritten.

rollback restore Process by which a snapshot is restored to its source or alternate location by using the capability of the storage array. A rollback restore destroys existing data on the target location.

rollover Backup of a snapshot to conventional storage media, such as disk or tape. Previously known as a live backup.

rollover-only backup Rollover whereupon the snapshot copy is deleted. Previously known as a serverless backup, live backup, or nonpersistent backup.

S

save set

1. Group of tiles or a file system copied to storage media by a backup or snapshot rollover operation.
2. NetWorker media database record for a specific backup or rollover.

save set ID (ssid) Internal identification number assigned to a save set.

scheduled backup Type of backup that is configured to start automatically at a specified time for a group of one or more NetWorker clients. A scheduled backup generates a bootstrap save set.

snapshot Point-in-time, read-only copy of specific data files, volumes, or file systems on an application host. Operations on the application host are momentarily suspended while the snapshot is created on a proxy host. Also called a PIT copy, image, or shadow copy.

snapshot backup Snapshot created on a storage array as a backup. Previously called instant backup.

snapshot mount host See [mount host](#).

snapshot policy	Sets of rules that control the life cycle of snapshots. These rule specify the frequency of snapshot creation, how long snapshots are retained, and which snapshots will be backed up to conventional storage media.
snapshot restore	Restore from a snapshot backup. Previously called instant restore.
snapshot save set	Group of files or other data included in a single snapshot. Previously called a snapset.
source LUN	LUN on the application host, whose production data can be copied by snapshot.
specific point-in-time (SPIT) copy	RecoverPoint term for PIT copy. See point-in-time copy (PIT copy) .
storage array	Integrated collection of subsystem disks, controllers, bus adapters, and software that provides storage services to one or more hosts.
storage node	Computer that manages physically attached storage devices or libraries, whose backup operations are administered from the controlling NetWorker server. Typically a “remote” storage node that resides on a host other than the NetWorker server.
T	
target client	NetWorker client on which data is to be restored This may be the same as the original source client from which the data was backed up, or it may be a different client.
V	
volume	<ol style="list-style-type: none"> 1. Unit of physical storage medium, such as a disk or magnetic tape, to which backup data is written. 2. Identifiable unit of data storage that may reside on one or more computer disks.
volume group	Logical group that contains sets of disks, as configured by a volume manager.

