

Dell EMC NetWorker

Version 19.2.x

Network Data Management Protocol User Guide

REV 02

May 2020

Copyright © 2015-2020 Dell Inc. or its subsidiaries. All rights reserved.

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 Technologies, Dell, EMC, 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

Figures	7
Tables	9
Preface	11
Chapter 1	Introduction 15
	Overview of NDMP.....16
	Components in a NetWorker NDMP environment.....16
	Configurations in a NetWorker NDMP environment.....16
	NDMP local backup.....17
	NDMP backups to non-NDMP devices (NDMP-DSA).....18
	Three-party backup.....21
	Pre-configuration requirements for NDMP data operations.....22
	Locale requirements with NDMP.....22
	Memory and space requirements for NDMP FH updates.....23
	Performance considerations.....23
	NDMP licensing requirements.....23
	NDMP feature requirements.....24
Chapter 2	PowerStore, Unity, Celerra, VNX, and VNXe 29
	Choosing a device type.....30
	Configuring devices for NDMP operations.....30
	NDMP device limitations.....30
	Determining NDMP device pathnames.....30
	Dynamic drive sharing.....31
	Configuring NDMP devices.....37
	Configuring NDMP-DSA devices.....40
	Configuring the Clone Storage Node.....40
	Pools requirements for NDMP.....40
	Configure NetWorker for NDMP backup and clone operations.....41
	Performing schedule backup and clone operations.....41
	Creating and configuring the NDMP client resource.....67
	Performing manual NDMP backups.....74
	Troubleshooting NDMP configuration and backup failures for PowerStore, Unity, Celerra, VNX, and VNXe.....76
	Monitoring NetWorker Server activities in the Administration window.....78
	About the Monitoring window.....79
	Customizing the Monitoring window.....80
	Policies/Actions pane.....81
	Sessions window.....85
	Alerts pane.....86
	Devices pane.....87
	Operations window.....88
	Log window.....90
	Recover window.....90

	Reporting NDMP Data.....	93
	Querying the NDMP volumes by backup type with the mminfo command.....	93
	Querying the NDMP save sets with the mminfo command.....	94
	Performing NDMP recoveries.....	94
	NDMP recovery requirements for Celerra and VNX.....	95
	Recover window.....	97
	Performing an NDMP index-based file-by-file data recovery.....	100
	Performing a Full or Directory Restore of NDMP data by using a save set recovery.....	103
	Troubleshooting NDMP recover.....	106
Chapter 3	Isilon	107
	Choosing a device type.....	108
	Configuring devices for NDMP operations.....	108
	NDMP device limitations.....	108
	Determining NDMP device pathnames.....	108
	Dynamic drive sharing.....	109
	Configuring NDMP on Isilon filer.....	114
	Configuring NDMP devices.....	115
	Configuring NDMP-DSA devices.....	118
	Configuring the Clone Storage Node.....	118
	Pools requirements for NDMP.....	118
	Configure NetWorker for NDMP backup and clone operations.....	119
	Performing schedule backup and clone operations.....	119
	Creating and configuring the NDMP client resource.....	145
	Performing manual NDMP backups.....	152
	Troubleshooting NDMP configuration and backup failures for Isilon...	154
	Monitoring NetWorker Server activities in the Administration window.....	156
	About the Monitoring window.....	157
	Customizing the Monitoring window.....	158
	Policies/Actions pane.....	159
	Sessions window.....	162
	Alerts pane.....	163
	Devices pane.....	164
	Operations window.....	165
	Log window.....	167
	Recover window.....	167
	Monitoring checkpoint-enabled backups.....	170
	Reporting NDMP Data.....	173
	Querying the NDMP volumes by backup type with the mminfo command.....	173
	Querying the NDMP save sets with the mminfo command.....	173
	Performing NDMP recoveries.....	174
	NDMP recovery requirements for Isilon.....	174
	Recovering data from partial save sets.....	176
	Recover window.....	176
	Performing an NDMP index-based file-by-file data recovery.....	179
	Performing a Full or Directory Restore of NDMP data by using a save set recovery.....	182
	Troubleshooting NDMP recover.....	185
Chapter 4	NetApp	187
	Choosing a device type.....	188

Configuring devices for NDMP operations.....	188
NDMP device limitations.....	188
Determining NDMP device pathnames.....	188
Dynamic drive sharing.....	189
Configuring NDMP devices.....	196
Configuring NDMP-DSA devices.....	199
Configuring the Clone Storage Node.....	199
Pools requirements for NDMP.....	199
Configure NetWorker for NDMP backup and clone operations.....	199
Creating and configuring the NDMP client resource.....	199
Performing schedule backup and clone operations.....	207
Performing manual NDMP backups.....	233
Troubleshooting NDMP configuration and backup failures for NetApp....	235
Monitoring NetWorker Server activities in the Administration window.....	238
About the Monitoring window.....	239
Customizing the Monitoring window.....	240
Policies/Actions pane.....	241
Sessions window.....	244
Alerts pane.....	246
Devices pane.....	246
Operations window.....	247
Log window.....	249
Recover window.....	250
Monitoring checkpoint-enabled backups.....	252
Reporting NDMP Data.....	255
Querying the NDMP volumes by backup type with the mminfo command.....	255
Querying the NDMP save sets with the mminfo command.....	255
Performing NDMP recoveries.....	256
NDMP recovery requirements for NetApp.....	256
Recovering data from partial save sets.....	258
Recover window.....	258
Performing an NDMP index-based file-by-file data recovery.....	261
Performing a Full or Directory Restore of NDMP data by using a save set recovery.....	264
Troubleshooting NDMP recover.....	267
 Chapter 5	
Other filers	269
Choosing a device type.....	270
Configuring devices for NDMP operations.....	270
NDMP device limitations.....	270
DinoStor-managed jukeboxes.....	270
Determining NDMP device pathnames.....	271
Dynamic drive sharing.....	272
Configuring NDMP devices.....	277
Configuring NDMP-DSA devices.....	280
Configuring the Clone Storage Node.....	280
Pools requirements for NDMP.....	280
Configure NetWorker for NDMP backup and clone operations.....	281
Creating and configuring the NDMP client resource.....	281
Performing schedule backup and clone operations.....	285
Performing manual NDMP backups.....	312
Troubleshooting NDMP configuration and backup failures for other filers.....	314

Monitoring NetWorker Server activities in the Administration window.....	316
About the Monitoring window.....	317
Customizing the Monitoring window.....	318
Policies/Actions pane.....	319
Sessions window.....	322
Alerts pane.....	323
Devices pane.....	324
Operations window.....	325
Log window.....	327
Recover window.....	327
Monitoring checkpoint-enabled backups.....	330
Reporting NDMP Data.....	333
Querying the NDMP volumes by backup type with the mminfo command.....	333
Querying the NDMP save sets with the mminfo command.....	333
Performing NDMP recoveries.....	334
NDMP recovery requirements for other filers.....	334
Recover window.....	336
Performing an NDMP index-based file-by-file data recovery.....	339
Performing a Full or Directory Restore of NDMP data by using a save set recovery.....	342
Troubleshooting NDMP recover.....	345

FIGURES

1	NDMP local backup configuration.....	18
2	Backup started from a NetWorker server with an attached storage device.....	19
3	NDMP backup that uses immediate save.....	20
4	Three-party NDMP backup to NDMP devices.....	22
5	Dynamic Drive Sharing.....	33
6	Data Protection Policy.....	42
7	Platinum policy configuration.....	43
8	Data protection policy example.....	44
9	Workflow path from a traditional backup action.....	49
10	Traditional backup workflow.....	66
11	Dynamic Drive Sharing.....	110
12	Data Protection Policy.....	120
13	Platinum policy configuration.....	121
14	Data protection policy example.....	122
15	Workflow path from a traditional backup action.....	127
16	Traditional backup workflow.....	144
17	Dynamic Drive Sharing.....	191
18	DDS with NDMP.....	194
19	Data Protection Policy.....	209
20	Platinum policy configuration.....	209
21	Data protection policy example.....	211
22	Workflow path from a traditional backup action.....	216
23	Traditional backup workflow.....	233
24	Dynamic Drive Sharing.....	273
25	Data Protection Policy.....	287
26	Platinum policy configuration.....	287
27	Data protection policy example.....	289
28	Workflow path from a traditional backup action.....	294
29	Traditional backup workflow.....	311

TABLES

1	Revision history.....	11
2	Style conventions.....	13
3	Distinctions between NDMP Device Backup and NDMP-DSA	17
4	NDMP features.....	24
5	Shared Devices attributes.....	36
6	Schedule icons.....	50
7	Schedule icons.....	54
8	Schedule icons.....	57
9	Schedule icons.....	63
10	Application information variable types.....	68
11	Celerra and VNX Application Information variables.....	68
12	Celerra and VNX Environment variables.....	71
13	NDMP backup.....	75
14	Monitoring window panel	79
15	Policy status icons.....	82
16	Sessions that can be stopped from NMC	85
17	Alerts window icons.....	86
18	Devices status icons	87
19	Operations window icons.....	88
20	Icons in the Log pane.....	90
21	Recovery toolbar options	91
22	Save recover configuration job status.....	92
23	Find options.....	93
24	Recovery toolbar options	97
25	Save recover configuration job status.....	98
26	Find options.....	99
27	Shared Devices attributes.....	113
28	Schedule icons.....	128
29	Schedule icons.....	131
30	Schedule icons.....	135
31	Schedule icons.....	140
32	Application information variable types.....	146
33	Isilon Information variables.....	146
34	NDMP backup.....	153
35	Monitoring window panel	157
36	Policy status icons.....	159
37	Sessions that can be stopped from NMC	163
38	Alerts window icons.....	164
39	Devices status icons	165
40	Operations window icons.....	166
41	Icons in the Log pane.....	167
42	Recovery toolbar options	168
43	Save recover configuration job status.....	169
44	Find options.....	170
45	New Checkpoint Restart media attributes.....	172
46	Checkpoint enabled save sets.....	172
47	Partial save sets for the checkpoint id.....	173
48	Recovery toolbar options	177
49	Save recover configuration job status.....	178
50	Find options.....	179
51	Shared Devices attributes.....	195
52	Application information variable types.....	200
53	Vendor-specific Application Information variables.....	201

54	NetApp Environment Variables.....	203
55	Schedule icons.....	217
56	Schedule icons.....	220
57	Schedule icons.....	224
58	Schedule icons.....	229
59	NDMP backup.....	235
60	Monitoring window panel	239
61	Policy status icons.....	241
62	Sessions that can be stopped from NMC	245
63	Alerts window icons.....	246
64	Devices status icons	247
65	Operations window icons.....	248
66	Icons in the Log pane.....	249
67	Recovery toolbar options	250
68	Save recover configuration job status.....	251
69	Find options.....	252
70	New Checkpoint Restart media attributes.....	254
71	Checkpoint enabled save sets.....	254
72	Partial save sets for the checkpoint id.....	255
73	Recovery toolbar options	259
74	Save recover configuration job status.....	260
75	Find options.....	261
76	Shared Devices attributes.....	276
77	Application information variable types.....	282
78	Vendor-specific Application Information variables.....	282
79	Schedule icons.....	295
80	Schedule icons.....	298
81	Schedule icons.....	302
82	Schedule icons.....	308
83	NDMP backup.....	313
84	Monitoring window panel	317
85	Policy status icons.....	319
86	Sessions that can be stopped from NMC	323
87	Alerts window icons.....	324
88	Devices status icons	325
89	Operations window icons.....	326
90	Icons in the Log pane.....	327
91	Recovery toolbar options	328
92	Save recover configuration job status.....	329
93	Find options.....	330
94	New Checkpoint Restart media attributes.....	332
95	Checkpoint enabled save sets.....	332
96	Partial save sets for the checkpoint id.....	333
97	Recovery toolbar options	336
98	Save recover configuration job status.....	338
99	Find options.....	339

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 describes how to use the NetWorker software to provide data protection for Network Data Management Protocol (NDMP) filers.

Each filer-specific chapter contains an end-to-end workflow that describes how to configure NetWorker to use the NDMP protocol to protect the filer data.

Audience

This guide is a part of the NetWorker documentation set, and is intended for use by system administrators who are responsible for setting up and maintaining backups on a network, and network-attached storage (NAS) filer administrators.

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
02	May 8, 2020	Changing instances of 'Unity, Celerra, VNX, and VNXe' to 'PowerStore, Unity, Celerra, VNX, and VNXe'.
01	November 15, 2019	First release of the document for NetWorker 19.2.

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.
- *NetWorkerSnapshot 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 19.2 with CloudBoost 19.2 Integration Guide*
Describes the integration of NetWorker with CloudBoost.
- *NetWorker 19.2 with CloudBoost 19.2Security 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.

 **Note:** References to Data Domain systems in this documentation, in the UI, and elsewhere in the product include PowerProtect DD systems and older Data Domain systems.

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: • 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 find 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

Introduction

This chapter includes the following topics:

- [Overview of NDMP](#) 16
- [Components in a NetWorker NDMP environment](#) 16
- [Configurations in a NetWorker NDMP environment](#) 16
- [Pre-configuration requirements for NDMP data operations](#) 22

Overview of NDMP

The network data management protocol (NDMP) is a TCP/IP-based protocol that specifies how network components communicate for the purpose of moving data across the network for backup and recovery.

The NDMP protocol addresses the problems that are associated with backing up data in heterogeneous environments when you use different operating systems, backup solutions, and Network Attached Storage (NAS) devices.

The NDMP enables disparate vendors to use a common NDMP protocol for the backup architecture. With the NetWorker NDMP interface, you can connect to hosts that have an active NDMP service or an NDMP data module installed. You do not install the NetWorker software on the NDMP host. NDMP allows a NAS device to back up data to other NDMP-controlled tape or disk devices that are on the network. NDMP passes control of the data and the file metadata to and from the NetWorker software.

By default, the NetWorker server tries to establish communications with a NAS filer by using NDMP version 4. If the NAS does not support or use NDMP version 4, communications are automatically negotiated to use the highest NDMP version that the NAS filer supports. NetWorker supports NDMP version 3 and later, however, some NetWorker features require a specific version of NDMP on the NAS.

The *NetWorker Hardware Compatibility Guide* on the Online Support website provides a list of NAS filers that the NetWorker software supports.

Components in a NetWorker NDMP environment

Three main components support NDMP data operations with the NetWorker software:

- **NDMP Data Server**—The system that contains the NDMP data. The NDMP Data Server is also known as the data mover, the NDMP client, or the NAS device. The NAS transfers the data to the primary storage devices through a data connection. You configure the NAS as a client of the NetWorker server. However, you do not install the NetWorker client software on the NAS.
- **NDMP Tape Server**—The host with the backup device to which NetWorker writes the NDMP data.
- **Data Management Agent (DMA)**—The NetWorker server is the DMA. The DMA performs the following tasks:
 - Starts the NDMP backup
 - Monitors the NDMP backup and recovery operations
 - Maintains the media database and the client file index entries for NDMP backups
 - Maintains the resource database information for the NAS and NDMP Tape Server

Configurations in a NetWorker NDMP environment

You can use three methods to configure the NDMP Data Server and the NDMP Tape Server to perform backups and recoveries.

The NAS device passes NDMP metadata to the NetWorker server in all the methods. You can customize the NetWorker environment to support NDMP data operations in each of the following scenarios:

- **NDMP local backup** on page 17—The NDMP Data Server and the NDMP Tape Server reside on the same physical host. The physical host is the NAS host.

- [NDMP backups to non-NDMP devices \(NDMP-DSA\)](#) on page 18—An NDMP Data Server Agent (NDMP-DSA) sends the NDMP data from the NDMP Data Server to non-NDMP devices. The NDMP Data Server and the NDMP Tape Server reside on different physical hosts. The NDMP Tape Server is always a NetWorker Server or a NetWorker Storage Node.
- [Three-party backup](#) on page 21—The NDMP Data Server and the NDMP Tape Server reside on different physical hosts. The NDMP Tape Server can be a NAS host, a NetWorker Server, a NetWorker Storage Node, or a Third party NDMP vendor. Three-party backup is also called Three-way backup.

 **Note:** NDMP-DSA is also a type of Three-party backup.

The following table summarizes the differences between the NDMP Device Backup and NDMP-DSA.

Table 3 Distinctions between NDMP Device Backup and NDMP-DSA

NDMP Device Backup	NDMP-DSA
Supports only the NDMP type of tape device	Supports any type of device that the NetWorker software supports
Does not support backup to disk	Supports backup to disk
Does not support multiplexing	Supports multiplexing
Does not support archiving	Supports archiving

NDMP local backup

In an NDMP local backup (Direct-NDMP), the NDMP Data Server (NAS) sends data to a locally attached tape device or library. The `nsrndmp_save` program runs on the NetWorker server, and only the metadata and the NDMP control information traverse the network between the NetWorker server and the NDMP host.

An advantage of NDMP local backup is the NDMP data does not traverse the network, which prevents network congestion.

The NDMP local backup has the following disadvantages:

- The `nsrndmp_save` program queries the NDMP Tape Server at consistent intervals to determine the status of a backup. These queries have an impact on backup performance.
- The NetWorker software does not multiplex NDMP save sets and writes NDMP data serially to the local device. As a result, backups are slower, but recoveries are faster.
- NDMP local backups are unsuitable when you need to back up many large file systems on an NAS filer.
- You cannot archive NDMP save sets.

The NetWorker server, or data management application (DMA), performs these tasks:

- Starts the backup or the recovery request through the NDMP connection.
- Receives the file history from the data server.


During a backup, the NAS filer is the NDMP Tape Server and the NDMP Data Server. The NAS filer performs the following tasks:

- Receives the backup requests.
- Reads the backup data on the NAS disks.
- Produces a data stream for the backup.

- Writes the data stream to the tape or file device.

The following figure illustrates a local backup configuration.

Figure 1 NDMP local backup configuration



NDMP backups to non-NDMP devices (NDMP-DSA)

In this scenario, the NetWorker software writes NDMP data to non-NDMP devices, including tape, virtual tape, AFTD, and Data Domain devices. NDMP-DSA backups are more flexible than backups to NDMP devices, because NDMP devices must be physical or virtual tape.

Use NDMP-DSA backups when any of the following scenarios are true:

- There are many small file systems to backup, and network throughput is not a concern.
- No physical or virtual autochanger is available for backups.
- The NAS system is not attached to a storage area network (SAN) for communication with a physical or virtual autochanger.

Also, directing NDMP staged and cloned data to a non-NDMP device is faster than sending the data to an NDMP device.

The NetWorker software uses the NDMP Data Server Agent (DSA) and the `nsrndmp_save` command to send NDMP data to a non-NDMP device. The `nsrdsa_save` process is associated with DSA.

Use NDMP_DSA when you want the following benefits:

- Write NDMP data to devices that also contain non-NDMP data.
- Multiplex NDMP save sets to improve backup speeds. Recovery speeds are slower.
- Stage save sets from the disk to tape.
- Archive NDMP save sets.

You can back up NDMP data to a non-NDMP device in one of two ways:

- NDMP data sent to non-NDMP devices that are local to the NetWorker server.
- NDMP data sent to non-NDMP devices that reside on a NetWorker storage node.

NDMP data sent to non-NDMP devices that are local to the NetWorker server


NDMP data that you send to non-NDMP devices that are local to the NetWorker servers has the following characteristics:

- The backup data traverses the network between the NetWorker server and the NDMP Data Server.

- The metadata, the NDMP control information, and the file history (FH) remains local to the NetWorker server and still traverses the network.

The following figure illustrates a NetWorker storage device that is attached directly to the NetWorker server. The NetWorker server starts an NDMP-DSA backup. The `nsrmmd` process on the NetWorker server processes the data and metadata. The `nsrndmp_2fh` and `nsrdmpix` processes on the NetWorker server process the FH data and then pass the FH data to the `nsrindexd` process.

Figure 2 Backup started from a NetWorker server with an attached storage device



To configure NDMP data sent to non-NDMP devices local to NetWorker server, set NetWorker server to the Storage Nodes field on the NDMP client properties, or set the non-NDMP devices to the media pool.

NDMP data sent to non-NDMP devices that reside on a NetWorker storage node

You can configure NDMP backups to write data to a NetWorker storage node in one of two ways.

Immediate save (`nsrdsa_save` runs on storage node)


When you configure an NDMP backup with the immediate save technology, the following actions occur:

- The `nsrdsa_save` backup command runs on the NetWorker storage node.
- The NetWorker software uses TCP/IP and shared memory to communicate between the `nsrdsa_save` and `nsrmmd` processes.
- The NetWorker server processes the backup data and sends the data to the non-NDMP device directly through the `nsrmmd` process on the storage node.

When the NetWorker software uses immediate save to send the NDMP data, the following actions occur:

- The `nsrindexd` process on the NetWorker server processes the file history.
- After the data backup completes and the sessions with the NDMP Data Server and the NetWorker server close, the NetWorker software commits the FH to the client file index that is associated with the NDMP client.

The following figure illustrates a NetWorker configuration that uses immediate save.

Figure 3 NDMP backup that uses immediate save

To configure immediate save:

1. Add the DSA host which runs `nsrdsa_save` to the remote access list of the NDMP client resource.
2. Add the `-M -P <proxy-host>` option for the `nsrndmp_save` command, or add `NSR_DSA_NODE=<proxy-host>` in the **Application Information** attribute in the Client Resource configuration on the NetWorker Server. Here the proxy-host is the DSA which runs `nsrdsa_save`. If both options are configured, the second one takes precedence.
3. Set the storage node attached with the non-NDMP devices to the **Storage Nodes** field on the NDMP client properties, or set the non-NDMP devices on the target storage node to the media pool. Please note the storage node here could be same with or different from the DSA which runs `nsrdsa_save` configured in the previous steps.

Examples

Below are two examples displaying the configuration for NDMP backup with DSA which is not on a NetWorker server.

Example 1: DSA and storage node on the same host

Host Type	Name	IP
NDMP filer	ndmp-server	10.0.0.1
NetWorker server	nw-server	10.0.0.2
DSA host/ Storage node	sn-server	10.0.0.3

Perform the following steps:

1. Create an NDMP client *ndmp-server*.
2. Add storage node *sn-server* and the non-NDMP devices attached on it.
3. Add *sn-server* to the Remote Access list of the client properties of *ndmp-server*.
4. Add *sn-server* to the **Storage Nodes** field of the client properties of *ndmp-server*.
5. Add `-M -P sn-server` to the option of the backup command `nsrndmp_save` for the client *ndmp-server* or add `NSR_DSA_NODE=sn-server` to the application information of the client *ndmp-server*.

Example 2: DSA and storage node on different hosts

Host Type	Name	IP
NDMP filer	ndmp-server	10.0.0.1

Host Type	Name	IP
NetWorker server	nw-server	10.0.0.2
DSA host	dsa-server	10.0.0.3
Storage node	sn-server	10.0.0.4

Perform the following steps:

1. Create an NDMP client *ndmp-server*.
2. Add storage node *sn-server* and the non-NDMP devices attached on it.
3. Add *dsa-server* to the Remote Access list of the client properties of *ndmp-server*.
4. Add *sn-server* to the Storage Nodes field of the client properties of *ndmp-server*.
5. Add `-M -P dsa-server` to the option of the backup command `nsrndmp_save` for the client *ndmp-server* or add `NSR_DSA_NODE=dsa-server` to the application information of the client *ndmp-server*.

Non-immediate save (nsrdsa_save runs on NetWorker Server)

By default, NDMP backups to a non-NDMP device use non-immediate save.

When you configure an NDMP backup to use non-immediate save, the following actions occur:

1. The `nsrdsa_save` backup command runs on the NetWorker server.
2. The `nsrdsa_save` process uses TCP/IP to read the data in a local buffer.
3. The `nsrdsa_save` process transmits the data to the `nsrmmmd` process on the storage node.
4. The `nsrmmmd` process writes the data to the storage device.

This approach is inefficient and has slow performance for the following reasons:

- Backup data traverses the network between the NetWorker server, the NDMP host, and the NetWorker storage node.
- Metadata and the NDMP control information traverse the network between the NetWorker server and the storage node.
- FH traverses the network between the NetWorker server and the NDMP Data Server.

To configure non-immediate save, set target storage node to the Storage Nodes field on the NDMP client properties or set the non-NDMP devices on the target storage node to the media pool.

Three-party backup

A three-party or three-way backup sends NDMP data to an NDMP Tape Server, but the NDMP Data Server and the NDMP Tape Server are not the same physical host.

There are mainly two types of three-party backups:

- In the first scenario, NetWorker sends the NDMP data to non-NDMP devices (NDMP-DSA). The NDMP Data Server and the NDMP Tape Server reside on different physical hosts. The NDMP Tape Server is always a NetWorker Server or a NetWorker Storage Node. Hence, NDMP-DSA is also a Three-way NDMP backup.
- In the second scenario, NetWorker sends NDMP data to NDMP devices. Here the data flows from the NDMP Data Server to the NDMP Tape Server, and then to a library that is locally attached to the NDMP Tape Server. In this configuration, you cannot archive the NDMP save sets.

In addition to using a NetWorker server or storage node as the NDMP Tape Server, you can use a third-party NDMP DinoStor Tape Server. This hardware connects one or more libraries to the

network and enables you to back up any NDMP host to one location instead of requiring a local backup device for each server.


The following figure shows a three-party backup configuration, which enables backup and recovery operations on an NDMP device that is attached to another NDMP server.

In this example:

- The first server is the data server.
- The second server is the tape server.
- The third party is the NetWorker server (DMA).

This configuration enables a similar flow of data between a NetWorker client and a NetWorker server or storage node, but you do not need to install the NetWorker software on either of the NDMP hosts. Data flows from the NDMP Data Server over the network to the NDMP Tape Server and then to tape. The NDMP Data Server sends the metadata to the NetWorker server.

Figure 4 Three-party NDMP backup to NDMP devices



Pre-configuration requirements for NDMP data operations

This section provides the requirements that you should review before you configure the NetWorker software for NDMP data operations.

Note: The *NetWorker Security Configuration Guide* describes how to determine port requirements for NDMP backups and recoveries when a firewall exists in the NetWorker datazone.

Locale requirements with NDMP

When running NDMP backups, ensure that you use consistent locale settings in the environment:

- NetWorker supports the UTF-8 format with CIFS clients. NetWorker supports NFS clients of a NetApp filer only when the NFS clients can generate UTF-8 encoded data. If you set the `UTF8=Y` variable during an NDMP backup and the backup contains path names with non-ASCII characters, index-based recoveries of the backup fail with an error message similar to the following:

```
RESTORE: could not create path pathname
```

- If you use the `UTF8=Y` variable and perform a backup, you must recover path names that contain non-ASCII characters by using either a save set recovery from the command line or an NDMP Directory restore.

- All UNIX locale settings on the NAS filer, including UTF-8, are the same.
- Configure the NAS filer to use UTF-8 character sets. Contact the NAS vendor for configuration requirements.
- Use only the UNIX NetWorker Management Console (NMC) clients that have the same locale setting as the NAS filer.
- You can perform backup and recovery operations in any locale. However, if you try to browse in a locale that is different from the original locale, the file names appear as random characters.
- A single save set supports data that belongs to only one code set. If you have data in multiple code sets, you must create multiple save groups.
- A save set can contain file names that belong to different languages if all characters in those languages belong to the same code set. For example, ISO 8859-1 and ISO 8859-15 include most Western European languages, such as French, Spanish, and Portuguese. NetWorker can back up file names from these languages in a single save set.

Memory and space requirements for NDMP FH updates

During an NDMP backup, the NDMP Data Server sends the file history (FH) metadata information to the NetWorker server. The NetWorker software does not verify or modify FH metadata that is received from the NAS. The NetWorker software uses the FH information to maintain appropriate indexes and media database entries for the NDMP client backups.

The `nsrndmp_2fh` and `nsrdmpix` binaries interact with the raw database, instead of virtual memory, to process the FH metadata. As a result, memory requirements for this process are minimal. The NetWorker server stores metadata updates in the `\nsr\tmp` directory and commits the metadata to the client file index after the NDMP client backup completes.

Use the following formula to determine the required physical space for the default `\nsr\tmp` directory:

$2 * (144 + \text{average file name length}) * \text{number of entries in the file system}$

 **Note:** This formula must be used for non-multistreaming only.

For example:

For one million file entries with an average file name length of 128, use this formula to compute the required temporary swap space:

$2 * (144 + 128) * 1,000,000 = 544 \text{ MB approximately}$

Performance considerations

Volume loading and positioning operations do not occur during a volume selection process because an information exchange occurs between the `nsrmmmd` process and the `nsrndmp_save` or `nsrndmp_recover` command. Backing up to a filetype device avoids the overhead that is associated with the exchange of information.

On NetApp filers with Data OnTap 6.4 and later, NetWorker reads all metadata from tape before recovering the files. For large save sets with 20 million files or more, the recovery time for a file can exceed 3 hours. This estimate also applies to backups because NetWorker records the metadata for the whole volume onto the tape during a single file backup.

NDMP licensing requirements

NetWorker with NDMP requires an additional license, separate from the NetWorker base product, because NetWorker uses a tiered or capacity-based licensing model. The *NetWorker Licensing Guide* provides more information.

NDMP feature requirements

Before you implement the IPv6, checkpoint restart, SMTape, iSCSI, vbb, and DAR/DDAR features in NetWorker, review the information in the following table.

Table 4 NDMP features

Feature	Information
Multistreaming	NetWorker 9.0.1 and later supports multistreaming for Isilon OneFS 8.0 and later backups, up to 32 streams. NetWorker uses the client parallelism value that is defined for an Isilon client to determine how many backups run concurrently. When a backup with multiple save sets start, NetWorker sends one stream to each available device. If the number of streams is greater than the number of devices, NetWorker distributes the remaining save streams evenly among the devices.
IPv6	NetWorker storage nodes support IPv6 communications with a NetWorker server. By default, NDMP backup and recovery operations use IPv6 to create the data connection between the NDMP data server and mover server, when the following requirements are met: • NDMP data server and mover server support the Connection Address Extension (CAE). • NDMP data server and mover server use NDMP v4. • NDMP data server and mover server have a configured IPv6 interface. • In an NDMP-DSA configuration, you use an 8.2 SP1 or later storage node.
CAB extension support	NetWorker support for NetApp c-mode requires the NDMP v4 Cluster Aware Backup (CAB) extension on the NetApp filer.
Checkpoint restart	The NetApp and Isilon filers create a snapshot of the file system before the backup. The save set is generated from the snapshot. A checkpoint restart: • Requires an NDMP v4 restartable backup extension. • Results in slower backups because NetWorker writes the checkpoint files at defined intervals. The more frequently NetWorker writes checkpoint files, the slower the backup occurs. • Supports the creation of partial save sets from a snapshot when the backup is interrupted, and later restarted. As a result, the partial backups provide an image of the file system from the point-in-time that the snapshot is taken.  Note: Checkpoint restart backups do not support multistreaming on Isilon filers. The <i>NetWorker Administration Guide</i> provides more information about checkpoint restarts.

Table 4 NDMP features (continued)

Feature	Information
Snapmirror to tape (SMTape)	SMTape provide the following benefits: • Performs block-level backup of SnapMirror volumes on NetApp filers. • Reduces the backup window when millions of files reside on the NetApp filer. Use the SMTape feature in instances where NDMP full backups become impractical. The SMTape option: • Copies large NetApp file systems to secondary storage instead of using the standard NDMP full or differential backups. • Supports a 240 KB block size. • Allows mirroring of backups to disk and to tape devices. • Supports only full volume backups and recoveries. You cannot use SMTape for file indexes or file restores. • Supports only save set recoveries. • Does not support incremental and differential backup levels. The NAS generates a snapshot of the file system at the beginning of the SMTape operation. Use environment variables to control the conditional call to retain or delete the snapshot. Creating and configuring the NDMP client resource on page 67 describes how to configure an NDMP client using SMTape.
iSCSI	NetWorker supports iSCSI LUNs on Celerra and NetApp filers. Celerra and VNX filers do not support NDMP-based backups and recoveries of iSCSI LUNS. NetApp filers support NDMP-based backups and recoveries of iSCSI LUNS, but you cannot perform an index-based recovery. To perform a full save set recovery to another volume, use a destination volume that is at least two and a half times as large as the source volume. NetApp recommends using SnapMirror instead of backups to safeguard iSCSI LUNS.
Volume Based Backup (vbb)	vbb supports: • Data Access in Real Time (DART) version 5.5 and later. • Index-based recoveries of a Celerra or VNX block level to the same volume or another location. Use: ■ Checkpoint configuration utility to configure checkpoint file systems on the Celerra or VNX before you perform a backup. ■ Full Destructive Restore (FDR) to perform a full save set recovery of a raw volume of equal or greater size than the backup.

Table 4 NDMP features (continued)

Feature	Information
	NetWorker performs a file-by-file recovery when you: - Recover data from a Celerra or VNX block-level backup. This recover requires disk space in the root directory of the target file system to store temporary recovery files. - Perform save set recoveries and NDMP Directory Restores to an existing file system. When you backup a volume that uses native Celerra deduplication, you cannot perform an index-based or NDMP Directory Restore of the backup. You can only perform an FDR restore from a level FULL save set. Configuring NDMP backups on Celerra and Using Celerra Data Deduplication documentation on the Support website provides detailed information about how to prepare the filer before you perform FDR.
Direct Access Recovery (DAR) and Dynamic Direct Access Recovery (DDAR)	DAR and DDAR send file information from the NAS filer to the NetWorker server. This action allows a single-file recovery or a directory recovery to position to the exact location of the data on the tape media. NetWorker does not read the file and record numbers sequentially to locate the data. You cannot use DDAR when you enable vbb. DAR and DDAR support DART version 5.5 or later and NetApp with OnTap version 6.4 and later. DAR and DDAR require NDMP version 3 or later. Recoveries on earlier NDMP versions fail.
Differential backup levels	Backups with levels 1, 2, ...9 are differential backup levels. You can only perform these backups through the CLI by adding the option "-l <backup_level>" to the <code>nsrndmp_save</code> command.

About partial save sets

The backup sequence of partial save sets is not the same as the backup sequence for complete backups. Each partial save set provides protection for part of the file system, but the completeness and consistency of the coverage of the whole file system cannot be guaranteed.

The checkpoint restart window is user-defined and can be large. If restarted hours apart, the partial backups might provide an image of the file system that is different from the state of the file system at any fixed point in time. The resulting file system backup is not guaranteed to be consistent.

NetWorker performs file and directory backups in alphabetical order. If a failure occurs, and you restart the backup, the backup operation starts alphabetically with the next file or folder that was not previously backed up. NetWorker does not review files or folder that were previously backed up for changes. If a previously backed up file or folder was edited or added after the backup failure, NetWorker does not back up the file or directory again.

Consider the following example in which a backup is interrupted while it is saving a directory and is restarted after the directory contents have changed:

1. A save set contains `/disk1/dir` with files `file_a`, `file_c` and `file_d`.

2. The backup of the save set is interrupted while `file_d` is backed up.

As a result, the first partial save set includes only `file_a` and `file_c`.

3. A user adds `file_b` to the file system.

4. The checkpoint restart is initiated for the save set.

The second partial save set contains `file_d` and `/disk1/dir`, which includes `file_a`, `file_b`, `file_c` and `file_d`. However, `file_b` is not in the save set.

CHAPTER 2

PowerStore, Unity, Celerra, VNX, and VNXe

This chapter includes the following topics:

• Choosing a device type	30
• Configuring devices for NDMP operations	30
• Configure NetWorker for NDMP backup and clone operations	41
• Monitoring NetWorker Server activities in the Administration window	78
• Reporting NDMP Data	93
• Performing NDMP recoveries	94

Choosing a device type

Network Data Management Protocol (NDMP) backups can be written to either an NDMP device, or if using NDMP-DSA, to a non-NDMP device.

Perform either of the following tasks:

- Configure devices for NDMP operations.
- Configure non-NDMP devices. If you are using NDMP-DSA, refer to the *NetWorker Administration Guide* for device configuration.

For a description of each configuration, refer to [Configurations in a NetWorker NDMP environment](#) on page 16.



Note:

PowerStore does not support local or direct NDMP backup.

Configuring devices for NDMP operations

Review this section for information about how to configure the NetWorker environment for Network Data Management Protocol (NDMP) data operations.

The *NetWorker Hardware Compatibility Guide* on the Support website provides a list of NDMP devices that the NetWorker software supports.

NDMP device limitations

Review these limitations before you configure Network Data Management Protocol (NDMP) devices:

- The timeout of the NetWorker server `nsrmmmd` resource attribute does not apply to NDMP devices, but it does apply to storage nodes devices.
- You cannot use the `jbexercise` utility with an NDMP autochanger.
- You cannot configure NDMP devices on a dedicated storage node.
- You must use a non-rewind device handle for the NDMP media device handle.
- You cannot configure advanced file type devices and file type devices as NDMP devices.
- You cannot configure an NDMP autochanger when the NDMP protocol is earlier than version 3. You must determine the NDMP device handles, then use the `jbconfig` command to configure the autochanger.

Determining NDMP device pathnames

To configure an NDMP stand-alone device or an NDMP jukebox, you must first determine the path names of the media devices. If the NAS filer does not support the NDMP_CONFIG interface or uses NDMP version 3, you must also determine the library device handle.

To determine the NDMP device path names and the library handle, use the `inquire` command or vendor-specific commands.

Determining the NDMP device path names using the `inquire` command

About this task

Use the `inquire` command to determine the path names and library handle.

Procedure

1. From a command prompt on the NetWorker server, type:

```
inquire -N NAS_hostname -T
```

2. When prompted, specify the NAS username and password.

NOTICE Use the `inquire` command with caution. When you run `inquire`, the command sends the `SCSI inquiry` command to all devices that are detected on the SCSI bus. If you use the `inquire` command during normal operations, unforeseen errors can occur, which might result in data loss.

Determining the NDMP device pathnames for Celerra and VNX

Before you begin

Before you configure an NDMP autochanger, determine the device path names of NDMP devices and of the robotic arm.

About this task

Use the Celerra or VNX Administrator program or manually query the `scsidevs` file to determine the device path names.

Procedure

- To manually query the `scsidevs` file, log in to the filer with the NDMP account and type the following command:

```
server_devconfig data_mover_name -p -s -n
```

The host responds with a list of media device names, for example:

```
server_2 :
Scsi device table
name addr type info
jbox1 c1t010 jbox ATL P1000 62200501.21
tape2 c1t410 tape QUANTUM DLT7000 245Fq_
tape3 c1t510 tape QUANTUM DLT7000 245Fq_
```

- To avoid tape drive issues, set the `ntape` parameter for every tape drive that you discover on a particular data mover. For example, if a data mover has five configured tape drives, set the parameter to `NDMP ntape=5`
- To modify the NDMP `ntape` parameter, edit the `/nas/server/slot_#/param` file, where `slot_#` correlates directly to the server number and restart the filer.

You cannot specify a value greater than 8 for `ntape`. Configuring NDMP on Celerra documentation on the Support website provides detailed information about configuring a Celerra filer.

Dynamic drive sharing

Dynamic Drive Sharing (DDS) is a feature that provides NetWorker software with the ability to recognize shared physical tape drives. DDS enables NetWorker software to perform the following operations:

- Skip the shared tape drives that are in use.
- Route the backups or recoveries to other available shared tape drives.

Introduction to DDS

DDS controls application requests for tape media and allows the NetWorker server and all storage nodes to access and share all attached devices.

A system administrator can configure DDS by setting a sharing policy for devices that are accessible from multiple storage nodes.

There are two terms that are central to the use of DDS are drive and device. Within the context of DDS, these terms are defined as follows:

- Drive—The physical backup object, such as a tape drive, disk, or file.
- Device—The access path to the physical drive.

 **Note:** NetWorker only supports DDS in a storage area network (SAN) Fibre Channel environment and not in a direct-connect SCSI environment.

Benefits of DDS


Enabling DDS on a NetWorker system provides these benefits:

- Reduces storage costs—You can share a single tape drive among several storage nodes. In fact, since NetWorker software uses the same open tape format for UNIX, Windows, NetWare and Linux, you can share the same tape between different platforms (assuming that respective save sets belong to the same pool).
- Reduces LAN traffic—You can configure clients as SAN storage nodes that can send save sets over the SAN to shared drives.
- Provides fault tolerance—Within a SAN environment, you can configure hardware to eliminate a single point of failure.
- Provides configuration over a greater distance—You can configure a system over a greater distance than with SCSI connections.

DDS configuration overview

The following figure illustrates the DDS process and potential device sharing configurations. This basic configuration consists of a server, two storage nodes, and a library with two tape drives.

Figure 5 Dynamic Drive Sharing



In this figure:

- Storage nodes sn_1 and sn_2 are attached to the library.
- Each storage node, on its own, has access to drive_1 and drive_2.
- With DDS enabled, both storage nodes have access to both drives and can recognize when a shared drive is in use.

This configuration requires two DDS licenses, one for each drive.

Note: Ensure that all applicable devices can be seen from each storage node by running the `inquire -l` command locally on each storage node.

DDS block-size compatibility between UNIX and Windows

With DDS enabled, drives can be shared between storage nodes on different platforms, such as UNIX and Microsoft Windows. For NetWorker software operations (such as backups and recoveries) to take place successfully, ensure that the block size is compatible between different platforms or hardware.

To ensure compatibility, make sure one of the following conditions is met:

- The various storage nodes sharing a drive support the same block sizes.
- When a tape is labeled on a drive, it is labeled with the block size defined on the storage nodes.

Block-size incompatibility between UNIX and Windows

Incompatible block-size settings between UNIX and Microsoft Windows storage nodes could result in any of these error scenarios:

- A backup taken on a UNIX node might not be recoverable on a Microsoft Windows node if the Windows node does not support large block sizes.
- A UNIX process labels and saves data to a tape and leaves the tape mounted. A Microsoft Windows process subsequently attempts to verify the label on this tape and fails because the label verification is done by reading a header from the data portion.
- A tape on a UNIX node is labeled with a large block size. The backup is started on a Microsoft Windows node and the Windows node attempts to write the backup by using the default block size. Internally, the backup on Windows is written by breaking down the big buffer of data into smaller segments of writable block sizes.

Attempting to recover a specific file on Windows in this situation fails due to positioning errors on the tape. The data is still recoverable from the Windows side, since the NetWorker software will switch from using file and block positioning to reading the tape from the beginning to reach the correct position. The data might not, however, be recoverable from the UNIX side.

Unintended Access to DDS device prevention

The Reserve/Release attribute has been added to the Device resource for tape devices to support Reserve/Release, including the Persistent Reserve commands.

Reserve/Release is a mechanism that uses SCSI commands to attempt to prevent unintended access to tape drives that are connected by using a shared-access technology, such as Fibre Channel, iSCSI, or SCSI multiplexers. It is a “cooperative” and host-based mechanism, which means that all applications should respect the reservations and not purposely break them. Access is granted based on the host system that reserved the device. Other applications that run on that host cannot be prevented from accessing a reserved device.

Reserve/Release cannot prevent a malicious or badly behaved application from accessing a reserved device. It also cannot prevent all problems caused by hardware issues (such as SCSI resets or FC LIPs) from interrupting data access.

The basic sequence requires that a host reserve a tape drive (using specific SCSI commands) before attempting to access the tape drive. If this “reservation” succeeds, then the host can use the drive. If the reservation fails (usually because the device is reserved by someone else), then the host attempting the reservation should not attempt to use the drive. When a host has finished using a reserved drive, that host must release the drive by using the appropriate SCSI commands.

The reservation is maintained by the drive itself. With older (called “Simple” in NetWorker software) Reserve/Release, the reservation is based on the SCSI ID of the system that issued the reserve command. For tape drives connected to Fibre Channel (FC) using FC-SCSI bridges, the mapping between FC host and reservation is done inside the bridge, since the initiator on the SCSI side is always the bridge itself, regardless which host actually issued the reserve command.

For Persistent Reserve, the reservation is associated with a 64-bit “key” that is registered by the host. Several keys can be registered with a given drive at any given time, but only one may hold the active reservation. NetWorker software uses the “exclusive” reservation method for Persistent Reserve. Only the host that holds the active reservation is allowed to access the drive.

The Reserve/Release attribute does not support file type or advanced file type devices.

The settings that relate to Reserve/Release and Persistent Reserve are found in a device’s **Properties** window, on the **Advanced** tab. They are visible only when diagnostic mode is turned on.

The default setting for Reserve/Release is None. Once any other Reserve/Release setting is selected, it works automatically, without further user intervention. The Reserve/Release attribute

is supported only on Common Device Interface (CDI) platforms, so if the CDI attribute in a device's **Properties** is set to Not Used, then Reserve/Release settings are ignored.

For newer hardware, once a Reserve/Release setting (other than None) has been selected, the appropriate Persistent Reserve commands are automatically issued before a device is opened for reading or writing, and before the device is closed. With older hardware, a SCSI-2 Reserve command is issued before opening the device, and a SCSI-2 Release command is issued after the device is closed.

Reserve/Release has these possible settings:

- None (the default)
- Simple
- Persistent Reserve
- Persistent Reserve + APTPL (Activate Persist Through Power Loss)

The Persistent Reserve Key attribute has also been added. It is used with Persistent Reservation calls.

Restrictions for use of the SCSI Reserve/Release setting

There are restrictions for using the SCSI Reserve or Release setting.

Consider the following:

- It is available on CDI platforms only. Consequently, since CDI is not supported within an NDMP environment, Reserve/Release is not supported with NDMP.
- Not all drives support persistent Reserve/Release. (All drives support at least simple reserve release. The code automatically drops back from Persistent +APTPL or Persistent to Simple on drives that do not support Persistent.)
- SCSI resets can clear Simple reservations at the device.
- Even with Reserve/Release, there is no guarantee against data loss.
- If the operating system has its own Reserve/Release feature, that feature must be disabled in order for the NetWorker Reserve/Release feature to work.
- Even if all of the enterprise's NetWorker storage nodes have this feature enabled, then it is possible that, on the storage node where a backup operation is run, data loss can be caused by the operating system's utilities or by third-party programs.

DDS on NDMP nodes in a SAN environment

You can configure shared drives between NDMP nodes in a SAN environment.

Ensure that:

- All the components of a SAN configuration are compatible when DDS is enabled with the NetWorker NDMP feature.
- The Fibre Channel switches are compatible with any NDMP hosts within a SAN.
- NDMP hosts and libraries in the SAN are compatible with each other.
- The NDMP nodes that will share the drives are homogeneous.

 **Note:** The current NDMP implementation does not allow the sharing of drives between non-homogeneous NDMP nodes. There is, however, no inherent limitation within DDS that would prevent this.

DDS attributes in the device properties

Configure the attributes that DDS uses, in the **Properties** window for a device.

The attributes include:

- Hardware ID
- Shared Devices

Hardware ID attribute

The Hardware ID attribute tracks the drives that are shared between multiple hosts. Device instances that share the same physical drive across multiple hosts have the same hardware ID. The device autoconfiguration process automatically assigns the Hardware ID to a device, or it is added when manually configuring a device. Users cannot edit the Hardware ID.

You can view the Hardware ID in the **Properties** window for a device, on the **General** tab, in the **Device Sharing** area.

NetWorker generates the Hardware ID when a device is scanned or configured. The Hardware ID consists of the following components:

- Hardware serial number
- Device type
- Worldwide part number (WWPN)
- Worldwide name (WWN)

Shared Devices attribute

The Shared Devices attribute appears on the **Operations** tab of a device's **Properties** window when in diagnostic mode. It features values that can be used to manipulate all shared instances of a drive simultaneously. This attribute enables or disables all devices that share the same Hardware ID with a single action. The following table lists allowed values and descriptions for the attribute.

Table 5 Shared Devices attributes

Value	Description
Enable All	When selected, enables all devices with the same Hardware ID.
Disable All	When selected, disables all the devices with the same Hardware ID.
Done	This value is the default setting. After the server has enabled or disabled all devices with the same Hardware ID, the attribute value is reset to Done.

You cannot configure the Shared Devices attribute with the `jbconfig` program.

Idle Device Timeout attribute and DDS

A tape might remain mounted in a drive after a backup completes. Other requests for the drive from another device path must wait during this timeout period. Use the Idle Device Timeout attribute to adjust the timeout value.

The Idle Device Timeout attribute is not specifically a DDS attribute, but is useful in configuring shared drives. This attribute appears on the device **Properties** window on the **Advanced** tab when displayed in Diagnostic Mode. The default value is 0 (zero) minutes, which means that the device never times out and you must manually eject the tape.

If the device belongs to a library, you can also specify the Idle Device Timeout value for all devices in the library. However, the library value will take effect only on those devices whose **Idle Device Timeout** value is 0. The Idle Device Timeout value for a library is located on the **Timer** tab of the library **Properties** window.

Max active devices

In a DDS environment, use the Max active devices attribute, on the **General** tab of the Storage Node resource to define the maximum number of active devices for a storage node.

This attribute sets the maximum number of devices that NetWorker may use from the storage node in a DDS configuration. In large environments with media libraries that have a large number of devices, storage nodes might not have the ability to optimize all the drives in the library. The Max active devices attribute allows you to limit the number of devices that the storage node uses at a specified time, which allows the storage node to have access to all the devices in the library, but does not limit the storage node to the number of devices it can fully optimize.

Configuring NDMP devices

You can back up NDMP data to an NDMP or non-NDMP device in a standalone or library configuration. You can also back up NDMP data to ACSLS controlled silos.

Configuring a standalone NDMP device

Use the NetWorker Management Console (NMC) to configure a standalone Network Data Management Protocol (NDMP) tape device for Direct NDMP backups.

Procedure

1. In the **Administration** window, click **Devices**.
2. In the navigation tree, right-click **Devices**, and then select **New**.
3. In the **Name** attribute, specify the NDMP device in the format:

```
rd=NAS_hostname:NAS_device_handle (NDMP)
```

where:

- *NAS_hostname* is the hostname of the NAS that has the NDMP device attached.
- *NAS_device_handle* is the path of the device.

 **Note:** Configure the NDMP device as a remote device and add (NDMP) after the pathname. Otherwise, you receive a message similar to the following:

```
NDMP device name shall be in rd=snode:devname (NDMP) format
```

4. In the **Media Type** attribute, specify the device type.
5. Specify a valid NAS administrator account in the **Remote User** attribute.

 **Note:** For the Celerra and VNX filers, specify the trusted account that was created for backup operations on each NDMP-Host Data Mover. Some Celerra versions require that you use a trusted account named *ndmp*. Configuring NDMP on Celerra on the Support website provides detailed information.
6. Specify the password for the NAS administrator account in the **Password** attribute.
7. On the **Configuration** tab:
 - a. Select the **NDMP** checkbox. You can only set this attribute when you create the device. You cannot change the NDMP attribute after you create the device. To change the device configuration, delete and re-create the device.
 - b. Set the **Target Sessions** attribute to 1. NDMP devices do not support multiplexing.
 - c. The **Dedicated Storage Node** attribute must remain at the default value: **no**.
8. Under the **Advanced** tab, the **CDI** attribute must remain at the default value: **Not used**.

9. (Optional) Change the block size that is used by the NDMP device.

By default, NDMP devices use a block size of 60 KB. If required, select a different block size in the **Device block size** field. When you configure the NDMP client, set the *NDMP_AUTO_BLOCK_SIZE* environment variable in the **Application Information** attribute.

10. Click **OK**.

Configuring an NDMP autochanger

You can use an NDMP autochanger to manage Direct NDMP or Three-party backups with NDMP devices. To configure an NDMP autochanger, use NMC or the `jbconfig` command.

Configuring an NDMP autochanger with NMC

When you configure an NDMP autochanger in NMC, the NetWorker software first detects the NDMP devices and then configures the library.

Procedure

1. In the **NetWorker Administration** window, click **Devices**.
2. Right-click the NetWorker Server, and then select **Configure All Libraries**.
3. On the **Provide General Configuration Information** window, accept the default library type, **SCSI/NDMP**, and then click **Next**.
4. On the **Select Target Storage Nodes** window, click **Create a new Storage Node**.
5. On the **Storage Node Name** field, specify the hostname of the NAS.
6. In the **Device Scan Type** attribute, select **NDMP**.
7. In the **NDMP User Name** and **NDMP Password** fields, specify the NAS administrator account. If DinoStor TapeServer manages the autochanger, specify the DinoStor username and password.
8. Click **Start Configuration**.
9. Click **Finish**.
10. Monitor the **Log** window for the status of the device scan.

When you specify an incorrect username and password combination:

- The Log status window reports:

```
No configured libraries detected on storage node
storage_node_name
```

- The `daemon.raw` file on the NetWorker server reports:

```
NDMP Service Debug: The process id for NDMP service is 0xb6c0b7b0
42597:dvdetect: connect auth: connection has not been authorized
42610:dvdetect: The NDMP connection is not successfully
authorized on host 'storage_node_name'
```

To resolve this issue, relaunch the **Configure All Libraries** wizard and correct the NDMP username and password combination.

Note: If the **Log** window reports that NetWorker cannot detect the serial numbers for the library, see [Configuring an NDMP autochanger by using the `jbconfig` command](#) on page 39 for detailed instructions.

Configuring an NDMP autochanger by using the `jbconfig` command

It is recommended that you use the NMC interface to configure an NDMP autochanger. Use the `jbconfig` command when you cannot configure the autochanger by using the NMC Configure Library wizard.

About this task

The *NetWorker Command Reference Guide* or the UNIX man page provides more information about the `jbconfig` command.

Procedure

1. Log in to the NetWorker server as root on UNIX, or Administrator on Windows.
2. At the command prompt, type `jbconfig`
3. At the **What kind of jukebox are you configuring** prompt, type 2 to configure an autodetected NDMP SCSI jukebox.
4. When prompted for an NDMP username, specify the NAS administrator account.
5. When prompted for an NDMP password, specify the NAS administrator password.
6. When prompted for the NDMP Tape Server Name, specify the NAS filer hostname.
7. At the **What name do you want to assign to this jukebox device** prompt, provide a name to identify the autochanger.
8. To enable auto-cleaning, accept the default value of **Yes**, otherwise type **no**.
9. At the **Is (any path of) any drive intended for NDMP use? (yes / no) [no]** prompt, type **yes**.
10. At the **Is any drive going to have more than one path defined? (yes / no) [no]** prompt, type **no** if you will not configure shared devices. Type **yes** to configure shared drives.
11. When prompted, for the first pathname for the NDMP devices in the jukebox, perform the following steps:
 - a. Specify the pathname in the following format:


```
NDMP_tape_server_name:device path
```

 where:
 - *NDMP_tape_server_name* is the hostname of the NDMP Server.
 - *device_path* is the first device path.
 - b. At the **Is this device configured as NDMP** prompt, type **yes**.
 - c. Repeat step a and step b for all NDMP devices in the autochanger.
 - d. When prompted, assign a hardware ID.
 - e. To use DDS:
 - Respond to the prompts as required so that the first host will have access to the shared drive.
 - When prompted to share this drive with another host, type **yes**.
 - When prompted, type the hostname and device path of the second host that will have access to the shared drive.
12. Complete the prompts for the second device.
13. In the **Enter the drive type of drive 1** prompt, specify the number that corresponds to the NDMP device type.

14. If each drive in the autochanger is the same model, then type **yes**. Otherwise, type **no**, and then specify the appropriate device types for each additional autochanger device.
15. When prompted to configure another autochanger, type **no**.

Changing the block size of an NDMP device

By default, the block size that is used to write data to an NDMP backup is 60KB. With the exception of Celerra, when you specify the *NDMP_AUTO_BLOCK_SIZE=Y* variable for an NDMP client, an NDMP device can use the value that is defined in its Device block size attribute.

About this task

To determine the block sizes that are supported by the NDMP filer before setting the block size for an NDMP device, consult the applicable vendor documentation.

To change the block size that is defined for the NDMP device, perform the following steps:

Procedure

1. From the **View** menu, select **Diagnostic Mode**.
2. In the **Devices** window, right-click the NDMP device, and then select **Properties**.
3. On the **Advanced** tab, select a value in the **Device block size** field.

 **Note:** The selected block size must not exceed the block size that is configured on the NAS filer.

4. Click **Ok**.

Configuring NDMP-DSA devices

When you use DSA, NetWorker sends the NDMP data to a NDMP-DSA device, which includes tape, virtual tape, AFTD, and Data Domain devices. The steps to configure a NDMP-DSA device for a specified device type is the same as configuring a non-NDMP device. The *NetWorker Administration Guide* provides detailed information.

Configuring the Clone Storage Node

When cloning NDMP data, specify the destination storage node, called the clone “write source” (the device that receives the clone data), in the Clone storage nodes attribute. The *NetWorker Administration Guide* provides details.

Pools requirements for NDMP

When you create a pool for non-NDMP devices, select only the devices that are required by the NDMP clients.

NetWorker cannot send bootstrap and index backups to an NDMP device. When you do not configure a non-NDMP devices or a non-NDMP device is not available to receive the index and bootstrap backups, the NDMP client backup appears to hang. Configure a separate pool to direct the index and bootstrap to a non-NDMP device.

Auto media verification in the Pool resource does not support NDMP.

When an NDMP client backup is a member of a clone-enabled group, configure a clone pool with non-NDMP devices that are local to the NetWorker server to receive the clone bootstrap and index.

Configure NetWorker for NDMP backup and clone operations

This section explains how to configure NetWorker for NDMP backup and clone operations.

Performing schedule backup and clone operations

Data Protection Policies provide you with the ability to schedule backup and clone operations, to protect NDMP data.

You can use the NDMP protocol to protect data on NAS devices.

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

Overview of protection policies

A protection policy allows you to design a protection solution for your environment at the data level instead of at the host level. With a data protection policy, each client in the environment is a backup object and not simply a host.

Data protection policies enable you to back up and manage data in a variety of environments, as well as to perform system maintenance tasks on the NetWorker server. You can use either the **NetWorker Management Web UI** or the NMC **NetWorker Administration** window to create your data protection policy solution.

A data protection policy solution encompasses the configuration of the following key NetWorker resources:

Policies

Policies provide you with a service-catalog approach to the configuration of a NetWorker datazone. Policies enable you to manage all data protection tasks and the data protection lifecycle from a central location.

Policies provide an organizational container for the workflows, actions, and groups that support and define the backup, clone, management, and system maintenance actions that you want to perform.

Workflows

The policy workflow defines a list of actions to perform sequentially or concurrently, a schedule window during which the workflow can run, and the protection group to which the workflow applies. You can create a workflow when you create a new policy, or you can create a workflow for an existing policy.

A workflow can be as simple as a single action that applies to a finite list of Client resources, or a complex chain of actions that apply to a dynamically changing list of resources. In a workflow, some actions can be set to occur sequentially, and others can occur concurrently.

You can create multiple workflows in a single policy. However, each workflow can belong to only one policy. When you add multiple workflows to the same policy, you can logically group data protection activities with similar service level provisions together, to provide easier configuration, access, and task execution.

Protection groups

Protection groups define a set of static or dynamic Client resources or save sets to which a workflow applies. There are also dedicated protection groups for backups in a VMware environment or for snapshot backups on a NAS device. Review the following information about protection groups:

- Create one protection group for each workflow. Each group can be assigned to only one workflow.
- You can add the same Client resources and save sets to more than one group at a time.
- You can create the group before you create the workflow, or you can create the group after you create the workflow and then assign the group to the workflow later.

Actions

Actions are the key resources in a workflow for a data protection policy and define a specific task (for example, a backup or clone) that occurs on the client resources in the group assigned to the workflow. NetWorker uses a work list to define the task. A work list is composed of one or several work items. Work items include client resources, virtual machines, save sets, or tags. You can chain multiple actions together to occur sequentially or concurrently in a workflow. All chained actions use the same work list.

When you configure an action, you define the days on which to perform the action, as well as other settings specific to the action. For example, you can specify a destination pool, a retention period, and a target storage node for the backup action, which can differ from the subsequent action that clones the data.


When you create an action for a policy that is associated with the virtual machine backup, you can select one of the following data protection action types:

- Backup — Performs a backup of virtual machines in vCenter to a Data Domain system. You can only perform one VMware backup action per workflow. The VMware backup action must occur before clone actions.
- Clone — Performs a clone of the VMware backup on a Data Domain system to any clone device that NetWorker supports (including Data Domain system or tape targets). You can specify multiple clone actions. Clone actions must occur after the Backup action.

You can create multiple actions for a single workflow. However, each action applies to a single workflow and policy.

The following figure provides a high level overview of the components that make up a data protection policy in a datazone.

Figure 6 Data Protection Policy



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 7 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.

Overview of configuring a new data protection policy

About this task

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 8 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*.

Strategies for traditional backups

The primary considerations for a traditional backup strategy are the groups of Client resources, the workflows that define the series of actions that are associated with the backup, and the schedule for the backup.

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 traditional backups

A protection groups for traditional backups identifies the client resources to back up.

Traditional backups support the following types of protection groups:

- **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.

Create multiple groups to perform different types of backups for different Client resources, or to perform backups on different schedules. For example:

- Create one group for backups of clients in the Accounting department, and another group for backups of clients in the Marketing department.
- Create one group for file system backups and one group for backups of Microsoft Exchange data with the NetWorker Module for Microsoft.
- Create one group for a workflow with backups actions that start at 11 p.m., and another group for a workflow with backup actions that start at 2 a.m.

 **Note:** A Client resource can belong to more than one group.

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.

About this task

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](#).

Actions sequences in traditional backup workflows

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


A workflow for a traditional backup can optionally include a probe or check connectivity action before the backup, and a clone action either concurrently with or after the backup.

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

Workflow path from a traditional backup action

The only action that can follow a traditional backup is a clone action.

Figure 9 Workflow path from a traditional backup action



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.

- 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.

- From the **Action Type** list, select **Check Connectivity**.
- 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, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
 - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
- 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.

Table 6 Schedule icons (continued)

Icon	Label	Description
	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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

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.
 - If a schedule is associated to an action, then override option is disabled.
20. Click **Next**.
The **Action Configuration Summary** page appears.
21. 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
 - Snapshot backup
-  **Note:** This option is not available for NAS snapshot backups.

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. In-built commands from NetWorker should be avoided as probe command.

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, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
 - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
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 7 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 action 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 action if any of the probes 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

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**.
i **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**.

i **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.
- If a schedule is associated to an action, then override option is disabled.

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 traditional backup action

A traditional backup is a scheduled backup of the save sets defined for the Client resources in the assigned group for the workflow.

Before you begin

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

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 **Backup**.
6. From the secondary action list, select the backup type, for example, **Traditional**.
7. (Optional) From the **Force Backup Level** list select a backup level.

For workflows that have more than one scheduled backup within a 24-hour period, use the **Force Backup Level** attribute to allow more than one backup to occur at two different backup levels in a 24-hour period. When you select a backup level in the **Force Backup Level** attribute, the first backup is performed at the scheduled backup level. Each subsequent occurrence of the backup action in the next 24 hours occurs at the level defined in the **Force Backup Level** attribute. For example, if the level defined by the schedule is Full

and the **Force Backup Level** attribute is set to Incr, the first backup started by the action occurs at a level full and subsequent backups, within 24 hours of the start of the full backup are incremental. By default this option is cleared, which means that if the action runs multiple backup operations in a 24 period, all the backups occur at the scheduled backup level.

8. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
9. 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.
10. Specify a weekly, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
 - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
11. To specify the backup level to perform when **Define** option under **Select Schedule** is selected, click the icon on each day. If it is **Select** option under **Select Schedule**, choose the customized NSR schedule using the drop down, after specifying the backup level within the NSR schedule resource.

The following table provides details about the backup level that each icon represents.

Table 8 Schedule icons

Icon	Label	Description
	Full	Perform a full backup on this day. Full backups include all files, regardless of whether the files changed.
	Incr	Perform an incremental backup on this day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
	Cumulative Incr	Perform a cumulative incremental backup. Cumulative incremental backups include files that have changed since the last full backup.
	Logs Only	Perform a backup of only database transaction logs.

Table 8 Schedule icons (continued)

Icon	Label	Description
	Incremental Synthetic Full  Note: Not supported for NDMP.	Perform an incremental synthetic backup on this day. An incremental synthetic full backup includes all data that changed since the last full backup and subsequent incremental backups to create a synthetic full backup.
	Skip	Do not perform a backup on this day.

To perform the same type of backup on each day, select the backup type from the list and click **Make All**.

NetWorker does not support the use of synthetic full backup levels for NDMP data.

Celerra, Isilon, VNX, Unity, and NetApp filers with NDMP version 4 or later support token-based backups (TBB) to perform NDMP full and incremental backups. NetWorker supports the same number of incremental levels that the NAS vendor supports. Celerra, Isilon, and NetApp documentation provide the maximum number of incremental levels that the TBB incremental backup can support.

When you configure TBB after you update the NetWorker server from 7.6 SP1 or earlier, the first incremental backup does not occur until after one complete full backup.

Filers that do not support TBB, do not support incremental backups. If you select the level Incr, the NetWorker server performs a full backup.

Verify that the NAS storage vendor supports NDMP incremental backups before you use this feature.

12. Click **Next**.

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

13. From the **Destination Storage Node** box, select the storage node with the devices on which to store the backup data.
14. From the **Destination Pool** box, select the media pool in which to store the backup data.
15. From the **Retention** boxes, specify the amount of time to retain the backup data.

After the retention period expires, the save set is removed from the client file index and marked as recyclable in the media database during an expiration server maintenance task.

When you define the retention policy an NDMP client, consider the amount of disk space that is required for the client file index. NDMP clients with several thousands of small files have significantly larger client file indexes on the NetWorker server than a non-NDMP client. A long retention policy for an NDMP client increases disk space requirements on the file system that contains the client file indexes.

16. From the **Client Override Behavior** box, specify how NetWorker uses certain client configuration attributes that perform the same function as attributes in the Action resource:
 - **Client Can Override**—The values in the Client resource for **Schedule**, **Pool**, **Retention policy**, and the **Storage Node** attributes take precedence over the values that are defined in the equivalent Action resource attributes.

Note: If the NetWorker policy action schedule is set to the `Skip` backup level, the **Client can Override** option is not honored. For NetWorker to consider the **Client can Override** option, change the action schedule to a different level.

- **Client Can Not Override**—The values in the Action resource for the **Schedule**, **Destination Pool**, **Destination Storage Node**, and the **Retention** attributes take precedence over the values that are defined in the equivalent Client resource attributes.
- **Legacy Backup Rules**—This value only appears in actions that are created by the migration process. The updating process sets the **Client Override Behavior** for the migrated backup actions to **Legacy Backup Rules**.

17. Click **Next**.

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

18. 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.

19. 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.

20. 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.

21. 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

22. 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.

23. 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.

24. 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.

25. (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.

26. (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.
- If a schedule is associated to an action, then override option is disabled.

27. 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**.

28. 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 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.

29. Click **Next**.

The **Action Configuration Summary** page appears.

30. 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 after the backup. A clone action is the only supported action after a backup action in a workflow.

Cloning NDMP save sets

You can clone Direct-NDMP and NDMP-DSA save sets by using the same methods used to clone non-NDMP save sets.

Before you clone NDMP save sets, review these requirements:

- To clone Direct-NDMP or Three-party backup data:
 - The source NAS must run NDMP version 3 or later.
 - The destination NAS can run any version of NDMP, but you cannot clone a volume cloned with NDMP earlier than version 3 to another volume.

- You cannot clone NDMP save sets to a non-NDMP device.
- You can clone NDMP tapes from one NDMP host to another NDMP host of the same type. For example, you can clone tapes from a NetApp filer with an attached library to another NetApp filer or to the same filer.
- You require two NDMP devices to clone the NDMP save sets, one device to perform the read operation and one device to perform the write operation.
- You must clone NDMP-DSA backups to non-NDMP devices. You can however, clone NDMP-DSA save from one type of tape device to another. For example you can clone save sets on a DLT device to an AIT device.
- Use the nsrclone program to clone NDMP save sets from a command prompt. The *NetWorker Command Reference Guide* or the UNIX man pages provide more information on nsrclone usage.

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

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 **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, monthly, or reference schedule for the action:
 - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
 - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.

- To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
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 9 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:
- From the **Destination Storage Node** list, select the storage node with the devices on which to store the cloned save sets.
 - 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.
 - From the **Destination Pool** list, select the target media pool for the cloned save sets.
 - 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:
- 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.
 - 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

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 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

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.
 - If a schedule is associated to an action, then override option is disabled.

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 traditional backup workflows

Figure 10 Traditional backup workflow



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.

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**.

Creating and configuring the NDMP client resource

Use the NMC Client Configuration wizard to create the NDMP client or create the client manually. It is recommended that you use the NMC Client Configuration wizard to create NDMP clients.

Using the Client Configuration wizard

About this task

Use the NMC Client Configuration wizard to create the NDMP client.

Procedure

1. From the **Administration** window in NMC, click **Protection**.
2. In the expanded left pane, select **Clients**, and then select **Protection > New Client Wizard**.
3. On the **Specify Client Information** window:
 - a. In the **Client Name** field, specify the hostname of the filer.
 - b. (Optional) Add comments in the **Comment** field.
 - c. (Optional) In the **Tag** field, specify the name of the tag for the dynamic group in which you want to add this client.
 - d. (Optional) In the **Groups** area, select an existing group, in which to add the client.
 - e. In the **Type** area, select **NDMP**, and then click **Next**.

For Celerra, you will typically use the CIFS server name configured on the Data Mover.

4. On the **Specify the NDMP Client Credentials** window:
 - a. In the **NDMP User Name** field, specify a valid NAS administrator account.
 - b. In the **NDMP Password** attribute, specify the password for the NAS administrator account, and then click **Next**.

For an Celerra and VNX filers, specify the trusted account that you created on each NDMP-Host Data Mover for backups. Some versions of Celerra require you to use an account called

ndmp. The Configuring NDMP on Celerra document on the Support website provides more information.

5. In the **Specify the NDMP Client Backup Options** window:

a. In the **NDMP backup type** attribute, select or specify the backup type:

- `dump`—Traverses a file tree in mixed width first and depth-first order.
- `vbb`—Used to back up the entire volume at the block level rather than at a file level. The vbb backup type reads data blocks more efficiently than traditional file-based backups. The vbb backup type does not support DDAR and TBB.
- `ts`—Enables a tape silvering backup.

b. In the **NDMP Array Name** field, specify the logical name that is assigned to the NDMP NAS array.

The **NDMP Array Name** field enables you to configure the same NAS device with multiple NDMP clients that have different host IDs.

 **Note:** NDMP clients that use the same NAS device must have the same NDMP array name.

c. Review the **App Info** options and disable options, as required. It is recommended that the default options remain enabled.

Table 10 Application information variable types

App Info Type	Description
HIST	Enables the backup of index data. If you do not select this option, you can only perform full recoveries of the backup data.
UPDATE	Enables the backup process to update the last backup dates in database on the NDMP client, after the backup completes. Only applies to NetApp and has no effect when backing up Celerra, VNX, or VNXe.
DIRECT	Enables DAR or DDAR support. DAR and DDAR on page 96 provides more information.
Use Token-Based Backup	Enables the NDMP backup to use last backup time tokens to decide what files to backup. Not all NDMP clients support token based backup. When you select this option and the NDMP data server on the client does not support token based backups, NetWorker performs the backup by using backup levels.

d. In the **Advanced App Info** field, specify additional NAS specific environments variables, one per line. The following table provides a list of the available **Application Information** environment variables for each NAS.

Table 11 Celerra and VNX Application Information variables

Variables	Definition
<code>DIRECT= n</code>	(Optional) When you use DAR or DDAR, you must set this value to <code>y</code> .

Table 11 Celerra and VNX Application Information variables (continued)

Variables	Definition
<i>EMC_EDIRnn=string</i>	(Optional) This string value identifies a directory to exclude from the backup. You can use asterisk (*) as a wildcard, but only when * is the last character in the string. To include multiple directories, increment the number. For example: - EMC_EDIR01=/fsX/DIRx - EMC_EDIR02=/fsX/DIRy Dart version 5.5 and later supports this variable. NetWorker ignores this variable when you perform a vbb backup.
<i>EMC_EFILEnn=string</i>	(Optional) This string value determines which files to exclude from the backup. You can use the asterisk (*) as a wildcard, but only when * is the first or last character in the string, or both. To include multiple files, increment the number. For example: - EMC_EFILE01=*mp3 - EMC_EFILE02=temp* Dart version 5.5 and later supports this variable. NetWorker ignores this variable when you perform a vbb backup.
<i>SNAPSURE=y</i>	(Required) This value ensures that a backup is performed on a SnapSure of the production file system to avoid any inconsistent problems of file access during backup window.
<i>USE_TBB_IF_AVAILABLE=n</i>	(Optional) The NetWorker software automatically enables TBB support for Celerra filers. Specify this variable and value to disable TBB support for incremental backups and when you use the vbb backup type. When you specify this value, the backup reverts to the native level-based backup of the NAS.
<i>ALLOW_SINGLE_FILE_BACKUP=y</i>	(Optional) Specify this variable only when you perform a single file backup.
<i>NDR_DSA_NODE</i>	When this NDR_DSA_NODE is used, nsrdsa_save binary is started on the specified sn-server during the backup. During recover, nsrdsa_recover binary is started on the specified sn-server.
<i>NSR_DSA_REC_SERVER</i>	During recover, nsrdsa_recover binary is started on the specified remote storage node.

6. Click **Next**.
7. On the **Select the NetWorker Client Properties** window:
 - a. In the **Priority** field, specify the order in which the NetWorker server contacts clients in a protection group for backup. The attribute can contain a value between 1 and 1,000. The lower the value, the higher the priority.
 - b. In the **Parallelism** attribute:

- For Direct-NDMP, set the **Parallelism** attribute to 1.
 - For NDMP-DSA, the parallelism value depends on the NAS capabilities and set parallelism to a value that is appropriate for the NAS. Parallelism values of 4 to 8 are common. In general, the best parallelism setting depends on filer configuration and the amount of installed RAM.
 - For Celerra using DartOS v.5 and earlier, the **Parallelism** attribute cannot exceed 4.
 - For VNX using DartOS v.6 and later, the maximum parallelism value is 8. The optimal parallelism value depends on the following factors:
 - Amount of physical memory on the Data Mover.
 - Amount of physical memory allocated to the NDMP PAX configuration.
 - Value defined for the `concurrentDataStreams` parameter on the filer.
- c. In the **Remote Access** attribute:
- Specify the root account on Linux/UNIX, and/or the administrator account on Windows, of any computer that you use to browse backups of the NAS. Specify each account on a separate line. For example:
- ```
administrator@windows_hostname
root@linux_hostname
```
- d. Select the **Data Domain Interface**. This option specifies the protocol to use if you send the backup data to a Data Domain Device. Available selections are Any, Fibre Channel, or IP.
- e. Do not select the **Block Based Backup** or **Client Direct** options, as they do not apply to NDMP backups.
8. On the **Specify the File System Objects** window, select or specify the objects to backup:
- Celerra, VNX, and VNXe do not support the NDMP Snapshot Management extension, and do not support browsing. Specify the savesets to backup.
    - To back up all the file systems on the client, type `ALL`.
    - When you do not use the `ALL` save set, specify the file system name, as configured on the NAS.
    - File system names in the **Save set** field are case sensitive.
    - You cannot specify a share name.
  - For Celerra backups, do not use the `ALL` save set. List the file systems, one per line excluding the root, or `/` file system. When you include the root file system, client index updates fail for hidden file systems (directories that start with a `."`) with the error:
 

```
Failed to store index entries.
```
  - For Celerra block-level backups, specify the entire file system mount point.
- To back up large client file systems, optionally schedule each file system to back up separately. For example, create two separate clients with the same name, but with different save sets.
9. Click **Next**.
10. On the **Client Configuration Summary** window, review the attributes, and then click **Create**.
11. On the **Client Configuration Results** window, review the results, and then click **Finish** to exit the wizard.

Troubleshooting NDMP configuration and backup failures for PowerStore, Unity, Celerra, VNX, and VNXe describes how to resolve errors that you may experience when you configure the NDMP client.

## Celerra and VNX Environment variables

Restart the NetWorker services after setting the following environment variables:

**Table 12** Celerra and VNX Environment variables

| Variables                                                                                                                                                                | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>NSR_NDMP_TMP_DIR=&lt;path to tmp&gt;</i>                                                                                                                              | Use this variable to redirect the tmp files going to /nsr/tmp (optional).<br><b>Note:</b> Path to tmp should be created manually.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <i>NSR_NDMP_RESTORE_LIMIT</i>                                                                                                                                            | Use this variable while recovering large NDMP savesets using DAR recovery. The <i>NSR_NDMP_RESTORE_LIMIT</i> environment variable is used to limit memory consumption during recoveries involving a large number of index entries (millions). This is specifically useful if there is not enough swap space or memory for the number of index entries selected for the recovery. If the variable is not set, the recovery might fail with an out of memory error message.<br><b>Note:</b> <i>NSR_NDMP_RESTORE_LIMIT</i> can be set to 50,000 or 1,000,000, but less than 3,000,000. |
| <i>NSR_NDMP_DDAR=y</i>                                                                                                                                                   | (Optional) Use this variable to recover the permissions of the parent directories of selected files. Specify this variable to perform a DDAR recovery when: <ul style="list-style-type: none"> <li>You set the <i>DIRECT=y</i> variable during the backup.</li> <li>The DART version is 5.5 and later.</li> </ul> <b>Note:</b> Do not specify <i>NSR_NDMP_DDAR</i> when you also use <i>NSR_NDMP_RECOVER_DIR</i> .                                                                                                                                                                  |
| <i>NSR_NDMP_RECOVER_NO_DAR=y</i>                                                                                                                                         | (Optional) Specify this variable to perform a non-DAR recovery when you set the <i>DIRECT=y</i> variable during the backup.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <i>NSR_NDMP_RECOVER_DIR=y</i><br><br>This environment variable must be set in the operating system before invoking either the <i>recover</i> or <i>winworkr</i> program. | (Optional) Specify this variable to perform a DAR recovery when: <ul style="list-style-type: none"> <li>When you set the <i>DIRECT=y</i> variable during the backup.</li> <li>The DART version is 5.5 and later.</li> </ul> <b>Note:</b> Do not use <i>NSR_NDMP_RECOVER_DIR</i> when you also use <i>NSR_NDMP_DDAR</i> .                                                                                                                                                                                                                                                            |

**Table 12** Celerra and VNX Environment variables (continued)

| Variables              | Definition                                                                                                                                           |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>NSR_NDMP_NDAR=y</i> | It can be used instead of <i>NSR_NDMP_DDAR</i> when both full and incremental backups are present. <i>NSR_NDMP_DDAR</i> works only with full backups |
| <i>TS=y</i>            | (Optional) Enables tape silvering.                                                                                                                   |

**Note:** To set environment variable on Linux NetWorker server edit `/opt/nsr/admin/networkerrc` and add the variable as shown below: ,

- `NSR_NDMP_DDAR=y`
- `export NSR_NDMP_DDAR`

To set environment variable on Windows NetWorker server navigate to Control Panel > System and Security > System > Advanced system settings > Environment Variables > System variables > New.

- Variable Name = `NSR_NDMP_TMP_DIR`
- Variable Value = `C:\NW\nsr\tmp`

## Performing post Client Configuration Wizard steps

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

### Modifying the Storage Node

#### About this task

On the **Globals (2 of 2)** tab, specify the storage node in the **Storage Nodes** attribute.

The attribute value depends on the type of backup:

- When you perform Direct-NDMP backups with NDMP devices, specify the hostname of the NAS that manages the tape device or autochanger.
- For three-party backups, specify the destination host first.
- For NDMP-DSA backups, specify the hostname of the storage node that manages the tape device or autochanger. If the NetWorker server is the storage node, specify *nsrserverhost*.

**NOTICE** For NDMP-DSA backups, the NetWorker software uses the **Storage Node** attribute field of the NDMP client to determine which host receives the backup data. The `nsrndmp_save` command does not require the `-M` and `-P` options. If you specify the `-M` and `-P` options, they override the **Storage Node** attribute value.

## Adding NDMP Client Properties

### Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Clients**.
3. Right-click a **Client**, and then select **New Client Properties**. The **Client Properties** screen appears.
4. Select the **Apps & Modules** tab. Select the **NDMP** attributes of the **Client**. The **NDMP** options are as follows:



- **NDMP**—Select this box to indicate whether this client is an NDMP client.
- **NDMP multistreams enabled**—Do not select this box. Only Isilon supports multistreaming.
- **NDMP log successful file recovery**—By default, NetWorker does not print each successfully recovered file name in the log messages, because this logging impacts performance and takes up space. To enable the logging of successful recoveries for each file, select this checkbox.
- **Disable IPv6**—Check this box to disable IPv6 on NDMP backup and recovery.
- **NDMP array name**—This name is the logical name that is assigned to the array in NDMP NAS array configurations.
- **NDMP vendor information**—This attribute contains NDMP client vendor information.

5. Click **OK**.

## Configuring the NDMP client manually

It is recommended that you create Network Data Management Protocol (NDMP) clients by using the Client Configuration wizard. If you create the NDMP client manually, then the configuration details for each attribute in the Client Configuration wizard apply when you create the client manually.

Review this information before you configure an NDMP client manually:

- For an NDMP configuration that includes Storage Node resources, configure a Client resource for each storage node that you define for an NDMP backup and clone operation.
- For NDMP three-party storage nodes that use NDMP devices, repeat these steps for each NDMP storage node.
- For NDMP-DSA storage nodes, create the NetWorker Client resources in the same manner as non-NDMP clients. The *NetWorker Administration Guide* provides details on how to create a non-NDMP Client resource.
- NDMP does not support the use of directives including AES encryption. The NetWorker software ignores any value that you define in the **Directives** attribute for an NDMP client.
- When you select **Checkpoint enabled** on the **General** tab, do not modify the **Checkpoint granularity** attribute. NDMP backups do not support checkpoint granularity and the NetWorker software ignores any value that you define for this attribute.
- If the NAS supports the NDMP snapshot management extension, then you can browse and mark individual file systems for backup instead of specifying the save sets in the **Save set** attribute. You cannot use the **Save set browse** icon to browse the NDMP file system until you:
  - Select the **NDMP** checkbox, on the **Apps & Modules** tab.
  - Specify the NDMP username and password in the **Remote user and password** fields on the **Apps and Modules** tab.

### **Note:**

- VNX does not support checkpoint restart.
- Celerra, VNX, VNXe, and NetApp C-mode do not support Snapshot Management Extension. Only NetApp 7-Mode supports Snapshot Management Extension.
- Isilon, Celerra, VNX, VNXe, and NetApp C-mode filers do not allow you to browse. Only NetApp 7-Mode allows you to browse.

## Performing manual NDMP backups

After you configure the NetWorker server for NDMP backup data operations, you can perform manual NDMP backups.

On Windows, you can manually back up NDMP data by using the NetWorker User program. The method to backup NDMP data is the same as a non-NDMP local backup. You cannot perform a three-party backup with the NetWorker User program.

On Windows and UNIX, you can perform a manual backup from a command prompt by using the `nsrndmp_save` command.

Before performing a manual backup by using the `nsrndmp_save` command or the NetWorker User program, review these requirements:

- You can only perform manual Direct-NDMP backups from a NetWorker server.
- You can start a manual NDMP-DSA backup from a NetWorker server, storage node, or client. When you do not start the NDMP-DSA backup from the NetWorker server, the `servers` file on the NetWorker server and storage node, must contain the hostname of the host that initiates the backup.
- Before you perform a manual backup, you must configure the NDMP client on the NetWorker server. Manual backups use client configuration information for example, the variables that are defined in the **Application Information** attribute of an NDMP client.
- Direct-NDMP and three-party NDMP backups support manual DAR backups when the NDMP client contains the `DIRECT=Y` and `HIST=Y` environment variables in the **Application Information** attribute for the NDMP client.

 **NOTICE** To use DAR, the NAS filer must use NDMP version 4. The *NetWorker E-LAB Navigator* describes how to determine if a particular NDMP vendor supports DAR.

## Performing an NDMP backup from the command line

### About this task

Use the `nsrndmp_save` command to perform a manual command line NDMP backup.

The `nsrndmp_save` command does not back up the bootstrap. Without the bootstrap, you cannot perform a disaster recovery of the NetWorker server. To back up the bootstrap, run the `nsrpolicy -G nsrpolicy policy_name start` command from the NetWorker server. The `nsrpolicy` command uses the attribute values specified for the policy. For example, the pool and schedule values.

To perform an NDMP backup from the command prompt, use the following syntax:

```
nsrndmp_save -T backup_type -s NetWorker_servername -c clientname -l
backup_level -t date_time -g nsrpolicy_path
```

where:

- *backup\_type* is a supported backup type for the NAS filer. The optimal backup types for Celerra NAS are `tar` or `dump`. However, you can use the `vbb` backup type to back up entire volumes at the block level.
- *backup\_level* is a full for a full backup, `incr` for an incremental backup. Each NAS supports full backups.
 

 **Note:** NetWorker only supports or allows Full, Incremental, and Differential level 1 NDMP backups of all NAS.
- *date\_time* is the date and time of the last backup, which is enclosed in double quotes. Specify this value for `incr` level backups. When you do not specify the date and time, the backup is a native NDMP level-based backup.

**NOTICE** During a NetWorker scheduled policy backup, the NetWorker software supplies the date and the time information, and incremental and level backups work as expected.

Use one of these methods to determine the date and time of the last NDMP backup:

- Review the `daemon.raw` file on the NetWorker server or the savegroup completion report for a line similar to the following:

```
42920:nsrndmp_save: browsable savetime=1296694621
```

Use the value after `savetime=` with the `-t` option.

- Specify the date and time of the last backup reported by the `mminfo` command for the NDMP save set.

**Note:** You can force the NDMP backup and recovery to ignore IPv6 and instead use IPv4 in one of two ways:

- Add the `-f` option to the `nsrndmp_save` or the `nsrndmp_recover` commands as required.
- On the Apps & Modules tab in the **Client Properties** window, select **Disable IPv6**.

## Example of an NDMP backup

To perform an incremental backup of a NetApp client that is named `mynetapp`, perform the following steps:

- Determine the time of the last full backup:

```
mminfo -v -c mynetapp
```

**Table 13** NDMP backup

| client   | date     | time     | size   | ssid       | fl   | lvl  | name      |
|----------|----------|----------|--------|------------|------|------|-----------|
| mynetapp | 08/16/15 | 15:23:58 | 1853MB | 3864812701 | cbNs | full | /.../set1 |
| mynetapp | 08/17/15 | 15:39:58 | 815MB  | 3848036430 | cbNs | incr | /.../set2 |

- Specify the last backup time in `nsrndmp_save` command:

```
nsrndmp_save -T dump -s my_nwserver -c mynetapp -l incr -t "02/16/11
15:23:58" -g mygroup path
```

For NDMP-DSA backups, the NetWorker software uses the Storage Node attribute field of the NDMP client to determine which host receives the backup data. The `nsrndmp_save` command does not require the `-M` and `-P` options. If you specify the `-M` and `-P` options, they override the **Storage Node** attribute value. The *NetWorker Command Reference Guide* and the `nsrndmp_save` man page on UNIX provide more information.

## Troubleshooting NDMP configuration and backup failures for PowerStore, Unity, Celerra, VNX, and VNXe

The following issues are a list of the possible causes and the resolutions for Network Data Management Protocol (NDMP) backup failures.

### Unable to connect to NDMP host *hostname*

This message appears when the NetWorker server cannot create or modify an NDMP client.

To resolve this issue ensure that the environment meets the following requirements:

- Username and password specified for the client is correct and has sufficient permissions to perform NDMP operations.
- NDMP service is running on the filer.

### NDMP Service Error: SnapSure file system creation fails

NDMP backup fails immediately with the error `nsrndmp_save: NDMP Service Error: SnapSure file system creation fails` due to the following reasons:

- There is not enough free disk space on the NAS to build the SnapSure checkpoint.
- The file system specified for backup is incorrect or does not exist.
- SnapSure creation fails if it does not get created within the specified timeout period. The default SnapSure timeout on the Celerra or VNX data mover is five minutes and can be extended up to thirty minutes.
- SnapSure creation fails if enabled during NDMP backup of a checkpoint file system.
- SnapSure creation fails due to a code defect in DartOS.
- SnapSure creation fails due to file system corruption or inode inconsistencies on the NAS disk back-end.

#### Resolution

Perform the following to troubleshoot SnapSure creation failures:

- Collect the Celerra or VNX data mover server log immediately following the failed NDMP backup. The log should provide a descriptive reason for the SnapSure creation failure.
- SnapSure can be disabled by removing the application information attribute "SNAPSURE=y" in the NDMP Client, or by setting it to "no" - "SNAPSURE=n".
- SnapSure can be set as enabled on the Celerra or VNX data mover using the command `server_param server_2 -f NDMP -m snapsure -v 1 .`
- SnapSure can be disabled on the data mover using the command `server_param server_2 -f NDMP -m snapsure -v 0 .`
- The setting in the NDMP client application information attributes for "SNAPSURE=y/n" overrides the setting on the Celerra or VNX data mover.

### Cannot perform NDMP backup after the NetWorker server licenses expire

If a NetWorker server running in evaluation mode expires before you authorize the server, NDMP devices remain disabled after the addition of the required licenses and authorization of the NetWorker server.

To re-enable NDMP devices, perform the following steps:

1. To connect to the NetWorker server, use NMC, and then click the **Devices** button.
2. In the **Devices** windows, right-click the NDMP device, and then select **Properties**.
3. Click the **Configuration** tab, and then set the **Target Sessions** attribute to **1**.
4. Click the **General** tab, and then in the **Enabled** section, select **Yes**.
5. Click **Ok**.

## No PAX threads available

This error message appears in the `server_log` on the NDMP Data Server when the client parallelism value for a Celerra client exceeds what the Celerra can support.

To resolve this issue, adjust the client parallelism attribute to a value that the Celerra supports:

- For a Celerra client that runs DartOS v5.0 or earlier, the client parallelism value cannot exceed 4.
- For a Celerra client that runs DartOS v6.0, the maximum parallelism value supported is 8, or the value defined in the `concurrentDataStreams` variable on the Celerra. By default, the `concurrentDataStreams` variable is 4.
- The maximum parallelism value also depends on the available physical memory and the amount of memory allocated to the PAX configuration. Configuring NDMP Backups on Celerra on the Support website provides more information.

## Failed to store index entries

This error message occurs in the `daemon.raw` file when an index backups fails due to an insufficient amount of swap space.

To resolve this issue, increase the amount of swap space available to the NetWorker server.

 **NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

## IO\_WritePage write failed - No space left on device (28): No space left on device

This error message appears in the `daemon.raw` file when the index backup fails. There is insufficient temporary space to store the index entries before the NetWorker software commits the information into the client file index.

To resolve this issue, specify a new the temp directory with sufficient disk space in one of the following ways:

- Define the `NSR_NDMP_TMP_DIR` environment variable in the Application Information attribute of the client.
- Define the `NSR_NDMP_TMP_DIR` as an operating system environment variable on the NetWorker server.

[Memory and space requirements for NDMP FH updates](#) on page 23 describes how to determine the amount of disk space the NetWorker software requires to temporarily store client files index entries.

 **NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

## Error reading the FH entries from save through stdin

This error message appears in the `daemon.raw` file of the NetWorker server when there is a communication error between the `nsrndmp_save` and `nsrndmp_2fh` processes.

Resolve any communication or connection issues, then retry the backup.

 **NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

Cannot find file history info for file name...You may still be able to recover this file with a save set recovery

This error message appears in the `daemon.raw` file of the NetWorker server when file history (FH) information is missing or corrupted for the file that is specified in the error message. For example, NetWorker cannot update the client file index (CFI) with FH information when a backup process interruption occurs during the failover of a clustered NetWorker environment.

You cannot perform an Network Data Management Protocol (NDMP) file-by-file recover or a save set recover when the CFI does not contain the associated FH information.

To recover this file, perform a save set recover from the command prompt. [Performing an NDMP save set recovery from the command prompt](#) on page 105 provides for further information.

 **NOTICE** The NetWorker server does not delete the FH files that are stored in the `tmp` directory when the CFI updates fail.

`nsrndmp_save: data connect: failed to establish connection`

This error message appears in the `daemon.raw` file of the NetWorker server for several reasons:

- Network connectivity or name resolution issues exist between the NetWorker server and the NDMP client.
- You specified an incorrect NDMP username or password specified for the NDMP client.
- The NDMP service is not started on the NAS filer.
- The NetWorker server cannot communicate with the NAS filer over port 10000.
- A free port in the NetWorker server's default port range (7937-9936) is not available during an NDMP-DSA backup.  
The *NetWorker Security Configuration Guide* provides more information about NDMP port requirements and configuration.
- A misconfigured loop router. For a Celerra filer, the server route command utility configures the loop router. For NetApp, the route utility configures loop back router. The value of this setup is network-specific and depends on the number of switches and hubs between the NAS filer, NetWorker server, and NetWorker storage node.
- On the host where DSA is running, if the hostname is present in the hosts file, the `nsrdsa_save` process uses this name during backup. The DSA host passes the loopback entry to the NDMP data server and the connection fails. To resolve this issue, remove the hostname from the localhost list.

Knowledge base articles on the Support website provides detailed troubleshooting information for this error message and other failed to establish connection failures that you might encounter during an NDMP backup.

## Monitoring NetWorker Server activities in the Administration window

The **Monitoring** window in the NetWorker **Administration** application enables you to monitor the activities of an individual NetWorker Server.

### About this task

The **Monitoring** window provides the following types of activity and status information:

- Data protection policies, workflows, and individual actions.
- Cloning, recovering, synthetic full backups, and browsing of client file indexes.
- Operations that are related to devices and jukeboxes.
- Alerts and log messages.

You can also perform some management operations from the **Monitoring** window, for example, starting, stopping, or restarting a data protection policy.

#### Procedure

1. From the **NMC Console** window, click **Enterprise**.
2. In the **Enterprise** view, right-click the NetWorker Server, and then select **Launch Application**.  
The **Administration** window appears.
3. To view the **Monitoring** window, click **Monitoring**.

## About the Monitoring window

On the **Administration** window taskbar, select **Monitoring** to view the details of current NetWorker server activities and status, such as:

- Policies and actions.
- Cloning, recovering, synthetic backups, checkpoint restart backups, and browsing of client file indexes.
- Alerts and log messages, and operations that are related to devices and jukeboxes.

While the **Monitoring** window is used primarily to monitor NetWorker server activities, it can also be used to perform certain operations. These operations include starting, stopping, or restarting a workflow.

The **Monitoring** window includes a docking panel that displays specific types of information. Select the types of information you want to view from the docking panel.

A portion of the **Monitoring** window, which is known as the task monitoring area, is always visible across all windows. A splitter separates the task monitoring area from the rest of the window. You can click and move the splitter to resize the task monitoring area. The arrow icon in the upper right corner of the **Monitoring** window allows you to select which tasks you want to appear in this view.

Smaller windows appear within the **Monitoring** window for each window. Each smaller window, once undocked, is a floating window and can be moved around the page to customize the view. You can select multiple types from the panel to create multiple floating windows that can be viewed simultaneously. The following table describes the various types of information available in the docking panel, and the details each one provides.

**Table 14** Monitoring window panel

| Window                  | Information provided                                                                                                                                                                                                                                                                         |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Policies/Actions</b> | The <b>Policies</b> tab provides you with status information about all configure policies and the associated workflows and actions. The <b>Actions</b> tab provides you with status information for all actions. <a href="#">Policies/Actions pane</a> on page 81 provides more information. |
| <b>Sessions</b>         | Allows you to customize whether to display all session types, or only certain session types. The information that is provided depends on which session type you select. For example, if you select <b>Save Sessions</b> , the window lists clients, save sets, groups, backup level,         |

**Table 14** Monitoring window panel (continued)

| Window            | Information provided                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                   | backup start time, duration of the backup, devices, rate, and size. <a href="#">Sessions window</a> on page 85 provides more information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Alerts</b>     | Lists the priority, category, time, and message of any alerts. <a href="#">Alerts pane</a> on page 86 provides more information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Devices</b>    | Lists devices, device status, storage nodes, libraries, volumes, pools, and related messages. <a href="#">Devices pane</a> on page 87 provides more information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Operations</b> | <p>Lists the status of all library and silo operations, including <code>nsrjb</code> operations that are run from the command prompt. Also lists user input, libraries, origin, operation data, operation start time, duration of the operation, progress messages, and error messages.</p> <p>When displaying Show Details from the <b>Operations</b> window, the length of time that the window is displayed depends on the value that is typed in the <b>Operation Lifespan</b> attribute on the <b>Timers</b> tab of the <b>Properties</b> dialog box for the corresponding library. To access library properties, click <b>Devices</b> in the taskbar. By default, this pane is hidden.</p> |
| <b>Log</b>        | Lists messages that are generated by the NetWorker server, including the priority of each message, the time the message was generated, the source of the message, and the category. <a href="#">Log window</a> on page 90 provides more information.                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## Customizing the Monitoring window

This section describes how to customize the **Monitoring** window in the **Administration** interface.

### Customizing tables

You can customize the organization and display of tabular information in the **Monitoring** window.

#### Sorting tables

You can change the display of tabular information that appears in the window. You can sort Table grids by column heading, and then by alphabetic or numeric order within those columns.

1. Drag and drop the column heading to its new position.
2. Click the column heading to sort the items into alphabetic and numeric order. An arrow appears in the column heading to indicate the sort order.

#### Sorting selected rows in a table

Selected rows are sorted to the top of the table. This sorting is particularly useful when you select **Highlight All** from the Find panel to select all rows matching the Find criteria and then moving all selected rows to the top of the table to view the results.

1. From the **Edit** menu, select **Find**, or press **Ctrl + F** to view the **Find** panel.
2. To select the rows, click each row or use the Find criteria.
3. Select **Sort Selected**.



### Sorting multiple columns in a table

You can select the column that you want to use as the tertiary sort key, the secondary sort key, and the primary sort key.

1. Click the column that you want to use as the last sort key.
2. Click the column that you want to use as the next-to-last sort key, and so on, until you select the primary column.

### Displaying columns in a table

You can select which columns to display in a table.

1. From the **View** menu, select **Choose Table Columns**.
2. Click a column name to select or clear the column and then click **OK**. You can also select the columns to display by right-clicking a table header and selecting **Add Column** from the drop-down.

## Displaying panes

You can choose to show or hide panes in the **Monitoring** window.

### About this task

Perform the following steps to hide or show a pane in the **Monitoring** window.

#### Procedure

1. From the **View** menu, select **Show**. A check mark appears beside the panes that appear in the **Monitoring** window.
2. To hide a pane, select a marked pane.  
A check mark does not appear beside the pane.
3. To show a pane, select an unmarked pane.  
A check mark appears beside the pane.

## Policies/Actions pane

The **Policies/Actions** pane provides you with the ability to review status information about policies and actions.

This pane has two tabs:

- **Policies**—Provides a navigation tree that displays all configured policies on the NetWorker Server. Expand each policy to display the workflows that are associated with each policy. Expand each workflow to display each action that is contained in the workflow.
- **Actions**—Provides a list of all Action resources.

### Policies pane

The **Monitoring** window in the **NetWorker Administration** window enables you to monitor activities for specific policies, workflows, and actions.

The **Policies/Actions** pane at the top of the **Monitoring** window lists the policies on the NetWorker Server by default. Click the + (plus) sign next to a policy in the list to view the workflows in the policy, and the + (plus) sign next to a workflow to view the actions for a workflow.

The **Policies** pane provides the following information for each item (where applicable):

- Overall status

The following table provides details on the status icons that may appear in the **Policies** pane.

**Table 15** Policy status icons

| Icon                                                                              | Status                                                |
|-----------------------------------------------------------------------------------|-------------------------------------------------------|
|  | Never run                                             |
|  | Running                                               |
|  | Succeeded                                             |
|  | Failed                                                |
|  | Probing                                               |
|  | Interrupted                                           |
|  | Queued                                                |
|  | Cloning                                               |
|  | Consolidating (NetWorker Server 8.2.x and lower only) |

 **Note:** When the schedule for an action is skip, the status of the action appears as Never Run and the status of the Workflow is Succeeded.

- Most recent start time.
- Duration of the most recent run.
- Next scheduled runtime.
- Name of the assigned save set.
- Device on which the save set is stored.
- Backup level.
- Data transfer rate.
- Size of the save set.
- Messages that resulted from an action.

Right-click an action in the **Policies** pane and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

When you sort the items on the **Policies/Actions** pane by using the **Status** column, NetWorker sorts the items in alphabetical order that is based on the label of the icon.

Consider the following when a policy/action is in a probing state:

- A message is sent when the group starts and finishes the probe operation.
- The results of the probe operation (run backup/do not run backup) are also logged.
- Probes do not affect the final status of the group, and the group status does not indicate the results of the probe.
- If probing indicates that a backup should not run, then the group status reverts to its state before the group running.

- Check the results of the probe in the **Log** window to ensure that the probe indicates that the backup can be taken.

### Actions pane

To view a list of all actions, click the **Actions** tab at the bottom of the **Policies** pane. The **Policies** pane becomes the **Actions** pane.

The **Actions** pane provides the following information for each action:

- Overall status

 **Note:** The **Actions** pane displays the same status icons as the **Policies** pane.

- Name
- Assigned policy
- Assigned workflow
- Type
- Date and time of the most recent run
- Duration of the most recent run
- Percent complete, for actions that are in progress
- Next scheduled runtime

Right-click an action in the **Actions** pane and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

## Workflow operations

This section describes how to use the **Monitoring** window to start, stop, and restart workflows.

### Starting, stopping, and restarting policies

The workflows in a policy can run automatically, based on a schedule. You can also manually start, stop, and restart specific workflows by using the the NMC **NetWorker Administration Monitoring** window.

#### About this task

You can restart any failed or canceled workflow. Note, however, that the restart must occur within the restart window that you specified for the workflow. Additionally, for a VMware backup, if you cancel a workflow from **NetWorker Administration** and then want to restart the backup, ensure that you restart the workflow from the **NetWorker Administration** window. If a workflow that was started from **NetWorker Administration** is restarted from the **vSphere Web Client**, the backup fails.

#### Procedure

1. In the **Monitoring** window, select the workflow or actions.
2. Right-click and then select **Start**, **Stop**, or **Restart**.

A confirmation message appears.

 **Note:** You cannot stop, restart, or start individual actions.

3. Click **Yes**.

## Viewing workflow backup details

Perform the following steps to view backup details for workflows.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Policies** in the docking panel, and expand the Policy that you want to monitor.
3. Right-click the workflow, and then select **Show Details**. The **Workflow Summary** window appears.
4. In the **Workflow runs** pane of the **Workflow Summary** window, select the workflow.
5. Click **Show Messages**. In the **Show Messages** window, select one of the following options:
  - Get Full Log—To display all messages.
  - Print—To print the log.
  - Save—To save the log to a local file.
  - OK—To close the **Show Messages** window.
6. Click **OK** to close the **Workflow Summary** window.

## Viewing action backup details

Perform the following steps to view backup details for actions.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Actions** in the docking panel.
3. In the **Actions** pane, right-click the action, and then select **Show Details**. The details window for the action appears.
4. Review the information in the **Actions Messages** pane. To display detailed information from the action log file, click **Show Action Logs**, and then select one of the following options:
  - Get Full Log—To display all messages.
  - Print—To print the log.
  - Save—To save the log to a local file.
  - OK—To close the **Show Messages** window.
5. In one of the Actions detail panes, for example, the **Completed successfully** pane, select the action that you want to review.
6. Click **Show Messages**. In the **Show Messages** window, select one of the following options:
  - Get Full Log—To display all messages.
  - Print—To print the log.
  - Save—To save the log to a local file.
  - OK—To close the **Show Messages** window.
7. Click **OK** to close the **Details** window.

## Sessions window

Use the **Sessions** window to view the sessions that are running on a NetWorker server. You can change the view of this window to display these sessions:

The **Sessions** pane below the **Policies/Actions** pane provides details on individual save, recover, clone, and synthetic full sessions by client.

To view all sessions or to limit the list of sessions by the session type, click the tabs at the bottom of the **Sessions** pane. Session types include:

- Save
- Recover
- Clone
- Browse
- Synthetic Full/Rehydrated Sessions
- All

To change the displayed session types go to **View > Show**, and select the type of sessions to display. To display all sessions currently running on the NetWorker Server, regardless of type, select **All Sessions**.

You can stop a session (backup, synthetic full backup, clone, and recovery sessions) from the **Monitoring** window, even if the session was started by running the `nsrworkflow` command.

To stop a session, right-click the workflow in the pane, and select **Stop** from the list box.

### Changing displayed session types

The column headings that are displayed on this window differ depending on the type of sessions you chose to display.

#### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Sessions** in the docking panel.
3. Select **View > Show** and then select the type of sessions to display. To display all sessions currently running on the NetWorker server, regardless of type, select **All Sessions**.

### Stopping a session

You can stop a session (backup, synthetic full backup, clone, and recovery sessions) from the **Monitoring** window, even if the session was started by running `savegrp`.

To stop a session, right-click the session in the window and select **Stop** from the drop-down.

The following table provides a list of actions that can be stopped from NMC.

**Table 16** Sessions that can be stopped from NMC

| Session type                | Stop from NMC? |
|-----------------------------|----------------|
| Save by Savegroup           | Yes            |
| Synthetic Full by Savegroup | Yes            |
| Clone by Savegroup          | Yes            |
| Schedule Clone              | Yes            |

**Table 16** Sessions that can be stopped from NMC (continued)

| Session type                             | Stop from NMC? |
|------------------------------------------|----------------|
| Manual Save                              | No             |
| Manual Clone via NMC                     | No             |
| Manual Clone via CLI                     | No             |
| Winworker and CLI Recovery               | No             |
| Recovery started from Recover wizard     | Yes            |
| VMware Backup Appliance Save and Recover | No             |

 **NOTICE** Stopping a session from NMC does not affect any other group operations running.

## Alerts pane

The **Alerts** pane displays alerts that are generated by a particular NetWorker server or Data Domain system that has devices that are configured on the NetWorker server. The **Alerts** pane includes priority, category, time, and message information.

An icon represents the priority of the alert. The following table lists and describes each icon.

**Table 17** Alerts window icons

| Icon                                                                                | Label        | Description                                                                                                                           |
|-------------------------------------------------------------------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
|  | Alert        | Error condition detected by the NetWorker server that should be fixed by a qualified operator.                                        |
|  | Critical     | Severe error condition that demands immediate attention.                                                                              |
|  | Emergency    | Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority. |
|  | Information  | Information about the current state of the server. This icon represents the lowest priority.                                          |
|  | Notification | Important information.                                                                                                                |
|  | Waiting      | The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.                                           |
|  | Warning      | A non-fatal error has occurred.                                                                                                       |

When items on the **Alerts** pane are sorted by the **Priority** column, they are sorted in alphabetical order based on the label of the icon.

## Removing alerts

Remove individual alert messages from the **Events** tables by removing them from the **Events** table. To delete a message in the **Events** table, right-click the message, and select **Dismiss**.

### About this task

 **Note:** The alert message remains in the **Log** window in the NetWorker **Administration** program.

## Devices pane

The **Devices** pane allows you to monitor the status of all devices, including NDMP devices. If the NetWorker server uses shared and logical devices, the window is adjusted dynamically to present a set of columns appropriate for the current configuration.

The **Devices** pane provides the following information:

- Status of the operation.
- Name of the device.
- Name of the storage node that contains the device.
- For tape devices, the name of the library that contains the device.
- Name of the volume in the device.
- Name of the pool that is associated with the volume.
- Last message generated for the device.
- Whether the operation requires user input.

For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape.

[Entering user input](#) on page 89 provides instructions on how to deal with a user input notification.

If the current server configuration includes a shared device, a **Shared Device Name** column appears on the **Devices** pane. The name of the shared device appears in the **Shared Device Name** column. If other devices for that configuration are not shared devices, then the **Shared Device Name** column is blank for those devices. Only a single device per hardware ID can be active at any particular moment. The information for inactive shared devices is filtered out, and as a result, only one device per hardware ID is presented on the window at any time.

An icon represents the device status. The following table lists and describes each icon.

**Table 18** Devices status icons

| Icon                                                                                | Label                       | Description                         |
|-------------------------------------------------------------------------------------|-----------------------------|-------------------------------------|
|  | Library device active       | The library device is active.       |
|  | Library device disabled     | The library device is disabled.     |
|  | Library device idle         | The library device is idle.         |
|  | Stand-alone device active   | The stand-alone device is active.   |
|  | Stand-alone device disabled | The stand-alone device is disabled. |

**Table 18** Devices status icons (continued)

| Icon                                                                              | Label                   | Description                     |
|-----------------------------------------------------------------------------------|-------------------------|---------------------------------|
|  | Stand-alone device idle | The stand-alone device is idle. |

When you sort items in the **Devices** pane by the **Status** column, NetWorker sorts the devices in alphabetical order based on the label name of the icon.

## Operations window

The **Operations** window displays information about device operations. It provides the following information:

- Status of the operation.
  - Name of the library.
  - Whether the operation requires user input.  
For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape. [Entering user input](#) on page 89 provides instructions on how to deal with a user input notification.
  - The origin, or source, of the operation.  
For example, the interface, nsrjb or the NetWorker server.
  - Time the operation started.
  - Type of operation.
  - Duration of the operation.
  - Status messages from the operation.
  - Any error messages.
-  **NOTICE** Only the last error message of the operation appears in the **Error Messages** column. Move the mouse pointer over the cell containing the last error message to display the entire list of error messages.

The operation status is represented by an icon. The following table lists and describes each of the icons.

**Table 19** Operations window icons

| Icon                                                                                | Label      | Description                                          |
|-------------------------------------------------------------------------------------|------------|------------------------------------------------------|
|  | Failed     | The operation failed.                                |
|  | Queued     | The operation is waiting in the queue to run.        |
|  | Retry      | The operation failed, but may work if you try again. |
|  | Running    | The operation is running.                            |
|  | Successful | The operation completed successfully.                |
|  | User Input | The operation requires user input.                   |



When items on the **Operations** window are sorted by the Status column, they are sorted in alphabetical order based on the label of the icon.

## Viewing operation details

The **Operation Details** dialog box opens, providing information about the completion of the operation. The **Completion Time** displays the time that the operation finished. The time that it took to complete the operation is the difference between the completion and start times of the operation.

### About this task

To save operation details to a file, click **Save** in the **Operation Details** dialog box. When prompted, identify a name and location for the file.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation, then select **Show Details**.

## Stopping an operation

Certain operations can be stopped from the **Operations** window.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation to stop, then select **Stop**.
4. Click **Yes** to confirm the stop.

 **Note:** Operations that were started from a command line program, such as the `nsrjb` command, cannot be stopped from the **Operations** window. To stop these operations, press `Ctrl-C` from the window where the command was started.

## Entering user input

If the system requires user input, select the labeling operation in slow/verbose mode and the **Supply User Input** icon appears.

### Procedure

1. Right-click the operation, then select **Supply Input**.
2. Confirm the requirement to supply input.

- If **Yes**, and input is supplied, the icon in the **User Input** column disappears.

 **Note:** If two users try to respond to the same user input prompt, the input of the first user takes precedence, and the second user receives an error message.

- If **No**, and input is not supplied, the operation will time out and fail.

## Log window

To view the most recent notification logs, click the **Log** window from the docking panel in the **Monitoring** window. The **Log** window provides the priority, time, source, category, and message for each log.

 **Note:** If a particular log file is no longer available, check the log file on the NetWorker server. The log files are located in `NetWorker_install_path\logs` directory.

An icon represents the priority of the log entry. The following table lists and describes each icon.

**Table 20** Icons in the Log pane

| Icon                                                                                | Label        | Description                                                                                                                           |
|-------------------------------------------------------------------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
|    | Alert        | Error condition that is detected by the NetWorker server that should be fixed by a qualified operator.                                |
|    | Critical     | Severe error condition that demands immediate attention.                                                                              |
|    | Emergency    | Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority. |
|   | Information  | Information about the current state of the server. This icon represents the lowest priority.                                          |
|  | Notification | Important information.                                                                                                                |
|  | Waiting      | The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.                                           |
|  | Warning      | Non-fatal error has occurred.                                                                                                         |

When you sort items on the **Log** pane by using the **Priority** column, NetWorker sorts the icons in alphabetical order based on the name of the label.

## Recover window

The **Recover** window displays information about recover configurations that are created with the NetWorker Management Console (NMC) Recovery wizard.

You can use this window to:

- Start the NMC Recovery wizard to create recover configurations or modify saved recover configurations.
- Identify the status of a recover configuration that is created with the NMC Recovery wizard.
- Start and stop a recover job.

The **Recover** window is divided into five sections:

- **Toolbar**—The toolbar is hidden by default. To display the recovery toolbar, select **View > Show toolbar**
- **Summary**
- **Configured Recovers**
- **Currently Running**

A splitter separates the **Configured Recovers** section from **Currently running** window. You can click and move the splitter to resize these two windows.

## Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. The following table summarizes the function of each toolbar button.

**Table 21** Recovery toolbar options

| Button                                                                              | Function                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|    | Starts the NMC <b>Recover</b> wizard to create recover configurations.                                                                                                                        |
|    | Displays the <b>Properties</b> window for the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                          |
|    | Deletes the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                                                            |
|    | Displays online help for the <b>Recover</b> window.                                                                                                                                           |
|   | Displays the <b>Find</b> window at the bottom of the <b>Recover</b> window. The <b>Find</b> window allows you to perform keyword searches for messages that appear in the <b>Logs</b> window. |
|  | Start the recover operation for a selected saved recover configuration. This option is only available for a recover configuration that has a Never run, or Failed status.                     |
|  | Stop in-progress recover operation that you selected in the <b>Currently Running</b> window.                                                                                                  |

 **Note:** The **Recover** toolbar does not appear by default. To display the **Recover** toolbar, select **View > Show toolbar**.

## Recover Summary

The Recover Summary section displays a high-level overview of recover jobs.

This section includes the following information:

- **Total Recovers**—The total number of successful recover jobs.
- **Since**—The number of successful recover jobs since this date.

## Configured Recovers

The **Configured Recovers** window displays a list of saved recover configurations in a tabular format. You can sort the information by column. The **Configured Recovers** table displays the following information for each saved recover configuration:

- **Status**—The job status of a saved recover configuration.
- **Name**
- **Source client**
- **Destination client**

- Recovery list
- Recover type—For example, file system or BBB.
- Comment
- OS—The operating system of the source host.
- Recover requestor—The Windows or UNIX account used to create the recover configuration.
- Start Time
- End Time
- Start date

**Table 22** Save recover configuration job status

| Icon                                                                               | Description                                        |
|------------------------------------------------------------------------------------|----------------------------------------------------|
|   | The last recover attempt failed.                   |
|   | The last recover attempt completed successfully.   |
|   | The recover job has never run.                     |
|   | The recover job is scheduled to run in the future. |
|  | The recover job has expired.                       |

## Currently running

The **Currently Running** window displays a list of in progress recover jobs in a tabular format. You can sort the information by column. The **Currently Running** table displays the following information for each job:

- Status
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB
- Volume
- Comment
- Device
- Size
- Total size
- % complete
- Rate (KB/s)
- Start time
- Duration

- Currently running

## Find

The **Find** section appears along the bottom of the **Recover** window, after you select the **Find** button on the **Recover** toolbar. **Find** allows you to search for keywords in the **Configured Recovers** window. The following table summarizes the available find options.

**Table 23** Find options

| Find option          | Description                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Find</b>          | Highlight the first saved recover configuration that contains the specified keyword.                                                         |
| <b>Prev</b>          | Highlight the previous saved recover configuration that contains the specified keyword.                                                      |
| <b>Highlight All</b> | Highlights each saved recover configuration that contains the specified keyword.                                                             |
| <b>Sort Selected</b> | Sorts each highlighted recover configuration in the Configured Recover table so that they appear at the top of the Configured Recover table. |
| <b>Match case</b>    | Make the keyword search case sensitive.                                                                                                      |

## Reporting NDMP Data

The NetWorker software reports information about NDMP clients, data, and volumes in two ways:

- The NetWorker Management Console (NMC) reporting feature—Reports NDMP data in the same manner as non-NDMP data. The *NetWorker Administration Guide* provides more information.
- The `mminfo` command. Use the `mminfo` program to query the media database for NDMP volume or save set information.

## Querying the NDMP volumes by backup type with the `mminfo` command

You can query save sets by backup format (NDMP or DSA) to display volume information.

For example:

- To query NDMP volumes, type `mminfo -q ndmp`. Output similar to the following appears:

```
volume client date size level name
005D0000 simlcifs1 6/22/2011 1036 MB full /fs1
005D0001 simlcifs1 6/22/2011 173 MB full /fs1
005D0001 simlcifs1 6/22/2011 862 MB full /fs1
005D0002 simlcifs1 6/22/2011 348 MB full /fs1
```

- To query NDMP -DSA volumes, type `mminfo -q dsa`. Output similar to the following appears:

```
volume client date size level name
NDMP.001 10.8.67.219 12/13/2011 644 MB full /vol/vol0
NDMP.001 10.8.67.219 12/13/2011 402 MB full /vol/vol1
```

```
NDMP.001 10.8.67.219 12/13/2011 402 MB full /vol/vol1
NDMP.001 10.8.67.219 12/13/2011 36 MB full /vol/vol2
```

## Querying the NDMP save sets with the mminfo command

To determine which save sets are Network Data Management Protocol (NDMP) save sets and the status of an NDMP save set in the media database, query the media database. NDMP save set status information is important when performing NDMP recoveries:

- To perform a browsable NDMP recover, the `ssflags (f1)` field for an NDMP save set must contain a `(b)`. The `b` value denotes a browsable save set.
- To perform a save set recover from the NetWorker User program, the `ssflags (f1)` field for an NDMP save set must contain either `(r)` or `(b)`.
- An NDMP save set contains an `N` attribute in the `ssflags (f1)` field.
- An NDMP-DSA save set contains an `s` attribute in the `ssflags (f1)` field.

In the following example, the NDMP save set status is recoverable `(r)`. To recover the data, you can only perform a save set recovery from a command line.

```
mminfo -av

volume type client date time size ssid fl lvl name
vol1 dlt clnt 6/22/2011 3:15:12 1036MB 3842140553 hrN full /fs1
```

In the following example, the NDMP-DSA save set status is browsable `(b)`. Recover the data by using the NetWorker User program, or from the command line. A browsable NDMP-DSA save set supports browsable and save set recoveries.

```
mminfo -av

volume type client date time size ssid fl lvl name
vol1 dlt clnt 6/22/2011 3:15:12 36MB 4259813785 cbNs full /fs1
```

## Performing NDMP recoveries

NetWorker uses the `nsrndmp_recover` program to coordinate recover operations between the NetWorker software and the Network Data Management Protocol (NDMP) client. The `nsrndmp_recover` program does not move data to the NDMP client. When the `nsrndmp_recover` program identifies an NDMP-DSA save set, `nsrndmp_recover` automatically runs the `nsrdsa_recover` program on the same host that runs the `nsrndmp_recover` command.

To recover NDMP data, you can run the `nsrndmp_recover` program from a command prompt, or use one of following programs, which automatically starts `nsrndmp_recover`:

- `recover`—The command line program on Windows and UNIX.
- `winworkr`—The NetWorker User GUI on Windows.
- The NMC Recovery wizard.

During the recovery process, the `nsrndmp_recover` program passes `nlist` information to the NDMP client. There are three methods to recover NDMP backups:

- Index-based file-by-file recover—The `nlist` includes file offset and ACL information. When you recover many files, the recover process uses a significant amount of system resources on both the NetWorker server and the NDMP client to build and process the `nlist` information.

- Full save set recovery—The nlist only includes the path to the recovery directory, down to and including the mount point. When you recover many files, the recover process uses less system resources than an index-based NDMP recover to build and process the nlist information.
- NDMP directory restore—A partial save set recovery of a single file or single directory.

For example, when the NetWorker software writes NDMP data to a remote storage node, start the `recover` program on the NetWorker storage node to prevent the data from traversing the network.

**Note:** When you start the `recover` program on the NetWorker server, the data flows from the storage node to the NetWorker server and from the NetWorker server to the NDMP client, over the network.

## NDMP recovery requirements for Celerra and VNX

The following list summarizes the requirements:

### scanner

You cannot use the `scanner` command with the `-i`, `-f` and `-r` options on an NDMP volume. You cannot use the `scanner` command on a volume that contains NDMP and non-NDMP save sets when you load the volume in an NDMP device. The *Scanner command usage* technical note provides more information about using the `scanner` command with NDMP data.

### Cross platform recoveries

You can recover NDMP data to different NDMP client however, you cannot perform a cross platform recover. Recover NDMP data to an NDMP client that is the same brand, a compatible model, and the same operating system as the original NDMP client.

### Devices

Recover Direct-NDMP and Three-party backups performed to an NDMP device from an NDMP device. To improve recover performance from an NDMP tape device, configure the tape device to support variable length records. Recover NDMP-DSA backups from a non-NDMP device.

### Localized environments

When recovering data in a localized NDMP environment, the Index Recover status window shows the process in English and not the localized language.

### NDMP-DSA

For better recovery performance, start the recover process on the NetWorker host where the backup volume resides.

### Immediate recoveries

Run the `nsrndmp_recover` program on the storage node with the locally attached backup device to perform an immediate recovery of NDMP-DSA data.

### Celerra and VNX

During recover operation the filer skips char and block special files. The following error message appears:

```
Warning: /fsl/SPE_REL/my.char_file has an unknown file type,
skipping
```

When you recover named pipes:

- If the recover directory contains 10,000 or more named pipes, then the recover operation will fail. Ensure that the recover directory contains less than 10,000 named pipes.

- The recover process changes the file permissions. The NetWorker software recovers named pipes as normal files.

The Configuring NDMP for VNX documentation on the Support website describes how to recover a tape silvering backup to a different data mover.

### vbb

When you set the backup type for an Celerra or VNX filer to vbb, the NetWorker software performs a block-based backup.

The NetWorker software recovers the data to a raw device. Use the Relocate data option to specify the raw device. When you use deduplication on the source or target file system, you cannot perform an index based file-by-file recover. When you perform a destructive save set recover, the recover process performs the following actions:

- Recovers the data to the original location or an alternate location.
- Overwrites existing data.
- Overlays the data at the file system level and reimposes the saved image on the file system.

Configuring NDMP backups on Celerra and Using Celerra Data Deduplication documentation on the Support website provides detailed information about how to prepare the filer before you perform FDR.

## DAR and DDAR

By default, the Network Data Management Protocol (NDMP) recover process reads an entire tape from start to finish. The recover process extracts the data as it encounters the data on the tape. For large backup images, recovery is slow.

The Direct Access Recovery (DAR) and Directory DAR (DDAR) recovery process:

- Provides the ability to recover a file or directory from the exact location on a tape.
- DDAR only passes the directory path to the NAS filer. DAR passes the paths of each file individually.
- Reduces the size of the nlist information that the recover process stores in memory. During the recover process, the NAS filer (DDAR) assumes that the directory path includes all cataloged files and directories. However, DAR mentions each file that it wants recovered.
- Does not sequentially read the file or record numbers on the tape to locate the data, which reduces the amount of time that you require to recover specific files from a backup.

 **Note:** [Creating and configuring the NDMP client resource](#) on page 67 describes how to configure the DAR and DDAR Application Information attributes for NDMP clients.

Use the `recover` command, the NetWorker User program, or Browse Recovery using the NMC Recovery Wizard to perform DAR and DDAR recoveries. You cannot use the `nsrndmp_recover` program to perform DAR/DDAR recoveries.

### When not to use DAR or DDAR

DAR and DDAR recoveries send multiple pathnames across the network to the NDMP Data Server and, in three-party configurations, to the NetWorker server. The recover process stores the pathnames in memory on the NDMP Data Server. Recoveries of a large amount of data from a large save set can negatively affect the network and the NDMP Data Server resources.

Do not use DAR and DDAR to recover the following objects:

- Several thousands of files in a single index-based recover operation.



- A specific directory structure containing several thousand or millions of files.

To perform a non-DAR-based recovery of a save set when you set the *DIRECT=y* at the time of backup, first define the *NSR\_NDMP\_RECOVER\_NO\_DAR=y* variable in the Application Information attribute of the NDMP client.

## Recover window

The **Recover** window displays information about recover configurations that are created with the NetWorker Management Console (NMC) Recovery wizard.

You can use this window to:

- Start the NMC Recovery wizard to create recover configurations or modify saved recover configurations.
- Identify the status of a recover configuration that is created with the NMC Recovery wizard.
- Start and stop a recover job.

The **Recover** window is divided into five sections:

- **Toolbar**—The toolbar is hidden by default. To display the recovery toolbar, select **View > Show toolbar**
- **Summary**
- **Configured Recovers**
- **Currently Running**

A splitter separates the **Configured Recovers** section from **Currently running** window. You can click and move the splitter to resize these two windows.

## Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. The following table summarizes the function of each toolbar button.

**Table 24** Recovery toolbar options

| Button                                                                              | Function                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Starts the NMC <b>Recover</b> wizard to create recover configurations.                                                                                                                        |
|  | Displays the <b>Properties</b> window for the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                          |
|  | Deletes the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                                                            |
|  | Displays online help for the <b>Recover</b> window.                                                                                                                                           |
|  | Displays the <b>Find</b> window at the bottom of the <b>Recover</b> window. The <b>Find</b> window allows you to perform keyword searches for messages that appear in the <b>Logs</b> window. |
|  | Start the recover operation for a selected saved recover configuration. This option is only available for a recover configuration that has a Never run, or Failed status.                     |

**Table 24** Recovery toolbar options (continued)

| Button                                                                            | Function                                                                                     |
|-----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
|  | Stop in-progress recover operation that you selected in the <b>Currently Running</b> window. |

 **Note:** The **Recover** toolbar does not appear by default. To display the **Recover** toolbar, select **View > Show toolbar**.

## Recover Summary

The Recover Summary section displays a high-level overview of recover jobs.

This section includes the following information:

- Total Recovers—The total number of successful recover jobs.
- Since—The number of successful recover jobs since this date.

## Configured Recovers

The **Configured Recovers** window displays a list of saved recover configurations in a tabular format. You can sort the information by column. The **Configured Recovers** table displays the following information for each saved recover configuration:

- Status—The job status of a saved recover configuration.
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB.
- Comment
- OS—The operating system of the source host.
- Recover requestor—The Windows or UNIX account used to create the recover configuration.
- Start Time
- End Time
- Start date

**Table 25** Save recover configuration job status

| Icon                                                                                | Description                                        |
|-------------------------------------------------------------------------------------|----------------------------------------------------|
|  | The last recover attempt failed.                   |
|  | The last recover attempt completed successfully.   |
|  | The recover job has never run.                     |
|  | The recover job is scheduled to run in the future. |
|  | The recover job has expired.                       |

## Currently running

The **Currently Running** window displays a list of in progress recover jobs in a tabular format. You can sort the information by column. The **Currently Running** table displays the following information for each job:

- Status
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB
- Volume
- Comment
- Device
- Size
- Total size
- % complete
- Rate (KB/s)
- Start time
- Duration
- Currently running

## Find

The **Find** section appears along the bottom of the **Recover** window, after you select the **Find** button on the **Recover** toolbar. **Find** allows you to search for keywords in the **Configured Recovers** window. The following table summarizes the available find options.

**Table 26** Find options

| Find option          | Description                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Find</b>          | Highlight the first saved recover configuration that contains the specified keyword.                                                         |
| <b>Prev</b>          | Highlight the previous saved recover configuration that contains the specified keyword.                                                      |
| <b>Highlight All</b> | Highlights each saved recover configuration that contains the specified keyword.                                                             |
| <b>Sort Selected</b> | Sorts each highlighted recover configuration in the Configured Recover table so that they appear at the top of the Configured Recover table. |
| <b>Match case</b>    | Make the keyword search case sensitive.                                                                                                      |

## Performing an NDMP index-based file-by-file data recovery

Perform an NDMP index based file-by-file recover in the same manner as a non-NDMP data recover. You can restore the data to the original NDMP client or perform a directed recovery to a different NDMP client.

Before you perform an index-based file-by-file recover, review the following information:

- Set the *HIST=y* in the Application Information attribute of the NDMP client at the time of the backup.
- The NDMP save set must be browsable. You cannot perform a browsable recover of a recoverable or recyclable save set. [Reporting NDMP Data](#) on page 93 describes how to determine the status of an NDMP save set.
- Do not use an index-based recovery to recover a large numbers of files or directories. For better recovery performance, use a save set recover. [Performing a Full or Directory Restore of NDMP data by using a save set recovery](#) on page 103 provides more information.
- To perform an index-based file-by-file recover:
  - Use the NetWorker User program on a Windows host. [Performing an NDMP index-based file-by-file recover using the NetWorker User program](#) on page 100 provides detailed information.
  - Use the `recover` program. [Performing an NDMP index-based file-by-file recover from a command prompt](#) on page 102 provides detailed information.

### Performing an NDMP index-based file-by-file recover using the NetWorker User program

On Windows, to recover data to the original NDMP client or to a different NDMP client, perform the following steps.

#### Procedure

1. Open the NetWorker User program and connect to the NetWorker server.

 **NOTICE** If you receive the error:

```
No file indexes were found for client client_name on server
server_name
```

Try connecting to a different NetWorker server and you selected the correct NetWorker server, then ensure that you selected a browsable save set. Alternatively, perform a save set recover.

2. Select **Recover** to open the **Source Client** window.
3. Select source NDMP client and click **OK**. The local client is the default selection.
4. Select the destination client for the recovered data and click **OK**. If the destination client is not the source client, ensure the NAS filer is the same brand, a compatible model and the same operating system as the source NDMP client.
5. (Optional) Recover the data from an earlier backup time. The **Recover** window appears with the latest version of the backup files. To recover data from an earlier backup, change the date and time of backup using one of the following methods:
  - a. Change the browse time for all files in the recover window:
    - From the **View** menu, select **Change Browse Time**.
    - In the **Change Browse Time** window, select a new day within the calendar. Select **Previous Month** or **Next Month** to change from the current month.

- In the **Time** field, change the time of day by typing an hour, a minute, and the letter a for A.M. or p for P.M. Use the 12-hour format.
  - Click **OK**.
- b. View all versions of the selected file system object:
- Highlight the file or directory for review.
  - From the **View** menu select **Versions**.
  - Once you locate the version to recover, change the browse time. To change the browse time, highlight the volume, directory, or file and click **Change Browse Time**. The **Version** window closes and the **Recover** window reflects the new browse time.
6. (Optional) Search for the files. To search for and recover the most recently backed-up version of a file or directory:
- a. From the **File** menu, select **Find**.
  - b. Type the name of the file or directory. Use wildcards to expand the search; without wildcards, partial filenames do not provide any results.
7. Mark the data to recover. To select file system objects to recover:
- a. In the left pane of the **Recover** window, click the appropriate directory folder.
  - b. Mark each directory or file to recover by selecting the checkbox next to each directory or file.
8. (Optional) Relocate the data to a different location. By default, the recover process recovers the selected files to the original location.

 **Note:** The NDMP protocol does not support name conflict resolutions. NetWorker will always overwrite existing files that have the same name as the recovered file. It is recommended that you recover the NDMP data to a different location, to avoid data loss.

To relocate the files to a different location:

- a. Select **Recover Options** from the **Options** menu.  
NDMP recoveries do not support the following options:
    - Rename recovered file
    - Discard recovered file
    - Prompt for every file conflict

NDMP recoveries will always overwrite existing files. It is recommended that you relocate the NDMP data to a different location, to avoid data loss.
  - b. In the **Relocate Recovered Data To** field, type the full path name of the target directory, click **OK**.  
  
The target directory is a literal string and must match the path as seen by the NAS filer in its native OS, exactly. Otherwise, the recover process uses the original location and overwrites existing files with the same name.
9. (Optional) To view the volumes required to recover the marked file system objects, from the **View** menu, select **Required Volumes**.
10. Click **Start** to begin the recovery. If any required volume is not available to the NetWorker server, a volume status warning appears.

When this warning appears:

- a. Click **No**.
- b. From the **View** menu, select **Required Volumes**.
- c. Ensure that the NetWorker software can mount each listed volume into an available device.
- d. Attempt the recover operation again.

The NetWorker server takes a few moments to recover the files, depending on file size, network traffic, server load, and tape positioning. During this time, messages appear so that you can monitor the progress of the recovery.

When the recovery completes successfully, a message similar to the following appears:

```
Received 1 file(S) from NSR server
server Recover completion time: Tue Jan 21 08:33:04 2009
```

## Performing an NDMP index-based file-by-file recover from a command prompt

This section applies to command line recoveries from a Windows and UNIX client.

### About this task

To avoid using the Windows version of `recover.exe` on Windows operating systems, perform one of the following actions:

- Specify the full path to the recover program. For example: `C:\Program Files\EMC NetWorker\nsr\bin\recover.exe`
- Ensure that the `$PATH` environment variable contains the `NetWorker_install_path\bin` directory before `%SystemRoot%\System32`

To recover Network Data Management Protocol (NDMP) data from a command prompt on a UNIX or Windows NetWorker host, perform the following steps.

### Procedure

1. From the command prompt, type:

```
recover -s NetWorker_servername -c client_name
```

where:

- `-s NetWorker_servername` specifies a particular NetWorker server on the network to use when recovering data.

When you do not use the `-s` option, the `recover` program tries to connect to the first computer listed in the servers file. When the servers file does not contain any servers, or lists more than one server, the **Change Server** window appears, and you can select the server.

- `-c client_name` specifies the source NDMP client.

2. When prompted, type the directory to browse, for example:

```
cd /mydirectory
```

3. Use the `add` command to add the required files or folders to the recover list. The *NetWorker Command Reference Guide* provides a complete list of options for the `recover` command.
4. When restoring NDMP data, it is recommended that you relocate the NDMP data to a different location.

**Note:** The NDMP protocol does not support name conflict resolutions. NetWorker will always overwrite existing files that have the same name as the recovered file. It is recommended that you recover the NDMP data to a different location, to avoid data loss.

- To relocate the data to a different directory, type:

```
relocate destination_directory_name
```

The target pathname for *destination\_directory\_name* is a literal string and must match the path as seen by the NAS filer in its native OS, exactly. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

- To recover the data to a different host, type:

```
relocate target_hostname::/mount_point
```

Data ONTAP may require you to add a backslash (\) after the mount point. For example, *target\_hostname::\mount\_point\*.

- After you add all of the required files, type:

```
recover
```

## Performing a Full or Directory Restore of NDMP data by using a save set recovery

You perform a Network Data Management Protocol (NDMP) save set recover in the same manner as a non-NDMP save set recovery. You can recover data to the original NDMP client or perform a directed recovery of the data to a different NDMP client of the same platform.

Before you perform a full save set recover, review the following information:

- Use a full save set recovery to recover all files and folders in an NDMP data save set, or to recover an entire directory within an NDMP save set. You cannot use the NetWorker User program to perform an NDMP Directory Restore.
- To use the NetWorker User program on Windows, a client file index entry for the save set must exist. When the index entry for the save set does not exist, the recover fails with an `index not found` error. When the client file index entries do not exist for the save set, use the `nsrndmp_recover` program with the `-v off` option.
- You cannot perform a save set recover from the NetWorker User program when the save set status is eligible for recycling (E). The recover process requires a recoverable (r) or browsable (b) save set status. The *NetWorker Administration Guide* provides information on how to change the status of a save set. A save set recover reads the entire tape, from beginning to end, to find and recover the requested files. The recovery process completes when the recover operations reads all required tapes in their entirety.
- As each file recovers, the file name appears on the target share but the file size is 0 KB. The actual file size update occurs after the recovery completes.

## Performing an NDMP save set recover by using the NetWorker User in Windows

### About this task

**NOTICE** When the recover operations fails with the error:

```
Failed to propagate handle <number> to child process: Access is denied
```

The save set is not in the client file index of the NDMP client. Perform a save set recover from a command prompt. [Performing an NDMP save set recovery from the command prompt](#) on page 105 provides more information.

## Procedure

1. Start the NetWorker User program.
2. On the **Change Server** window, select the NetWorker server and click **OK**.
3. Select **Options > Recover Save Sets**.
4. On the **Source Client** window, select the appropriate NDMP client, and then click **OK**.
5. On the **Save Sets** window, select the name of the save set.
6. Select the version of the save set, if there are multiple versions. You can also select the cloned version of a save set, if applicable.
7. To recover specific files and directories instead of the entire save set:
  - a. Click **Files**.
  - b. Specify the files and directories, one per line.
  - c. Click **OK**.

 **NOTICE** Do not use this method to mark tens of thousands of files. Instead, perform an NDMP Directory Restore. Marking many files and directories generates a large nlist and requires intensive resources on both the NetWorker server and the NAS filer.

8. Click **Recover Options**.

An NDMP data recovery does not support the following options:

- Rename recovered file
- Discard recovered file
- Prompt for every file conflict

 **NOTICE** It is recommended that you relocate the NDMP data to a different location. NDMP recoveries always overwrite existing files.

9. To recover the data to a pathname that is different from the original backup location, in the **Relocate Recovered Data To** field, type the full pathname of the destination directory, then click **Ok**.

For NDMP data recoveries, the target pathname is a literal string and must exactly match the path as seen by the native OS on the NAS filer. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

10. To recover the data to a different NDMP client, specify the name of the client to receive the NDMP data in the **Destination Client** field.
11. To view the volumes that are required to perform the recover, select **View > Required Volumes**
12. Click **OK** to begin the recovery. The recovery status appears in the **Recover Status** window.



## Performing an NDMP save set recovery from the command prompt

To perform a save set recovery to the original Network Data Management Protocol (NDMP) client or to a different NDMP client, use the `nsrndmp_recover` command.

### About this task

For example:

```
nsrndmp_recover -s NetWorker_server -c source_ndmp_client -S ssid/cloneid -v off
-m target_ndmp_client::/target_path /source_path
```

where:

- `source_ndmp_client` is the hostname of the source NDMP client.
- `target_ndmp_client` is the hostname of the destination NDMP client.
- `/source_path` is the original location of the data.
- `/target_path` is the location to recover the data.

**NOTICE** It is recommended that you relocate the NDMP data to a different location. NDMP recoveries always overwrite existing files. The `/target_path` is a literal string and must exactly match the path as seen by the native OS on the NAS filer. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

- `-v off` allows you to restore data when client file index of the NDMP client does not contain information about the NDMP save set.

In the following examples, the NetWorker server is mars and the backup client is venus:

- To recover a mount point `/mnt` from a backup of NDMP host venus to a directory `/newmnt` on NDMP host jupiter, type:

```
nsrndmp_recover -s mars -c venus -S 123456789 -v off -m jupiter::/newmnt
```

- To recover a mount point `/mnt` from a backup of NDMP host venus to NDMP host pluto, type:

```
nsrndmp_recover -s mars -c venus -R pluto -S 123456789 -v off -m /mnt
```

Data ONTAP may require that you to add a slash (/) after the mount point. For example, `target_hostname::/ mount_point/`.

## Performing destructive save set recoveries for vbb backups

Use the `nsrndmp_recover` command with the `-r raw_device` and `-m mount_point` options to perform a destructive save set recovery.

**NOTICE** Do not perform an NDMP Directory Restore from a vbb backup of a Celerra deduplicated file system.

For example:

On a Windows system to perform a destructive save set recovery to the `/data` drive:

```
nsrndmp_recover -s mars -c venus -m /data -r raw_device_name -S 2674606849
```

On a UNIX system, the following command performs a destructive save set recovery to the `/dev/c1t1d0s0` device, mounted at the `/` file system:

```
nsrndmp_recover -s mars -c venus -r /dev/c1t1d0s0 -S 2674606849 -m /
```

The *NetWorker Command Reference Guide* or the UNIX man page provides more information about the `nsrndmp_recover` command.

If you do not specify the `-r` option when you use the `-m`, the recover operation:

- Is nondestructive.
- Operates at the file or directory level, rather than the file system level.

This nondestructive restore overwrites existing files on the destination that have the same names as those in the recovery list. Other data remains untouched on the file system.

Use this nondestructive method to:

- Perform a directory level recovery on a high density file system.
- Recover many files in one directory.

## Troubleshooting NDMP recover

This section provides a list of the possible causes and the possible resolutions for NDMP recovery issues.

### RESTORE: could not create path *pathname*

This error message appears when restoring NetApp data. This error, when encountered, appears in the `daemon.raw` file of the NetWorker server and the recovery output.

To resolve this issue:

- Ensure that you specify a source and a target path during the recover that exists on the target filer.
- If you set the `UTF8=Y` application information variable during an NDMP client backup and the backup contains path names with non-ASCII characters, then perform a save set recover. Index-based recoveries will fail with this error message.

### These files were not restored (Restore failed with error, or file/directory specified but not found in backup)

This error message appears in the `daemon.raw` file of the NetWorker server and the in the recovery output.

To resolve this issue:

- Ensure that the file or directory specified during the recover, exists in the save set.
- Ensure that the pathname specified to relocate the data exists on the destination filer. For NDMP data recoveries, the target pathname is a literal string and must exactly match the path as seen by the native OS on the NAS filer.

# CHAPTER 3

## Isilon

This chapter includes the following topics:

|                                                                                             |     |
|---------------------------------------------------------------------------------------------|-----|
| • <a href="#">Choosing a device type</a> .....                                              | 108 |
| • <a href="#">Configuring devices for NDMP operations</a> .....                             | 108 |
| • <a href="#">Configure NetWorker for NDMP backup and clone operations</a> .....            | 119 |
| • <a href="#">Monitoring NetWorker Server activities in the Administration window</a> ..... | 156 |
| • <a href="#">Reporting NDMP Data</a> .....                                                 | 173 |
| • <a href="#">Performing NDMP recoveries</a> .....                                          | 174 |

## Choosing a device type

Network Data Management Protocol (NDMP) backups can be written to either an NDMP device, or if using NDMP-DSA, to a non-NDMP device.

Perform either of the following tasks:

- Configure devices for NDMP operations.
- Configure non-NDMP devices. If you are using NDMP-DSA, refer to the *NetWorker Administration Guide* for device configuration.

For a description of each configuration, refer to [Configurations in a NetWorker NDMP environment](#) on page 16.



**Note:**

PowerStore does not support local or direct NDMP backup.

## Configuring devices for NDMP operations

Review this section for information about how to configure the NetWorker environment for Network Data Management Protocol (NDMP) data operations.

The *NetWorker Hardware Compatibility Guide* on the Support website provides a list of NDMP devices that the NetWorker software supports.

### NDMP device limitations

Review these limitations before you configure Network Data Management Protocol (NDMP) devices:

- The timeout of the NetWorker server `nsrmmmd` resource attribute does not apply to NDMP devices, but it does apply to storage nodes devices.
- You cannot use the `jbexercise` utility with an NDMP autochanger.
- You cannot configure NDMP devices on a dedicated storage node.
- You must use a non-rewind device handle for the NDMP media device handle.
- You cannot configure advanced file type devices and file type devices as NDMP devices.
- You cannot configure an NDMP autochanger when the NDMP protocol is earlier than version 3. You must determine the NDMP device handles, then use the `jbconfig` command to configure the autochanger.

### Determining NDMP device pathnames

To configure an NDMP stand-alone device or an NDMP jukebox, you must first determine the path names of the media devices. If the NAS filer does not support the NDMP\_CONFIG interface or uses NDMP version 3, you must also determine the library device handle.

To determine the NDMP device path names and the library handle, use the `inquire` command or vendor-specific commands.

### Determining the NDMP device path names using the `inquire` command

#### About this task

Use the `inquire` command to determine the path names and library handle.

## Procedure

1. From a command prompt on the NetWorker server, type:

```
inquire -N NAS_hostname -T
```

2. When prompted, specify the NAS username and password.

**NOTICE** Use the `inquire` command with caution. When you run `inquire`, the command sends the `SCSI inquiry` command to all devices that are detected on the SCSI bus. If you use the `inquire` command during normal operations, unforeseen errors can occur, which might result in data loss.

## Determining the NDMP device path names for Isilon

Before you configure an NDMP autochanger you must determine the device path names of NDMP devices and the robotic arm. The following table provides vendor-specific information that you can use to determine the device path names.

### About this task

For an NDMP local backup, configure Backup Accelerator:

## Procedure

- Use the `isi fc list` command to ensure that the state of each fibre channel ports is enabled.
- Use the `isi tape rescan --reconcile` command to scan for tape devices.  
`--reconcile` deletes the device entries for devices that a Backup Accelerator node no longer manages.
- Use the `isi tape ls -v` to display a list of current devices.

## Dynamic drive sharing

Dynamic Drive Sharing (DDS) is a feature that provides NetWorker software with the ability to recognize shared physical tape drives. DDS enables NetWorker software to perform the following operations:

- Skip the shared tape drives that are in use.
- Route the backups or recoveries to other available shared tape drives.

## Introduction to DDS

DDS controls application requests for tape media and allows the NetWorker server and all storage nodes to access and share all attached devices.

A system administrator can configure DDS by setting a sharing policy for devices that are accessible from multiple storage nodes.

There are two terms that are central to the use of DDS are drive and device. Within the context of DDS, these terms are defined as follows:

- Drive—The physical backup object, such as a tape drive, disk, or file.
  - Device—The access path to the physical drive.
- Note:** NetWorker only supports DDS in a storage area network (SAN) Fibre Channel environment and not in a direct-connect SCSI environment.

## Benefits of DDS


Enabling DDS on a NetWorker system provides these benefits:

- Reduces storage costs—You can share a single tape drive among several storage nodes. In fact, since NetWorker software uses the same open tape format for UNIX, Windows, NetWare and Linux, you can share the same tape between different platforms (assuming that respective save sets belong to the same pool).
- Reduces LAN traffic—You can configure clients as SAN storage nodes that can send save sets over the SAN to shared drives.
- Provides fault tolerance—Within a SAN environment, you can configure hardware to eliminate a single point of failure.
- Provides configuration over a greater distance—You can configure a system over a greater distance than with SCSI connections.

## DDS configuration overview

The following figure illustrates the DDS process and potential device sharing configurations. This basic configuration consists of a server, two storage nodes, and a library with two tape drives.

**Figure 11** Dynamic Drive Sharing



In this figure:

- Storage nodes sn\_1 and sn\_2 are attached to the library.
- Each storage node, on its own, has access to drive\_1 and drive\_2.
- With DDS enabled, both storage nodes have access to both drives and can recognize when a shared drive is in use.

This configuration requires two DDS licenses, one for each drive.

**Note:** Ensure that all applicable devices can be seen from each storage node by running the `inquire -l` command locally on each storage node.

## DDS block-size compatibility between UNIX and Windows

With DDS enabled, drives can be shared between storage nodes on different platforms, such as UNIX and Microsoft Windows. For NetWorker software operations (such as backups and recoveries) to take place successfully, ensure that the block size is compatible between different platforms or hardware.

To ensure compatibility, make sure one of the following conditions is met:

- The various storage nodes sharing a drive support the same block sizes.
- When a tape is labeled on a drive, it is labeled with the block size defined on the storage nodes.

## Block-size incompatibility between UNIX and Windows

Incompatible block-size settings between UNIX and Microsoft Windows storage nodes could result in any of these error scenarios:

- A backup taken on a UNIX node might not be recoverable on a Microsoft Windows node if the Windows node does not support large block sizes.
- A UNIX process labels and saves data to a tape and leaves the tape mounted. A Microsoft Windows process subsequently attempts to verify the label on this tape and fails because the label verification is done by reading a header from the data portion.
- A tape on a UNIX node is labeled with a large block size. The backup is started on a Microsoft Windows node and the Windows node attempts to write the backup by using the default block size. Internally, the backup on Windows is written by breaking down the big buffer of data into smaller segments of writable block sizes.

Attempting to recover a specific file on Windows in this situation fails due to positioning errors on the tape. The data is still recoverable from the Windows side, since the NetWorker software will switch from using file and block positioning to reading the tape from the beginning to reach the correct position. The data might not, however, be recoverable from the UNIX side.

## Unintended Access to DDS device prevention

The Reserve/Release attribute has been added to the Device resource for tape devices to support Reserve/Release, including the Persistent Reserve commands.

Reserve/Release is a mechanism that uses SCSI commands to attempt to prevent unintended access to tape drives that are connected by using a shared-access technology, such as Fibre Channel, iSCSI, or SCSI multiplexers. It is a “cooperative” and host-based mechanism, which means that all applications should respect the reservations and not purposely break them. Access is granted based on the host system that reserved the device. Other applications that run on that host cannot be prevented from accessing a reserved device.

Reserve/Release cannot prevent a malicious or badly behaved application from accessing a reserved device. It also cannot prevent all problems caused by hardware issues (such as SCSI resets or FC LIPs) from interrupting data access.

The basic sequence requires that a host reserve a tape drive (using specific SCSI commands) before attempting to access the tape drive. If this “reservation” succeeds, then the host can use the drive. If the reservation fails (usually because the device is reserved by someone else), then the host attempting the reservation should not attempt to use the drive. When a host has finished using a reserved drive, that host must release the drive by using the appropriate SCSI commands.

The reservation is maintained by the drive itself. With older (called “Simple” in NetWorker software) Reserve/Release, the reservation is based on the SCSI ID of the system that issued the reserve command. For tape drives connected to Fibre Channel (FC) using FC-SCSI bridges, the mapping between FC host and reservation is done inside the bridge, since the initiator on the SCSI side is always the bridge itself, regardless which host actually issued the reserve command.

For Persistent Reserve, the reservation is associated with a 64-bit “key” that is registered by the host. Several keys can be registered with a given drive at any given time, but only one may hold

the active reservation. NetWorker software uses the “exclusive” reservation method for Persistent Reserve. Only the host that holds the active reservation is allowed to access the drive. The Reserve/Release attribute does not support file type or advanced file type devices.

The settings that relate to Reserve/Release and Persistent Reserve are found in a device’s **Properties** window, on the **Advanced** tab. They are visible only when diagnostic mode is turned on.

The default setting for Reserve/Release is None. Once any other Reserve/Release setting is selected, it works automatically, without further user intervention. The Reserve/Release attribute is supported only on Common Device Interface (CDI) platforms, so if the CDI attribute in a device’s **Properties** is set to Not Used, then Reserve/Release settings are ignored.

For newer hardware, once a Reserve/Release setting (other than None) has been selected, the appropriate Persistent Reserve commands are automatically issued before a device is opened for reading or writing, and before the device is closed. With older hardware, a SCSI-2 Reserve command is issued before opening the device, and a SCSI-2 Release command is issued after the device is closed.

Reserve/Release has these possible settings:

- None (the default)
- Simple
- Persistent Reserve
- Persistent Reserve + APTPL (Activate Persist Through Power Loss)

The Persistent Reserve Key attribute has also been added. It is used with Persistent Reservation calls.

### Restrictions for use of the SCSI Reserve/Release setting

There are restrictions for using the SCSI Reserve or Release setting.

Consider the following:

- It is available on CDI platforms only. Consequently, since CDI is not supported within an NDMP environment, Reserve/Release is not supported with NDMP.
- Not all drives support persistent Reserve/Release. (All drives support at least simple reserve release. The code automatically drops back from Persistent +APTPL or Persistent to Simple on drives that do not support Persistent.)
- SCSI resets can clear Simple reservations at the device.
- Even with Reserve/Release, there is no guarantee against data loss.
- If the operating system has its own Reserve/Release feature, that feature must be disabled in order for the NetWorker Reserve/Release feature to work.
- Even if all of the enterprise’s NetWorker storage nodes have this feature enabled, then it is possible that, on the storage node where a backup operation is run, data loss can be caused by the operating system’s utilities or by third-party programs.

### DDS on NDMP nodes in a SAN environment

You can configure shared drives between NDMP nodes in a SAN environment.

Ensure that:

- All the components of a SAN configuration are compatible when DDS is enabled with the NetWorker NDMP feature.
- The Fibre Channel switches are compatible with any NDMP hosts within a SAN.
- NDMP hosts and libraries in the SAN are compatible with each other.



- The NDMP nodes that will share the drives are homogeneous.

**Note:** The current NDMP implementation does not allow the sharing of drives between non-homogeneous NDMP nodes. There is, however, no inherent limitation within DDS that would prevent this.

## DDS attributes in the device properties

Configure the attributes that DDS uses, in the **Properties** window for a device.

The attributes include:

- Hardware ID
- Shared Devices

### Hardware ID attribute

The Hardware ID attribute tracks the drives that are shared between multiple hosts. Device instances that share the same physical drive across multiple hosts have the same hardware ID. The device autoconfiguration process automatically assigns the Hardware ID to a device, or it is added when manually configuring a device. Users cannot edit the Hardware ID.

You can view the Hardware ID in the **Properties** window for a device, on the **General** tab, in the **Device Sharing** area.

NetWorker generates the Hardware ID when a device is scanned or configured. The Hardware ID consists of the following components:

- Hardware serial number
- Device type
- Worldwide part number (WWPN)
- Worldwide name (WWN)

### Shared Devices attribute

The Shared Devices attribute appears on the **Operations** tab of a device's **Properties** window when in diagnostic mode. It features values that can be used to manipulate all shared instances of a drive simultaneously. This attribute enables or disables all devices that share the same Hardware ID with a single action. The following table lists allowed values and descriptions for the attribute.

**Table 27** Shared Devices attributes

| Value       | Description                                                                                                                                              |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable All  | When selected, enables all devices with the same Hardware ID.                                                                                            |
| Disable All | When selected, disables all the devices with the same Hardware ID.                                                                                       |
| Done        | This value is the default setting. After the server has enabled or disabled all devices with the same Hardware ID, the attribute value is reset to Done. |

You cannot configure the Shared Devices attribute with the `jbconfig` program.

## Idle Device Timeout attribute and DDS

A tape might remain mounted in a drive after a backup completes. Other requests for the drive from another device path must wait during this timeout period. Use the Idle Device Timeout attribute to adjust the timeout value.

The Idle Device Timeout attribute is not specifically a DDS attribute, but is useful in configuring shared drives. This attribute appears on the device **Properties** window on the **Advanced** tab when displayed in Diagnostic Mode. The default value is 0 (zero) minutes, which means that the device never times out and you must manually eject the tape.

If the device belongs to a library, you can also specify the Idle Device Timeout value for all devices in the library. However, the library value will take effect only on those devices whose **Idle Device Timeout** value is 0. The Idle Device Timeout value for a library is located on the **Timer** tab of the library **Properties** window.

## Max active devices

In a DDS environment, use the Max active devices attribute, on the **General** tab of the Storage Node resource to define the maximum number of active devices for a storage node.

This attribute sets the maximum number of devices that NetWorker may use from the storage node in a DDS configuration. In large environments with media libraries that have a large number of devices, storage nodes might not have the ability to optimize all the drives in the library. The Max active devices attribute allows you to limit the number of devices that the storage node uses at a specified time, which allows the storage node to have access to all the devices in the library, but does not limit the storage node to the number of devices it can fully optimize.

## Configuring NDMP on Isilon filer

### About this task

Before you perform a Network Data management Protocol (NDMP) backup for an Isilon client, you must configure OneFS NDMP.

#### 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
```

5. To access the global view, type the following command:

```
isi ndmp settings global view
```

The system displays the following NDMP settings:

```
Service: True
Port: 10000
Dma: generic
Bre Max Num Contexts: 64
Msb Context Retention Duration: 300
Msr Context Retention Duration: 600
```

## Configuring NDMP devices

You can back up NDMP data to an NDMP or non-NDMP device in a standalone or library configuration. You can also back up NDMP data to ACSLS controlled silos.

### Configuration for an Isilon attached tape library

For Isilon attached with a tape library, tape drivers are only attached to the backup Accelerator node. The tape drivers are not accessible from other data nodes. To enable backups to the tape drivers, you must configure a NAS client with the Accelerator node on the NetWorker server. Meaning you must configure Device and storage pool with the Accelerator client.

### Configuring a standalone NDMP device

Use the NetWorker Management Console (NMC) to configure a standalone Network Data Management Protocol (NDMP) tape device for Direct NDMP backups.

#### Procedure

1. In the **Administration** window, click **Devices**.
2. In the navigation tree, right-click **Devices**, and then select **New**.
3. In the **Name** attribute, specify the NDMP device in the format:

```
rd=NAS_hostname:NAS_device_handle (NDMP)
```

where:

- *NAS\_hostname* is the hostname of the NAS that has the NDMP device attached.
- *NAS\_device\_handle* is the path of the device.

 **Note:** Configure the NDMP device as a remote device and add (NDMP) after the pathname. Otherwise, you receive a message similar to the following:

```
NDMP device name shall be in rd=snode:devname (NDMP) format
```

4. In the **Media Type** attribute, specify the device type.
5. Specify a valid NAS administrator account in the **Remote User** attribute.
6. Specify the password for the NAS administrator account in the **Password** attribute.
7. On the **Configuration** tab:
  - a. Select the **NDMP** checkbox. You can only set this attribute when you create the device. You cannot change the NDMP attribute after you create the device. To change the device configuration, delete and re-create the device.
  - b. Set the **Target Sessions** attribute to 1. NDMP devices do not support multiplexing.
  - c. The **Dedicated Storage Node** attribute must remain at the default value: **No**.

8. Under the **Advanced** tab, the **CDI** attribute must remain at the default value: `Not used`.
9. (Optional) Change the block size that is used by the NDMP device.  
By default, NDMP devices use a block size of 60 KB. If required, select a different block size in the **Device block size** field. When you configure the NDMP client, set the `NDMP_AUTO_BLOCK_SIZE` environment variable in the **Application Information** attribute.
10. Click **OK**.

## Configuring an NDMP autochanger

You can use an NDMP autochanger to manage Direct NDMP or Three-party backups with NDMP devices. To configure an NDMP autochanger, use NMC or the `jbconfig` command.

### Configuring an NDMP autochanger with NMC

When you configure an NDMP autochanger in NMC, the NetWorker software first detects the NDMP devices and then configures the library.

#### Procedure

1. In the **NetWorker Administration** window, click **Devices**.
2. Right-click the NetWorker Server, and then select **Configure All Libraries**.
3. On the **Provide General Configuration Information** window, accept the default library type, **SCSI/NDMP**, and then click **Next**.
4. On the **Select Target Storage Nodes** window, click **Create a new Storage Node**.
5. On the **Storage Node Name** field, specify the hostname of the NAS.
6. In the **Device Scan Type** attribute, select **NDMP**.
7. In the **NDMP User Name** and **NDMP Password** fields, specify the NAS administrator account. If DinoStor TapeServer manages the autochanger, specify the DinoStor username and password.
8. Click **Start Configuration**.
9. Click **Finish**.
10. Monitor the **Log** window for the status of the device scan.

When you specify an incorrect username and password combination:

- The Log status window reports:

```
No configured libraries detected on storage node
storage_node_name
```

- The `daemon.raw` file on the NetWorker server reports:

```
NDMP Service Debug: The process id for NDMP service is 0xb6c0b7b0
42597:dvdetect: connect auth: connection has not been authorized
42610:dvdetect: The NDMP connection is not successfully
authorized on host 'storage_node_name'
```

To resolve this issue, relaunch the **Configure All Libraries** wizard and correct the NDMP username and password combination.

**Note:** If the **Log** window reports that NetWorker cannot detect the serial numbers for the library, see [Configuring an NDMP autochanger by using the `jbconfig` command](#) on page 39 for detailed instructions.

## Configuring an NDMP autochanger by using the `jbconfig` command

It is recommended that you use the NMC interface to configure an NDMP autochanger. Use the `jbconfig` command when you cannot configure the autochanger by using the NMC Configure Library wizard.

### About this task

The *NetWorker Command Reference Guide* or the UNIX man page provides more information about the `jbconfig` command.

### Procedure

1. Log in to the NetWorker server as root on UNIX, or Administrator on Windows.
2. At the command prompt, type `jbconfig`
3. At the **What kind of jukebox are you configuring** prompt, type 2 to configure an autodetected NDMP SCSI jukebox.
4. When prompted for an NDMP username, specify the NAS administrator account.
5. When prompted for an NDMP password, specify the NAS administrator password.
6. When prompted for the NDMP Tape Server Name, specify the NAS filer hostname.
7. At the **What name do you want to assign to this jukebox device** prompt, provide a name to identify the autochanger.
8. To enable auto-cleaning, accept the default value of **Yes**, otherwise type **no**.
9. At the **Is (any path of) any drive intended for NDMP use? (yes / no) [no]** prompt, type **yes**.
10. At the **Is any drive going to have more than one path defined? (yes / no) [no]** prompt, type **no** if you will not configure shared devices. Type **yes** to configure shared drives.
11. When prompted, for the first pathname for the NDMP devices in the jukebox, perform the following steps:
  - a. Specify the pathname in the following format:
 

```
NDMP_tape_server_name:device path
```

 where:
    - *NDMP\_tape\_server\_name* is the hostname of the NDMP Server.
    - *device\_path* is the first device path.
  - b. At the **Is this device configured as NDMP** prompt, type **yes**.
  - c. Repeat step a and step b for all NDMP devices in the autochanger.
  - d. When prompted, assign a hardware ID.
  - e. To use DDS:
    - Respond to the prompts as required so that the first host will have access to the shared drive.
    - When prompted to share this drive with another host, type **yes**.
    - When prompted, type the hostname and device path of the second host that will have access to the shared drive.
12. Complete the prompts for the second device.
13. In the **Enter the drive type of drive 1** prompt, specify the number that corresponds to the NDMP device type.

14. If each drive in the autochanger is the same model, then type **yes**. Otherwise, type **no**, and then specify the appropriate device types for each additional autochanger device.
15. When prompted to configure another autochanger, type **no**.

## Changing the block size of an NDMP device

By default, the block size that is used to write data to an NDMP backup is 60KB. With the exception of Celerra, when you specify the `NDMP_AUTO_BLOCK_SIZE=Y` variable for an NDMP client, an NDMP device can use the value that is defined in its Device block size attribute.

### About this task

To determine the block sizes that are supported by the NDMP filer before setting the block size for an NDMP device, consult the applicable vendor documentation.

To change the block size that is defined for the NDMP device, perform the following steps:

### Procedure

1. From the **View** menu, select **Diagnostic Mode**.
2. In the **Devices** window, right-click the NDMP device, and then select **Properties**.
3. On the **Advanced** tab, select a value in the **Device block size** field.

 **Note:** The selected block size must not exceed the block size that is configured on the NAS filer.

4. Click **Ok**.

## Message displayed when CDI enabled on NDMP or file type device

If you enable the CDI feature for an NDMP tape device or file type device (FTD), a message similar to the following appears:

```
nsrd: media notice: The CDI attribute for device "/dev/rmt/3cbn" has
been changed to "Not used".
```

To avoid this message, do not enable the CDI attribute for these device types.

## Configuring NDMP-DSA devices

When you use DSA, NetWorker sends the NDMP data to a NDMP-DSA device, which includes tape, virtual tape, AFTD, and Data Domain devices. The steps to configure a NDMP-DSA device for a specified device type is the same as configuring a non-NDMP device. The *NetWorker Administration Guide* provides detailed information.

## Configuring the Clone Storage Node

When cloning NDMP data, specify the destination storage node, called the clone “write source” (the device that receives the clone data), in the Clone storage nodes attribute. The *NetWorker Administration Guide* provides details.

## Pools requirements for NDMP

When you create a pool for non-NDMP devices, select only the devices that are required by the NDMP clients.

NetWorker cannot send bootstrap and index backups to an NDMP device. When you do not configure a non-NDMP devices or a non-NDMP device is not available to receive the index and bootstrap backups, the NDMP client backup appears to hang. Configure a separate pool to direct the index and bootstrap to a non-NDMP device.

Auto media verification in the Pool resource does not support NDMP.

When an NDMP client backup is a member of a clone-enabled group, configure a clone pool with non-NDMP devices that are local to the NetWorker server to receive the clone bootstrap and index.

## Configure NetWorker for NDMP backup and clone operations

This section explains how to configure NetWorker for NDMP backup and clone operations.

### Performing schedule backup and clone operations

Data Protection Policies provide you with the ability to schedule backup and clone operations, to protect NDMP data.

You can use the NDMP protocol to protect data on NAS devices.

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

### Overview of protection policies

A protection policy allows you to design a protection solution for your environment at the data level instead of at the host level. With a data protection policy, each client in the environment is a backup object and not simply a host.

Data protection policies enable you to back up and manage data in a variety of environments, as well as to perform system maintenance tasks on the NetWorker server. You can use either the **NetWorker Management Web UI** or the NMC **NetWorker Administration** window to create your data protection policy solution.

A data protection policy solution encompasses the configuration of the following key NetWorker resources:

#### Policies

Policies provide you with a service-catalog approach to the configuration of a NetWorker datazone. Policies enable you to manage all data protection tasks and the data protection lifecycle from a central location.

Policies provide an organizational container for the workflows, actions, and groups that support and define the backup, clone, management, and system maintenance actions that you want to perform.

#### Workflows

The policy workflow defines a list of actions to perform sequentially or concurrently, a schedule window during which the workflow can run, and the protection group to which the workflow applies. You can create a workflow when you create a new policy, or you can create a workflow for an existing policy.

A workflow can be as simple as a single action that applies to a finite list of Client resources, or a complex chain of actions that apply to a dynamically changing list of resources. In a workflow, some actions can be set to occur sequentially, and others can occur concurrently.

You can create multiple workflows in a single policy. However, each workflow can belong to only one policy. When you add multiple workflows to the same policy, you can logically group data protection activities with similar service level provisions together, to provide easier configuration, access, and task execution.

## Protection groups

Protection groups define a set of static or dynamic Client resources or save sets to which a workflow applies. There are also dedicated protection groups for backups in a VMware environment or for snapshot backups on a NAS device. Review the following information about protection groups:

- Create one protection group for each workflow. Each group can be assigned to only one workflow.
- You can add the same Client resources and save sets to more than one group at a time.
- You can create the group before you create the workflow, or you can create the group after you create the workflow and then assign the group to the workflow later.

## Actions

Actions are the key resources in a workflow for a data protection policy and define a specific task (for example, a backup or clone) that occurs on the client resources in the group assigned to the workflow. NetWorker uses a work list to define the task. A work list is composed of one or several work items. Work items include client resources, virtual machines, save sets, or tags. You can chain multiple actions together to occur sequentially or concurrently in a workflow. All chained actions use the same work list.

When you configure an action, you define the days on which to perform the action, as well as other settings specific to the action. For example, you can specify a destination pool, a retention period, and a target storage node for the backup action, which can differ from the subsequent action that clones the data.


When you create an action for a policy that is associated with the virtual machine backup, you can select one of the following data protection action types:

- Backup — Performs a backup of virtual machines in vCenter to a Data Domain system. You can only perform one VMware backup action per workflow. The VMware backup action must occur before clone actions.
- Clone — Performs a clone of the VMware backup on a Data Domain system to any clone device that NetWorker supports (including Data Domain system or tape targets). You can specify multiple clone actions. Clone actions must occur after the Backup action.

You can create multiple actions for a single workflow. However, each action applies to a single workflow and policy.

The following figure provides a high level overview of the components that make up a data protection policy in a datazone.

**Figure 12** Data Protection Policy



## 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 13** 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.

## Overview of configuring a new data protection policy

### About this task

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

**Figure 14** 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*.

## Strategies for traditional backups

The primary considerations for a traditional backup strategy are the groups of Client resources, the workflows that define the series of actions that are associated with the backup, and the schedule for the backup.

### 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 traditional backups

A protection groups for traditional backups identifies the client resources to back up.

Traditional backups support the following types of protection groups:

- **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.

Create multiple groups to perform different types of backups for different Client resources, or to perform backups on different schedules. For example:

- Create one group for backups of clients in the Accounting department, and another group for backups of clients in the Marketing department.
- Create one group for file system backups and one group for backups of Microsoft Exchange data with the NetWorker Module for Microsoft.
- Create one group for a workflow with backups actions that start at 11 p.m., and another group for a workflow with backup actions that start at 2 a.m.

 **Note:** A Client resource can belong to more than one group.

### 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.

### About this task

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](#).

## Actions sequences in traditional backup workflows

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


A workflow for a traditional backup can optionally include a probe or check connectivity action before the backup, and a clone action either concurrently with or after the backup.

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

### Workflow path from a traditional backup action

The only action that can follow a traditional backup is a clone action.

**Figure 15** Workflow path from a traditional backup action



### Creating a check connectivity action

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

#### Before you begin

Create the policy and 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, monthly, or reference schedule for the action:
  - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
  - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
  - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
9. Click the icon on each day to specify whether to check connectivity with the client.

The following table provides details on the icons.

**Table 28** Schedule icons

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

To check connectivity every day, select **Execute** from the list, and then click **Make All**.

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 successful connectivity is achieved with all clients in the assigned group, select the **Succeed only after all clients succeed** checkbox.
  - To specify that the connectivity check is successful if connectivity is achieved with one or more clients in the assigned group, clear the checkbox.
12. Click **Next**.  
The **Specify the Advanced Options** page appears.
13. Optionally, 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.  
  
For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.  
  
If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.  
  
When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.  
  
Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.  
  
For Isilon clients that use checkpoint restart, on the **Advanced** tab, set the value of the **Client retries** attribute to a number greater than 0.
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) 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.
- If a schedule is associated to an action, then override option is disabled.

19. Click **Next**.

The **Action Configuration Summary** page appears.

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

#### After you finish

Optionally, 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 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. In-built commands from NetWorker should be avoided as probe command.

#### 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, monthly, or reference schedule for the action:
    - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
    - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
    - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
  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 29** 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 action 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 action if any of the probes 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

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.
- If a schedule is associated to an action, then override option is disabled.

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 traditional backup action

A traditional backup is a scheduled backup of the save sets defined for the Client resources in the assigned group for the workflow.

#### Before you begin

- Create the policy and workflow that contain the action.

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

### 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 **Backup**.
6. From the secondary action list, select the backup type, for example, **Traditional**.
7. (Optional) From the **Force Backup Level** list select a backup level.

For workflows that have more than one scheduled backup within a 24-hour period, use the **Force Backup Level** attribute to allow more than one backup to occur at two different backup levels in a 24-hour period. When you select a backup level in the **Force Backup Level** attribute, the first backup is performed at the scheduled backup level. Each subsequent occurrence of the backup action in the next 24 hours occurs at the level defined in the **Force Backup Level** attribute. For example, if the level defined by the schedule is Full and the **Force Backup Level** attribute is set to Incr, the first backup started by the action occurs at a level full and subsequent backups, within 24 hours of the start of the full backup are incremental. By default this option is cleared, which means that if the action runs multiple backup operations in a 24 period, all the backups occur at the scheduled backup level.

8. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
9. 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.
10. Specify a weekly, monthly, or reference schedule for the action:

- To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
  - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
  - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
11. To specify the backup level to perform when **Define** option under **Select Schedule** is selected, click the icon on each day. If it is **Select** option under **Select Schedule**, choose the customized NSR schedule using the drop down, after specifying the backup level within the NSR schedule resource.

The following table provides details about the backup level that each icon represents.

**Table 30** Schedule icons

| Icon                                                                                | Label                                                                                                                                                         | Description                                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|    | <b>Full</b>                                                                                                                                                   | Perform a full backup on this day. Full backups include all files, regardless of whether the files changed.                                                                                                               |
|    | <b>Incr</b>                                                                                                                                                   | Perform an incremental backup on this day. Incremental backups include files that have changed since the last backup of any type (full or incremental).                                                                   |
|  | <b>Cumulative Incr</b>                                                                                                                                        | Perform a cumulative incremental backup. Cumulative incremental backups include files that have changed since the last full backup.                                                                                       |
|  | <b>Logs Only</b>                                                                                                                                              | Perform a backup of only database transaction logs.                                                                                                                                                                       |
|  | <b>Incremental Synthetic Full</b><br> <b>Note:</b> Not supported for NDMP. | Perform an incremental synthetic backup on this day. An incremental synthetic full backup includes all data that changed since the last full backup and subsequent incremental backups to create a synthetic full backup. |
|  | <b>Skip</b>                                                                                                                                                   | Do not perform a backup on this day.                                                                                                                                                                                      |

To perform the same type of backup on each day, select the backup type from the list and click **Make All**.

NetWorker does not support the use of synthetic full backup levels for NDMP data.

Celerra, Isilon, VNX, Unity, and NetApp filers with NDMP version 4 or later support token-based backups (TBB) to perform NDMP full and incremental backups. NetWorker supports

the same number of incremental levels that the NAS vendor supports. Celerra, Isilon, and NetApp documentation provide the maximum number of incremental levels that the TBB incremental backup can support.

When you configure TBB after you update the NetWorker server from 7.6 SP1 or earlier, the first incremental backup does not occur until after one complete full backup.

Filers that do not support TBB, do not support incremental backups. If you select the level **Incr**, the NetWorker server performs a full backup.

Verify that the NAS storage vendor supports NDMP incremental backups before you use this feature.

12. Click **Next**.

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

13. From the **Destination Storage Node** box, select the storage node with the devices on which to store the backup data.
14. From the **Destination Pool** box, select the media pool in which to store the backup data.
15. From the **Retention** boxes, specify the amount of time to retain the backup data.

After the retention period expires, the save set is removed from the client file index and marked as recyclable in the media database during an expiration server maintenance task.

When you define the retention policy an NDMP client, consider the amount of disk space that is required for the client file index. NDMP clients with several thousands of small files have significantly larger client file indexes on the NetWorker server than a non-NDMP client. A long retention policy for an NDMP client increases disk space requirements on the file system that contains the client file indexes.

16. From the **Client Override Behavior** box, specify how NetWorker uses certain client configuration attributes that perform the same function as attributes in the Action resource:
  - **Client Can Override**—The values in the Client resource for **Schedule**, **Pool**, **Retention policy**, and the **Storage Node** attributes take precedence over the values that are defined in the equivalent Action resource attributes.
    -  **Note:** If the NetWorker policy action schedule is set to the **Skip** backup level, the **Client can Override** option is not honored. For NetWorker to consider the **Client can Override** option, change the action schedule to a different level.
  - **Client Can Not Override**—The values in the Action resource for the **Schedule**, **Destination Pool**, **Destination Storage Node**, and the **Retention** attributes take precedence over the values that are defined in the equivalent Client resource attributes.
  - **Legacy Backup Rules**—This value only appears in actions that are created by the migration process. The updating process sets the **Client Override Behavior** for the migrated backup actions to **Legacy Backup Rules**.

17. Click **Next**.

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

18. 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.



19. 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.

20. 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.

21. 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

22. 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.

23. 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.

24. 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.

25. (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.

26. (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.
- If a schedule is associated to an action, then override option is disabled.

27. 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**.

28. 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 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.

29. Click **Next**.

The **Action Configuration Summary** page appears.

30. 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 after the backup. A clone action is the only supported action after a backup action in a workflow.

### Cloning NDMP save sets

You can clone Direct-NDMP and NDMP-DSA save sets by using the same methods used to clone non-NDMP save sets.

Before you clone NDMP save sets, review these requirements:

- To clone Direct-NDMP or Three-party backup data:
  - The source NAS must run NDMP version 3 or later.
  - The destination NAS can run any version of NDMP, but you cannot clone a volume cloned with NDMP earlier than version 3 to another volume.
  - You cannot clone NDMP save sets to a non-NDMP device.
  - You can clone NDMP tapes from one NDMP host to another NDMP host of the same type. For example, you can clone tapes from a NetApp filer with an attached library to another NetApp filer or to the same filer.
- You require two NDMP devices to clone the NDMP save sets, one device to perform the read operation and one device to perform the write operation.
- You must clone NDMP-DSA backups to non-NDMP devices. You can however, clone NDMP-DSA save from one type of tape device to another. For example you can clone save sets on a DLT device to an AIT device.
- Use the `nsrclone` program to clone NDMP save sets from a command prompt. The *NetWorker Command Reference Guide* or the UNIX man pages provide more information on `nsrclone` usage.

#### 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

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 **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, monthly, or reference schedule for the action:
  - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
  - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
  - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
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 31** 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 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

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 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.

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.
- If a schedule is associated to an action, then override option is disabled.

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 traditional backup workflows

**Figure 16** Traditional backup workflow



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.

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**.

## Creating and configuring the NDMP client resource

To protect an NDMP filer, a client resource is necessary. This resource stores configuration information for the filer. Use the NMC Client Configuration wizard to create the NDMP client or create the client manually. It is recommended that you use the NMC Client Configuration wizard to create NDMP clients. Isilon SmartConnect hostnames are supported in the Client Name field. If SmartConnect hostname is not available, use a hostname that is associated with a non-aggregated (non-teamed) network connection to one of the Isilon cluster nodes that has a fixed IP address.

### Isilon SmartConnect

NetWorker NDMP backups can utilize the Isilon SmartConnect feature. Isilon SmartConnect allows the filer to distribute the backup sessions across the backup accelerator nodes. The distribution algorithm is configured and controlled by the filer. NetWorker backups perform a SmartConnect lookup for both single and multi-stream backups.

### Using the Client Configuration wizard

#### About this task

Use the NMC Client Configuration wizard to create the NDMP client.

#### Procedure

1. From the **Administration** window in NMC, click **Protection**.
2. In the expanded left pane, select **Clients**, and then select **Protection > New Client Wizard**.
3. On the **Specify Client Information** window:
  - a. In the **Client Name** field, specify the hostname of the filer.
  - b. (Optional) Add comments in the **Comment** field.
  - c. (Optional) In the **Tag** field, specify the name of the tag for the dynamic group in which you want to add this client.
  - d. (Optional) In the **Groups** area, select an existing group, in which to add the client.
  - e. In the **Type** area, select **NDMP**, and then click **Next**.
4. On the **Specify the NDMP Client Credentials** window:
  - a. In the **NDMP User Name** field, specify a valid NAS administrator account.
  - b. In the **NDMP Password** attribute, specify the password for the NAS administrator account, and then click **Next**.

The *OneFS Users Guide* on Support site describes how to create NDMP administrators.

5. In the **Specify the NDMP Client Backup Options** window:
  - a. In the **NDMP backup type** attribute, select or specify the backup type:
    - `tar`—Use this backup type for Checkpoint Restart.

 **Note:** Isilon supports only `tar` backups.
  - b. In the **NDMP Array Name** field, specify the logical name that is assigned to the NDMP NAS array.

The **NDMP Array Name** field enables you to configure the same NAS device with multiple NDMP clients that have different host IDs.

**Note:** NDMP clients that use the same NAS device must have the same NDMP array name.

- c. Review the **App Info** options and disable options, as required. It is recommended that the default options remain enabled.

**Table 32** Application information variable types

| App Info Type          | Description                                                                                                                                                                                                        |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HIST                   | Enables the backup of index data. Leaving this option selected allows for browsing of individual files and directories during recovery. If you do not select this option, you can only perform saveset recoveries. |
| UPDATE                 | Enables the backup process to update the last backup dates in database on the NDMP client, after the backup completes. Only applies to NetApp and has no effect when backing up Isilon.                            |
| DIRECT                 | Enables DAR or DDAR support. <a href="#">NDMP recovery requirements for Isilon</a> on page 174 and <a href="#">DAR and DDAR</a> on page 175 provide more information.                                              |
| Use Token-Based Backup | Enables the NDMP backup to use last backup time tokens to decide what files to backup. Token-Based backup allows for Incremental level backups. Isilon supports Token-Based backup.                                |

- d. In the **Advanced App Info** field, specify additional NAS specific environments variables, one per line. The following table provides a list of the available **Application Information** environment variables for each NAS.

**NOTICE** Environment variables are case-sensitive. Use an equal (=) sign to separate the environment variable name from its value.

**Table 33** Isilon Information variables

| Variables                       | Definition                                                                                                                                                                                                                                       |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>DIRECT=y</i>                 | Required.                                                                                                                                                                                                                                        |
| <i>HIST=F/D</i>                 | Set this variable to <i>F</i> for Checkpoint Restart backups to ensure the use of path based file history. For all other backups, the value of HIST can be either <i>F</i> or <i>D</i> .                                                         |
| <i>FILES=pattern</i>            | Optional. Use this variable to back up only files that match the defined pattern. You can use wildcards in the pattern definition.                                                                                                               |
| <i>PER_DIRECTORY_MATCHING=y</i> | Optional. Use this variable along with the <i>FILES=pattern</i> variable. Matches the pattern defined by FILES across directories.                                                                                                               |
| <i>USE_TBB_IF_AVAILABLE=n</i>   | Optional. The NetWorker software automatically enables TBB for Isilon filers. Specify this variable to disable TBB support for incremental backups. When you specify this value, the backup reverts to the native level-based backup of the NAS. |

**Table 33** Isilon Information variables (continued)

| Variables                                              | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>NDMP_AUTO_BLOCK_SIZE=Y</i>                          | Optional. Specify this variable to override the default block size of 60 KB when writing NDMP backups to an NDMP device. Uses the block size value defined in the Device block size attribute when you labeled the NDMP volume. <a href="#">Configuring NDMP devices</a> on page 37 provides more information.                                                                                                                                                                                                                         |
| <i>BACKUP_MODE=SNAPSHOT</i><br><i>BACKUP_OPTIONS=7</i> | Optional. Specify these variables to enable NetWorker to perform a Fast Incremental when the level defined for a backup is incremental. Fast Incremental has no impact on throughput. However, it saves time at the start of Incremental level backups, which lowers the overall backup time for the save set. Fast Incremental will help most with Incremental level backups of large, dense save sets (many millions of files) that change frequently. It will not help with save sets containing a few large files and static data. |
| <i>NDR_DSA_NODE</i>                                    | When this <i>NDR_DSA_NODE</i> is used, nsrdsa_save binary is started on the specified sn-server during the backup. During recover, nsrdsa_recover binary is started on the specified sn-server.                                                                                                                                                                                                                                                                                                                                        |
| <i>NSR_DSA_REC_SERVER</i>                              | During recover, nsrdsa_recover binary is started on the specified remote storage node.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

6. Click **Next**.

7. On the **Select the NetWorker Client Properties** window:

a. In the **Priority** field, specify the order in which the NetWorker server contacts clients in a protection group for backup. The attribute can contain a value between 1 and 1,000. The lower the value, the higher the priority.

b. In the **Parallelism** attribute:

- For Direct-NDMP, set the **Parallelism** attribute to 1.
- For NDMP-DSA, the parallelism value depends on the NAS capabilities and set parallelism to a value that is appropriate for the NAS. Parallelism values of 4 to 8 are common. In general, the best parallelism setting depends on filer configuration and the amount of installed RAM.

c. In the **Remote Access** attribute:

Specify the root account on Linux/UNIX, and/or the administrator account on Windows, of any computer that you will use to browse backups of the Isilon. Specify each account on a separate line. For example:

```
administrator@windows_hostname
```

```
root@linux_hostname
```

- d. Select the **Data Domain Interface**. This option specifies the protocol to use if you send the backup data to a Data Domain Device. Available selections are Any, Fibre Channel, or IP.
- e. Do not select the **Block Based Backup** or **Client Direct** options, as they do not apply to NDMP backups.
8. On the **Specify the File System Objects** window, specify objects to backup:
  - List the directories under `/ifs`, one per line.
  - Use the directory names as configured on the Isilon. You can use the Isilon OneFS web administration interface to browse `/ifs`.
  - Names are case sensitive.
  - Do not list the top-level `/ifs` directory.
  - You cannot specify a share name.
  - Isilon does not currently support browsing from NetWorker.
9. Click **Next**.
10. On the **Client Configuration Summary** window, review the attributes, and then click **Create**.
11. On the **Client Configuration Results** window, review the results, and then click **Finish** to exit the wizard.

[Troubleshooting NDMP configuration and backup failures for PowerStore, Unity, Celerra, VNX, and VNXe](#) describes how to resolve errors that you may experience when you configure the NDMP client.

## Isilon Environment Variables

Restart the NetWorker services after setting the following environment variables:

### Isilon Environment variables and definitions

| Variables                                                                                            | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>NSR_NDMP_TMP_DIR=&lt;path to tmp&gt;</code>                                                    | Use this variable to redirect the tmp files going to <code>/nsr/tmp</code> (optional).<br> <b>Note:</b> Path to tmp should be created manually.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <code>NSR_NDMP_RESTORE_LIMIT</code>                                                                  | Use this variable while recovering large NDMP savesets using DAR recovery. The <code>NSR_NDMP_RESTORE_LIMIT</code> environment variable is used to limit memory consumption during recoveries involving a large number of index entries (millions). This is specifically useful if there is not enough swap space or memory for the number of index entries selected for the recovery. If the variable is not set, the recovery might fail with an <code>out of memory</code> error message.<br> <b>Note:</b> <code>NSR_NDMP_RESTORE_LIMIT</code> can be set to 50,000 or 1,000,000, but less than 3,000,000. |
| <code>NSR_NDMP_DDAR=y</code><br>This environment variable must be set in the operating system before | Use this variable to recover the permissions of the parent directories of selected files. Specify this                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

| Variables                                                                                                                                              | Definition                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| invoking either the recover or winworkr program.                                                                                                       | variable to perform a DDAR recovery when (optional): <ul style="list-style-type: none"> <li>You set the <i>DIRECT=y</i> variable during the backup</li> <li>The DART version is 5.5 and later</li> </ul> <p><b>Note:</b> Do not specify <i>NSR_NDMP_DDAR</i> when you also use <i>NSR_NDMP_RECOVER_DIR</i>.</p>              |
| <i>NSR_NDMP_RECOVER_NO_DAR=y</i>                                                                                                                       | Define this variable to perform an non-DAR recovery when you set the <i>DIRECT=y</i> variable during the backup (optional).                                                                                                                                                                                                  |
| <i>NSR_NDMP_RECOVER_DIR=y</i><br>This environment variable must be set in the operating system before invoking either the recover or winworkr program. | Specify this variable to perform a DAR recovery when (optional): <ul style="list-style-type: none"> <li>When you set the <i>DIRECT=y</i> variable during the backup</li> <li>The DART version is 5.5 and later</li> </ul> <p><b>Note:</b> Do not use <i>NSR_NDMP_RECOVER_DIR</i> when you also use <i>NSR_NDMP_DDAR</i>.</p> |
| <i>NSR_NDMP_NDAR=y</i>                                                                                                                                 | It can be used instead of <i>NSR_NDMP_DDAR</i> when both full and incremental backups are present. <i>NSR_NDMP_DDAR</i> works only with full backups                                                                                                                                                                         |

**Note:** To set environment variable on Linux NetWorker server edit `/opt/nmr/admin/networkerrc` and add the variable as shown below: ,

- `NSR_NDMP_DDAR=y`
- `export NSR_NDMP_DDAR`

To set environment variable on Windows NetWorker server navigate to Control Panel > System and Security > System > Advanced system settings > Environment Variables > System variables > New.

- Variable Name = `NSR_NDMP_TMP_DIR`
- Variable Value = `C:\NW\nsr\tmp`

## Performing post Client Configuration Wizard steps

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

## Modifying the Storage Node

### About this task

On the **Globals (2 of 2)** tab, specify the storage node in the **Storage Nodes** attribute.

The attribute value depends on the type of backup:

- When you perform Direct-NDMP backups with NDMP devices, specify the hostname of the NAS that manages the tape device or autochanger.
- For three-party backups, specify the destination host first.
- For NDMP-DSA backups, specify the hostname of the storage node that manages the tape device or autochanger. If the NetWorker server is the storage node, specify *nsrserverhost*.

**NOTICE** For NDMP-DSA backups, the NetWorker software uses the **Storage Node** attribute field of the NDMP client to determine which host receives the backup data. The `nsrndmp_save` command does not require the `-M` and `-P` options. If you specify the `-M` and `-P` options, they override the **Storage Node** attribute value.

## Enabling Checkpoint Restart

To allow a failed backup for a client to restart from a known good point, you must enable Checkpoint Restart for the NetWorker client resource and configure the number of automatic retries for the backup action in the data protection policy.

### Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Clients**.
3. Right-click the client resource and select **Properties**.

The **Client Properties** dialog box appears.

4. On the **General** tab, select the **Checkpoint enabled** checkbox.

**Note:** When you enable **Diagnostic Mode** the **Checkpoint granularity** attribute appears. This attribute does not apply to NDMP.

5. Define the interval at which checkpoints are written during the backup:
  - a. Select the **Apps & Modules** tab on the **Client Properties** dialog box.
  - b. In the **Application information** attribute, specify the `CHECKPOINT_INTERVAL_IN_BYTES` variable.

For example, to write a checkpoint after every 1,000,000 bytes, type:

```
CHECKPOINT_INTERVAL_IN_BYTES=1000000
```

To define the value by using a different multiplier, specify the multiplier with the numeric value. Supported multipliers include KB, MB, GB, TB, kb, mb, gb, and tb. For example, to write a checkpoint after every 1 GB, type:

```
CHECKPOINT_INTERVAL_IN_BYTES=1GB
```

**Note:** The checkpoint interval value is automatically rounded up to a multiple of the tape block size.

6. Click **OK** on the **Client Properties** dialog box.

## Adding NDMP Client Properties

### Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Clients**.
3. Right-click a **Client** and select **New Client Properties**. The **Client Properties** screen appears.
4. Select the **Apps & Modules** tab. Select the **NDMP** attributes of the **Client**. The **NDMP** options are as follows:
  - **NDMP**—Select this box to indicate whether this client is an NDMP client.
  - **NDMP multistreams enabled**—Select this box to indicate whether the Isilon client is using the NDMP multistream feature.

**Note:** Increasing the number of streams does not always increase performance. Depending on your environment, performance might be reduced when the number of streams is increased, especially for a high density file system backup and recover.

If the Isilon has one or more large directories under `/ifs`, consider creating two or more clients with the same name but with different save sets, and schedule them separately.

- To start a multistream backup from the command line interface, add option `-A multistreams=<n>` to the command `nsrndmp_save` where `n`=Number of streams. For example:

```
nsrndmp_save -T <Backup_type> -c <Isilon_client> -A multistreams=<n>
<Backup_Path>
```

- To start a multistream backup from NMC GUI or a scheduled backup, set the **Parallelism** value in the **Globals (1 of 2)** tab.
  - **NDMP log successful file recovery**—By default, NetWorker does not print each successfully recovered file name in the log messages, because this logging impacts performance and takes up space. To enable the logging of successful recoveries for each file, select this checkbox.
  - **Disable IPv6**—Check this box to disable IPv6 on NDMP backup and recovery.
  - **NDMP array name**—This name is the logical name that is assigned to the array in NDMP NAS array configurations.
  - **NDMP vendor information**—This attribute contains NDMP client vendor information.
5. Click **OK**.

**Note:** Client parallelism feature is supported on Isilon filer and is limited to the NDMP multi-stream backups. This feature does not support Celerra, VNX, and VNXe NAS. To use it, the user must provide the `NMS:streams_per_ss` parameter in Save operations: box in the Apps and Modules UI window.

```
NMS:streams_per_ss=4,/....>
NMS:streams_per_ss=4,ALL
NMS:streams_per_ss=4,all
NMS:streams_per_ss=4,*
```

## Configuring the NDMP client manually

It is recommended that you create Network Data Management Protocol (NDMP) clients by using the Client Configuration wizard. If you create the NDMP client manually, then the configuration details for each attribute in the Client Configuration wizard apply when you create the client manually.

Review this information before you configure an NDMP client manually:

- For an NDMP configuration that includes Storage Node resources, configure a Client resource for each storage node that you define for an NDMP backup and clone operation.
- For NDMP three-party storage nodes that use NDMP devices, repeat these steps for each NDMP storage node.
- For NDMP-DSA storage nodes, create the NetWorker Client resources in the same manner as non-NDMP clients. The *NetWorker Administration Guide* provides details on how to create a non-NDMP Client resource.
- NDMP does not support the use of directives including AES encryption. The NetWorker software ignores any value that you define in the **Directives** attribute for an NDMP client.

- When you select **Checkpoint enabled** on the **General** tab, do not modify the **Checkpoint granularity** attribute. NDMP backups do not support checkpoint granularity and the NetWorker software ignores any value that you define for this attribute.
- If the NAS supports the NDMP snapshot management extension, then you can browse and mark individual file systems for backup instead of specifying the save sets in the **Save set** attribute. You cannot use the **Save set browse** icon to browse the NDMP file system until you:
  - Select the **NDMP** checkbox, on the **Apps & Modules** tab.
  - Specify the NDMP username and password in the **Remote user and password** fields on the **Apps and Modules** tab.

**Note:**

- Celerra, VNX, VNXe, and NetApp C-mode do not support Snapshot Management Extension. Only NetApp 7-Mode supports Snapshot Management Extension.
- Isilon, Celerra, VNX, VNXe, and NetApp C-mode filers do not allow you to browse. Only NetApp 7-Mode allows you to browse.

## Performing manual NDMP backups

After you configure the NetWorker server for NDMP backup data operations, you can perform manual NDMP backups.

On Windows, you can manually back up NDMP data by using the NetWorker User program. The method to backup NDMP data is the same as a non-NDMP local backup. You cannot perform a three-party backup with the NetWorker User program.

On Windows and UNIX, you can perform a manual backup from a command prompt by using the `nsrndmp_save` command.

Before performing a manual backup by using the `nsrndmp_save` command or the NetWorker User program, review these requirements:

- You can only perform manual Direct-NDMP backups from a NetWorker server.
- You can start a manual NDMP-DSA backup from a NetWorker server, storage node, or client. When you do not start the NDMP-DSA backup from the NetWorker server, the `servers` file on the NetWorker server and storage node, must contain the hostname of the host that initiates the backup.
- Before you perform a manual backup, you must configure the NDMP client on the NetWorker server. Manual backups use client configuration information for example, the variables that are defined in the **Application Information** attribute of an NDMP client.
- Direct-NDMP and three-party NDMP backups support manual DAR backups when the NDMP client contains the `DIRECT=Y` and `HIST=Y` environment variables in the **Application Information** attribute for the NDMP client.

**NOTICE** To use DAR, the NAS filer must use NDMP version 4. The *NetWorker E-LAB Navigator* describes how to determine if a particular NDMP vendor supports DAR.

## Performing an NDMP backup from the command line

### About this task

Use the `nsrndmp_save` command to perform a manual command line NDMP backup.

The `nsrndmp_save` command does not back up the bootstrap. Without the bootstrap, you cannot perform a disaster recovery of the NetWorker server. To back up the bootstrap, run the `nsrpolicy -G nsrpolicy policy_name start` command from the NetWorker server. The



`nsrpolicy` command uses the attribute values specified for the policy. For example, the pool and schedule values.

To perform an NDMP backup from the command prompt, use the following syntax:

```
nsrndmp_save -T backup_type -s NetWorker_servername -c clientname -l
backup_level -t date_time -g nsrpolicy path
where:
```

- *backup\_type* is a supported backup type for the NAS filer:  
Isilon accepts `tar` and `dump` backup types. However, if you select `dump` as the backup type, Isilon stores the backup data only in `tar` format.
- *backup\_level* is a full for a full backup, `incr` for an incremental backup. Each NAS supports full backups.  
**Note:** NetWorker only supports or allows Full, Incremental, and Differential level 1 NDMP backups of all NAS.
- *date\_time* is the date and time of the last backup, which is enclosed in double quotes. Specify this value for `incr` level backups. When you do not specify the date and time, the backup is a native NDMP level-based backup.  
**NOTICE** During a NetWorker scheduled policy backup, the NetWorker software supplies the date and the time information, and incremental and level backups work as expected.

Use one of these methods to determine the date and time of the last NDMP backup:

- Review the `daemon.raw` file on the NetWorker server or the savegroup completion report for a line similar to the following:

```
42920:nsrndmp_save: browsable savetime=1296694621
```

Use the value after `savetime=` with the `-t` option.

- Specify the date and time of the last backup reported by the `mminfo` command for the NDMP save set.

**Note:** You can force the NDMP backup and recovery to ignore IPv6 and instead use IPv4 in one of two ways:

1. Add the `-f` option to the `nsrndmp_save` or the `nsrndmp_recover` commands as required.
2. On the Apps & Modules tab in the **Client Properties** window, select **Disable IPv6**.

## Example of an NDMP backup

To perform an incremental backup of a NetApp client that is named `mynetapp`, perform the following steps:

1. Determine the time of the last full backup:

```
mminfo -v -c mynetapp
```

**Table 34** NDMP backup

| client   | date     | time     | size   | ssid       | fl   | lvl  | name      |
|----------|----------|----------|--------|------------|------|------|-----------|
| mynetapp | 08/16/15 | 15:23:58 | 1853MB | 3864812701 | cbNs | full | /.../set1 |
| mynetapp | 08/17/15 | 15:39:58 | 815MB  | 3848036430 | cbNs | incr | /.../set2 |

2. Specify the last backup time in `nsrndmp_save` command:

```
nsrndmp_save -T dump -s my_nwserver -c mynetapp -l incr -t "02/16/11
15:23:58" -g mygroup path
```

For NDMP-DSA backups, the NetWorker software uses the Storage Node attribute field of the NDMP client to determine which host receives the backup data. The `nsrndmp_save` command does not require the `-M` and `-P` options. If you specify the `-M` and `-P` options, they override the **Storage Node** attribute value. The *NetWorker Command Reference Guide* and the `nsrndmp_save` man page on UNIX provide more information.

## Troubleshooting NDMP configuration and backup failures for Isilon

This section provides a list of the possible causes and the resolutions for NDMP backup failures.

### Unable to connect to NDMP host *hostname*

This message appears when the NetWorker server cannot create or modify an NDMP client.

To resolve this issue ensure that the environment meets the following requirements:

- Username and password specified for the client is correct and has sufficient permissions to perform NDMP operations.
- NDMP service is running on the filer.

### Cannot perform NDMP backup after the NetWorker server licenses expire

If a NetWorker sever running in evaluation mode expires before you authorize the server, NDMP devices remain disabled after the addition of the required licenses and authorization of the NetWorker server.

To re-enable NDMP devices, perform the following steps:

1. To connect to the NetWorker server, use NMC, and then click the **Devices** button.
2. In the **Devices** windows, right-click the NDMP device, and then select **Properties**.
3. Click the **Configuration** tab, and then set the **Target Sessions** attribute to 1.
4. Click the **General** tab, and then in the **Enabled** section, select **Yes**.
5. Click **Ok**.

### Failed to store index entries

This error message occurs in the `daemon.raw` file when an index backups fails due to an insufficient amount of swap space.

To resolve this issue, increase the amount of swap space available to the NetWorker server.

 **NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

### IO\_WritePage write failed - No space left on device (28): No space left on device

This error message appears in the `daemon.raw` file when the index backup fails. There is insufficient temporary space to store the index entries before the NetWorker software commits the information into the client file index.

To resolve this issue, specify a new the temp directory with sufficient disk space in one of the following ways:

- Define the `NSR_NDMP_TMP_DIR` environment variable in the Application Information attribute of the client.

- Define the `NSR_NDMP_TMP_DIR` as an operating system environment variable on the NetWorker server.

[Memory and space requirements for NDMP FH updates](#) on page 23 describes how to determine the amount of disk space the NetWorker software requires to temporarily store client files index entries.

**NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

## Error reading the FH entries from save through stdin

This error message appears in the `daemon.raw` file of the NetWorker server when there is a communication error between the `nsrndmp_save` and `nsrndmp_2fh` processes.

Resolve any communication or connection issues, then retry the backup.

**NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

## Cannot find file history info for file name...You may still be able to recover this file with a save set recovery

This error message appears in the `daemon.raw` file of the NetWorker server when file history (FH) information is missing or corrupted for the file that is specified in the error message. For example, NetWorker cannot update the client file index (CFI) with FH information when a backup process interruption occurs during the failover of a clustered NetWorker environment.

You cannot perform an Network Data Management Protocol (NDMP) file-by-file recover or a save set recover when the CFI does not contain the associated FH information.

To recover this file, perform a save set recover from the command prompt. [Performing an NDMP save set recovery from the command prompt](#) on page 105 provides for further information.

**NOTICE** The NetWorker server does not delete the FH files that are stored in the `tmp` directory when the CFI updates fail.

## nsrndmp\_save: data connect: failed to establish connection

This error message appears in the `daemon.raw` file of the NetWorker server for several reasons:

- Network connectivity or name resolution issues exist between the NetWorker server and the NDMP client.
- You specified an incorrect NDMP username or password specified for the NDMP client.
- The NDMP service is not started on the NAS filer.
- The NetWorker server cannot communicate with the NAS filer over port 10000.
- A free port in the NetWorker server's default port range (7937-9936) is not available during an NDMP-DSA backup.  
The *NetWorker Security Configuration Guide* provides more information about NDMP port requirements and configuration.
- A misconfigured loop router. For a Celerra filer, the server route command utility configures the loop router. For NetApp, the route utility configures loop back router. The value of this setup is network-specific and depends on the number of switches and hubs between the NAS filer, NetWorker server, and NetWorker storage node.
- On the host where DSA is running, if the hostname is present in the hosts file, the `nsrdsa_save` process uses this name during backup. The DSA host passes the loopback entry to the NDMP data server and the connection fails. To resolve this issue, remove the hostname from the localhost list.

Knowledge base articles on the Support website provides detailed troubleshooting information for this error message and other failed to establish connection failures that you might encounter during an NDMP backup.

### nsrndmp\_save: get extension list: communication failure

This message appears during a NDMP local backup when NetWorker cannot determine the filer name.

To resolve this issue, perform the following steps:

1. From a command prompt on the NetWorker server, type:

```
ndmpsup -c NDMP_hostname -o output_filename
```

For example:

```
ndmpsup -c myfiler.mnd.com -o ndmpsup.txt
```

2. Edit the output file that the `ndmpsup` command generates and search for the string **Vendor Name**. Make note of the reported Vendor Name.

For example:

```
Vendor Name = BlueArc Corp
```

3. Change to the `/nsr/debug` directory on UNIX or the `NetWorker_installation_dir\nsr\debug` directory on Windows.

4. Create new empty file and name it with the following format:

```
ndmpgettextlist_disable_VENDOR_NAME
```

where you replace `VENDOR_NAME` with the vendor name of the filer reported in the `ndmpsup` output file.

For example, to create this file for a BlueArc filer on UNIX, type:

```
touch "ndmpgettextlist_disable_BlueArc Corp"
```

## Monitoring NetWorker Server activities in the Administration window

The **Monitoring** window in the NetWorker **Administration** application enables you to monitor the activities of an individual NetWorker Server.

### About this task

The **Monitoring** window provides the following types of activity and status information:

- Data protection policies, workflows, and individual actions.
- Cloning, recovering, synthetic full backups, and browsing of client file indexes.
- Operations that are related to devices and jukeboxes.
- Alerts and log messages.

You can also perform some management operations from the **Monitoring** window, for example, starting, stopping, or restarting a data protection policy.

### Procedure

1. From the **NMC Console** window, click **Enterprise**.
2. In the **Enterprise** view, right-click the NetWorker Server, and then select **Launch Application**.

The **Administration** window appears.

3. To view the **Monitoring** window, click **Monitoring**.

## About the Monitoring window

On the **Administration** window taskbar, select **Monitoring** to view the details of current NetWorker server activities and status, such as:

- Policies and actions.
- Cloning, recovering, synthetic backups, checkpoint restart backups, and browsing of client file indexes.
- Alerts and log messages, and operations that are related to devices and jukeboxes.

While the **Monitoring** window is used primarily to monitor NetWorker server activities, it can also be used to perform certain operations. These operations include starting, stopping, or restarting a workflow.

The **Monitoring** window includes a docking panel that displays specific types of information. Select the types of information you want to view from the docking panel.

A portion of the **Monitoring** window, which is known as the task monitoring area, is always visible across all windows. A splitter separates the task monitoring area from the rest of the window. You can click and move the splitter to resize the task monitoring area. The arrow icon in the upper right corner of the **Monitoring** window allows you to select which tasks you want to appear in this view.

Smaller windows appear within the **Monitoring** window for each window. Each smaller window, once undocked, is a floating window and can be moved around the page to customize the view. You can select multiple types from the panel to create multiple floating windows that can be viewed simultaneously. The following table describes the various types of information available in the docking panel, and the details each one provides.

**Table 35** Monitoring window panel

| Window                  | Information provided                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Policies/Actions</b> | The <b>Policies</b> tab provides you with status information about all configure policies and the associated workflows and actions. The <b>Actions</b> tab provides you with status information for all actions. <a href="#">Policies/Actions pane</a> on page 81 provides more information.                                                                                                                                   |
| <b>Sessions</b>         | Allows you to customize whether to display all session types, or only certain session types. The information that is provided depends on which session type you select. For example, if you select <b>Save Sessions</b> , the window lists clients, save sets, groups, backup level, backup start time, duration of the backup, devices, rate, and size. <a href="#">Sessions window</a> on page 85 provides more information. |
| <b>Alerts</b>           | Lists the priority, category, time, and message of any alerts. <a href="#">Alerts pane</a> on page 86 provides more information.                                                                                                                                                                                                                                                                                               |
| <b>Devices</b>          | Lists devices, device status, storage nodes, libraries, volumes, pools, and related messages. <a href="#">Devices pane</a> on page 87 provides more information.                                                                                                                                                                                                                                                               |
| <b>Operations</b>       | Lists the status of all library and silo operations, including <code>nsrjb</code> operations that are run from the command prompt. Also lists user input, libraries, origin, operation data, operation start time, duration of the operation, progress messages, and error messages.                                                                                                                                           |

**Table 35** Monitoring window panel (continued)

| Window     | Information provided                                                                                                                                                                                                                                                                                                                                                                          |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|            | When displaying Show Details from the <b>Operations</b> window, the length of time that the window is displayed depends on the value that is typed in the <b>Operation Lifespan</b> attribute on the <b>Timers</b> tab of the <b>Properties</b> dialog box for the corresponding library. To access library properties, click <b>Devices</b> in the taskbar. By default, this pane is hidden. |
| <b>Log</b> | Lists messages that are generated by the NetWorker server, including the priority of each message, the time the message was generated, the source of the message, and the category. <a href="#">Log window</a> on page 90 provides more information.                                                                                                                                          |

## Customizing the Monitoring window

This section describes how to customize the **Monitoring** window in the **Administration** interface.

### Customizing tables

You can customize the organization and display of tabular information in the **Monitoring** window.

#### Sorting tables

You can change the display of tabular information that appears in the window. You can sort Table grids by column heading, and then by alphabetic or numeric order within those columns.

1. Drag and drop the column heading to its new position.
2. Click the column heading to sort the items into alphabetic and numeric order. An arrow appears in the column heading to indicate the sort order.

#### Sorting selected rows in a table

Selected rows are sorted to the top of the table. This sorting is particularly useful when you select **Highlight All** from the Find panel to select all rows matching the Find criteria and then moving all selected rows to the top of the table to view the results.

1. From the **Edit** menu, select **Find**, or press **Ctrl + F** to view the **Find** panel.
2. To select the rows, click each row or use the Find criteria.
3. Select **Sort Selected**.

#### Sorting multiple columns in a table

You can select the column that you want to use as the tertiary sort key, the secondary sort key, and the primary sort key.

1. Click the column that you want to use as the last sort key.
2. Click the column that you want to use as the next-to-last sort key, and so on, until you select the primary column.

#### Displaying columns in a table

You can select which columns to display in a table.

1. From the **View** menu, select **Choose Table Columns**.

2. Click a column name to select or clear the column and then click **OK**. You can also select the columns to display by right-clicking a table header and selecting **Add Column** from the drop-down.

## Displaying panes

You can choose to show or hide panes in the **Monitoring** window.

### About this task

Perform the following steps to hide or show a pane in the **Monitoring** window.

### Procedure

1. From the **View** menu, select **Show**. A check mark appears beside the panes that appear in the **Monitoring** window.
2. To hide a pane, select a marked pane.  
A check mark does not appear beside the pane.
3. To show a pane, select an unmarked pane.  
A check mark appears beside the pane.

## Policies/Actions pane

The **Policies/Actions** pane provides you with the ability to review status information about policies and actions.

This pane has two tabs:

- **Policies**—Provides a navigation tree that displays all configured policies on the NetWorker Server. Expand each policy to display the workflows that are associated with each policy. Expand each workflow to display each action that is contained in the workflow.
- **Actions**—Provides a list of all Action resources.

### Policies pane

The **Monitoring** window in the **NetWorker Administration** window enables you to monitor activities for specific policies, workflows, and actions.

The **Policies/Actions** pane at the top of the **Monitoring** window lists the policies on the NetWorker Server by default. Click the + (plus) sign next to a policy in the list to view the workflows in the policy, and the + (plus) sign next to a workflow to view the actions for a workflow.

The **Policies** pane provides the following information for each item (where applicable):

- Overall status

The following table provides details on the status icons that may appear in the **Policies** pane.

**Table 36** Policy status icons

| Icon                                                                                | Status    |
|-------------------------------------------------------------------------------------|-----------|
|  | Never run |
|  | Running   |
|  | Succeeded |

**Table 36** Policy status icons (continued)

| Icon                                                                              | Status                                                |
|-----------------------------------------------------------------------------------|-------------------------------------------------------|
|  | Failed                                                |
|  | Probing                                               |
|  | Interrupted                                           |
|  | Queued                                                |
|  | Cloning                                               |
|  | Consolidating (NetWorker Server 8.2.x and lower only) |

 **Note:** When the schedule for an action is skip, the status of the action appears as Never Run and the status of the Workflow is Succeeded.

- Most recent start time.
- Duration of the most recent run.
- Next scheduled runtime.
- Name of the assigned save set.
- Device on which the save set is stored.
- Backup level.
- Data transfer rate.
- Size of the save set.
- Messages that resulted from an action.

Right-click an action in the **Policies** pane and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

When you sort the items on the **Policies/Actions** pane by using the **Status** column, NetWorker sorts the items in alphabetical order that is based on the label of the icon.

Consider the following when a policy/action is in a probing state:

- A message is sent when the group starts and finishes the probe operation.
- The results of the probe operation (run backup/do not run backup) are also logged.
- Probes do not affect the final status of the group, and the group status does not indicate the results of the probe.
- If probing indicates that a backup should not run, then the group status reverts to its state before the group running.
- Check the results of the probe in the **Log** window to ensure that the probe indicates that the backup can be taken.

### Actions pane

To view a list of all actions, click the **Actions** tab at the bottom of the **Policies** pane. The **Policies** pane becomes the **Actions** pane.

The **Actions** pane provides the following information for each action:



- Overall status

 **Note:** The **Actions** pane displays the same status icons as the **Policies** pane.

- Name
- Assigned policy
- Assigned workflow
- Type
- Date and time of the most recent run
- Duration of the most recent run
- Percent complete, for actions that are in progress
- Next scheduled runtime

Right-click an action in the **Actions** pane and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

## Workflow operations

This section describes how to use the **Monitoring** window to start, stop, and restart workflows.

### Starting, stopping, and restarting policies

The workflows in a policy can run automatically, based on a schedule. You can also manually start, stop, and restart specific workflows by using the the NMC **NetWorker Administration Monitoring** window.

#### About this task

You can restart any failed or canceled workflow. Note, however, that the restart must occur within the restart window that you specified for the workflow. Additionally, for a VMware backup, if you cancel a workflow from **NetWorker Administration** and then want to restart the backup, ensure that you restart the workflow from the **NetWorker Administration** window. If a workflow that was started from **NetWorker Administration** is restarted from the **vSphere Web Client**, the backup fails.

#### Procedure

1. In the **Monitoring** window, select the workflow or actions.
2. Right-click and then select **Start**, **Stop**, or **Restart**.

A confirmation message appears.

 **Note:** You cannot stop, restart, or start individual actions.

3. Click **Yes**.

### Viewing workflow backup details

Perform the following steps to view backup details for workflows.

#### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Policies** in the docking panel, and expand the Policy that you want to monitor.
3. Right-click the workflow, and then select **Show Details**. The **Workflow Summary** window appears.
4. In the **Workflow runs** pane of the **Workflow Summary** window, select the workflow.
5. Click **Show Messages**. In the **Show Messages** window, select one of the following options:

- Get Full Log—To display all messages.
- Print—To print the log.
- Save—To save the log to a local file.
- OK—To close the **Show Messages** window.

6. Click **OK** to close the **Workflow Summary** window.

## Viewing action backup details

Perform the following steps to view backup details for actions.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Actions** in the docking panel.
3. In the **Actions** pane, right-click the action, and then select **Show Details**. The details window for the action appears.
4. Review the information in the **Actions Messages** pane. To display detailed information from the action log file, click **Show Action Logs**, and then select one of the following options:
  - Get Full Log—To display all messages.
  - Print—To print the log.
  - Save—To save the log to a local file.
  - OK—To close the **Show Messages** window.
5. In one of the Actions detail panes, for example, the **Completed successfully** pane, select the action that you want to review.
6. Click **Show Messages**. In the **Show Messages** window, select one of the following options:
  - Get Full Log—To display all messages.
  - Print—To print the log.
  - Save—To save the log to a local file.
  - OK—To close the **Show Messages** window.
7. Click **OK** to close the **Details** window.

## Sessions window

Use the **Sessions** window to view the sessions that are running on a NetWorker server. You can change the view of this window to display these sessions:

The **Sessions** pane below the **Policies/Actions** pane provides details on individual save, recover, clone, and synthetic full sessions by client.

To view all sessions or to limit the list of sessions by the session type, click the tabs at the bottom of the **Sessions** pane. Session types include:

- Save
- Recover
- Clone
- Browse
- Synthetic Full/Rehydrated Sessions
- All

To change the displayed session types go to **View > Show**, and select the type of sessions to display. To display all sessions currently running on the NetWorker Server, regardless of type, select **All Sessions**.

You can stop a session (backup, synthetic full backup, clone, and recovery sessions) from the **Monitoring** window, even if the session was started by running the `nsrworkflow` command.

To stop a session, right-click the workflow in the pane, and select **Stop** from the list box.

## Changing displayed session types

The column headings that are displayed on this window differ depending on the type of sessions you chose to display.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Sessions** in the docking panel.
3. Select **View > Show** and then select the type of sessions to display. To display all sessions currently running on the NetWorker server, regardless of type, select **All Sessions**.

## Stopping a session

You can stop a session (backup, synthetic full backup, clone, and recovery sessions) from the **Monitoring** window, even if the session was started by running `savegrp`.

To stop a session, right-click the session in the window and select **Stop** from the drop-down.

The following table provides a list of actions that can be stopped from NMC.

**Table 37** Sessions that can be stopped from NMC

| Session type                             | Stop from NMC? |
|------------------------------------------|----------------|
| Save by Savegroup                        | Yes            |
| Synthetic Full by Savegroup              | Yes            |
| Clone by Savegroup                       | Yes            |
| Schedule Clone                           | Yes            |
| Manual Save                              | No             |
| Manual Clone via NMC                     | No             |
| Manual Clone via CLI                     | No             |
| Winworker and CLI Recovery               | No             |
| Recovery started from Recover wizard     | Yes            |
| VMware Backup Appliance Save and Recover | No             |

 **NOTICE** Stopping a session from NMC does not affect any other group operations running.

## Alerts pane

The **Alerts** pane displays alerts that are generated by a particular NetWorker server or Data Domain system that has devices that are configured on the NetWorker server. The **Alerts** pane includes priority, category, time, and message information.

An icon represents the priority of the alert. The following table lists and describes each icon.

**Table 38** Alerts window icons

| Icon                                                                              | Label        | Description                                                                                                                           |
|-----------------------------------------------------------------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
|  | Alert        | Error condition detected by the NetWorker server that should be fixed by a qualified operator.                                        |
|  | Critical     | Severe error condition that demands immediate attention.                                                                              |
|  | Emergency    | Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority. |
|  | Information  | Information about the current state of the server. This icon represents the lowest priority.                                          |
|  | Notification | Important information.                                                                                                                |
|  | Waiting      | The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.                                           |
|  | Warning      | A non-fatal error has occurred.                                                                                                       |

When items on the **Alerts** pane are sorted by the **Priority** column, they are sorted in alphabetical order based on the label of the icon.

## Removing alerts

Remove individual alert messages from the **Events** tables by removing them from the **Events** table. To delete a message in the **Events** table, right-click the message, and select **Dismiss**.

### About this task

 **Note:** The alert message remains in the **Log** window in the NetWorker **Administration** program.

## Devices pane

The **Devices** pane allows you to monitor the status of all devices, including NDMP devices. If the NetWorker server uses shared and logical devices, the window is adjusted dynamically to present a set of columns appropriate for the current configuration.

The **Devices** pane provides the following information:

- Status of the operation.
- Name of the device.
- Name of the storage node that contains the device.
- For tape devices, the name of the library that contains the device.
- Name of the volume in the device.
- Name of the pool that is associated with the volume.
- Last message generated for the device.
- Whether the operation requires user input.

For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape.

[Entering user input](#) on page 89 provides instructions on how to deal with a user input notification.

If the current server configuration includes a shared device, a **Shared Device Name** column appears on the **Devices** pane. The name of the shared device appears in the **Shared Device Name** column. If other devices for that configuration are not shared devices, then the **Shared Device Name** column is blank for those devices. Only a single device per hardware ID can be active at any particular moment. The information for inactive shared devices is filtered out, and as a result, only one device per hardware ID is presented on the window at any time.

An icon represents the device status. The following table lists and describes each icon.

**Table 39** Devices status icons

| Icon                                                                                | Label                       | Description                         |
|-------------------------------------------------------------------------------------|-----------------------------|-------------------------------------|
|    | Library device active       | The library device is active.       |
|    | Library device disabled     | The library device is disabled.     |
|    | Library device idle         | The library device is idle.         |
|    | Stand-alone device active   | The stand-alone device is active.   |
|   | Stand-alone device disabled | The stand-alone device is disabled. |
|  | Stand-alone device idle     | The stand-alone device is idle.     |

When you sort items in the **Devices** pane by the **Status** column, NetWorker sorts the devices in alphabetical order based on the label name of the icon.

## Operations window

The **Operations** window displays information about device operations. It provides the following information:

- Status of the operation.
- Name of the library.
- Whether the operation requires user input.  
For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape. [Entering user input](#) on page 89 provides instructions on how to deal with a user input notification.
- The origin, or source, of the operation.  
For example, the interface, nsrjb or the NetWorker server.
- Time the operation started.
- Type of operation.
- Duration of the operation.
- Status messages from the operation.
- Any error messages.

 **NOTICE** Only the last error message of the operation appears in the **Error Messages** column. Move the mouse pointer over the cell containing the last error message to display the entire list of error messages.

The operation status is represented by an icon. The following table lists and describes each of the icons.

**Table 40** Operations window icons

| Icon                                                                              | Label      | Description                                          |
|-----------------------------------------------------------------------------------|------------|------------------------------------------------------|
|  | Failed     | The operation failed.                                |
|  | Queued     | The operation is waiting in the queue to run.        |
|  | Retry      | The operation failed, but may work if you try again. |
|  | Running    | The operation is running.                            |
|  | Successful | The operation completed successfully.                |
|  | User Input | The operation requires user input.                   |

When items on the **Operations** window are sorted by the Status column, they are sorted in alphabetical order based on the label of the icon.

## Viewing operation details

The **Operation Details** dialog box opens, providing information about the completion of the operation. The **Completion Time** displays the time that the operation finished. The time that it took to complete the operation is the difference between the completion and start times of the operation.

### About this task

To save operation details to a file, click **Save** in the **Operation Details** dialog box. When prompted, identify a name and location for the file.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation, then select **Show Details**.

## Stopping an operation

Certain operations can be stopped from the **Operations** window.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation to stop, then select **Stop**.
4. Click **Yes** to confirm the stop.

 **Note:** Operations that were started from a command line program, such as the `nsrjb` command, cannot be stopped from the **Operations** window. To stop these operations, press `Ctrl-C` from the window where the command was started.

## Entering user input

If the system requires user input, select the labeling operation in slow/verbose mode and the **Supply User Input** icon appears.

### Procedure

1. Right-click the operation, then select **Supply Input**.
2. Confirm the requirement to supply input.
  - If **Yes**, and input is supplied, the icon in the **User Input** column disappears.
  -  **Note:** If two users try to respond to the same user input prompt, the input of the first user takes precedence, and the second user receives an error message.
  - If **No**, and input is not supplied, the operation will time out and fail.

## Log window

To view the most recent notification logs, click the **Log** window from the docking panel in the **Monitoring** window. The **Log** window provides the priority, time, source, category, and message for each log.

 **Note:** If a particular log file is no longer available, check the log file on the NetWorker server. The log files are located in `NetWorker_install_path\logs` directory.

An icon represents the priority of the log entry. The following table lists and describes each icon.

**Table 41** Icons in the Log pane

| Icon                                                                                | Label        | Description                                                                                                                           |
|-------------------------------------------------------------------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
|  | Alert        | Error condition that is detected by the NetWorker server that should be fixed by a qualified operator.                                |
|  | Critical     | Severe error condition that demands immediate attention.                                                                              |
|  | Emergency    | Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority. |
|  | Information  | Information about the current state of the server. This icon represents the lowest priority.                                          |
|  | Notification | Important information.                                                                                                                |
|  | Waiting      | The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.                                           |
|  | Warning      | Non-fatal error has occurred.                                                                                                         |

When you sort items on the **Log** pane by using the **Priority** column, NetWorker sorts the icons in alphabetical order based on the name of the label.

## Recover window

The **Recover** window displays information about recover configurations that are created with the NetWorker Management Console (NMC) Recovery wizard.

You can use this window to:

- Start the NMC Recovery wizard to create recover configurations or modify saved recover configurations.
- Identify the status of a recover configuration that is created with the NMC Recovery wizard.
- Start and stop a recover job.

The **Recover** window is divided into five sections:

- **Toolbar**—The toolbar is hidden by default. To display the recovery toolbar, select **View > Show toolbar**
- **Summary**
- **Configured Recovers**
- **Currently Running**

A splitter separates the **Configured Recovers** section from **Currently running** window. You can click and move the splitter to resize these two windows.

## Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. The following table summarizes the function of each toolbar button.

**Table 42** Recovery toolbar options

| Button                                                                              | Function                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|   | Starts the NMC <b>Recover</b> wizard to create recover configurations.                                                                                                                        |
|  | Displays the <b>Properties</b> window for the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                          |
|  | Deletes the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                                                            |
|  | Displays online help for the <b>Recover</b> window.                                                                                                                                           |
|  | Displays the <b>Find</b> window at the bottom of the <b>Recover</b> window. The <b>Find</b> window allows you to perform keyword searches for messages that appear in the <b>Logs</b> window. |
|  | Start the recover operation for a selected saved recover configuration. This option is only available for a recover configuration that has a Never run, or Failed status.                     |
|  | Stop in-progress recover operation that you selected in the <b>Currently Running</b> window.                                                                                                  |

 **Note:** The **Recover** toolbar does not appear by default. To display the **Recover** toolbar, select **View > Show toolbar**.

## Recover Summary

The Recover Summary section displays a high-level overview of recover jobs.

This section includes the following information:



- Total Recovers—The total number of successful recover jobs.
- Since—The number of successful recover jobs since this date.

## Configured Recovers

The **Configured Recovers** window displays a list of saved recover configurations in a tabular format. You can sort the information by column. The **Configured Recovers** table displays the following information for each saved recover configuration:

- Status—The job status of a saved recover configuration.
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB.
- Comment
- OS—The operating system of the source host.
- Recover requestor—The Windows or UNIX account used to create the recover configuration.
- Start Time
- End Time
- Start date

**Table 43** Save recover configuration job status

| Icon                                                                                | Description                                        |
|-------------------------------------------------------------------------------------|----------------------------------------------------|
|  | The last recover attempt failed.                   |
|  | The last recover attempt completed successfully.   |
|  | The recover job has never run.                     |
|  | The recover job is scheduled to run in the future. |
|  | The recover job has expired.                       |

## Currently running

The **Currently Running** window displays a list of in progress recover jobs in a tabular format. You can sort the information by column. The **Currently Running** table displays the following information for each job:

- Status
- Name
- Source client
- Destination client
- Recovery list

- Recover type—For example, file system or BBB
- Volume
- Comment
- Device
- Size
- Total size
- % complete
- Rate (KB/s)
- Start time
- Duration
- Currently running

## Find

The **Find** section appears along the bottom of the **Recover** window, after you select the **Find** button on the **Recover** toolbar. **Find** allows you to search for keywords in the **Configured Recovers** window. The following table summarizes the available find options.

**Table 44** Find options

| Find option          | Description                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Find</b>          | Highlight the first saved recover configuration that contains the specified keyword.                                                         |
| <b>Prev</b>          | Highlight the previous saved recover configuration that contains the specified keyword.                                                      |
| <b>Highlight All</b> | Highlights each saved recover configuration that contains the specified keyword.                                                             |
| <b>Sort Selected</b> | Sorts each highlighted recover configuration in the Configured Recover table so that they appear at the top of the Configured Recover table. |
| <b>Match case</b>    | Make the keyword search case sensitive.                                                                                                      |

## Monitoring checkpoint-enabled backups

You can view detailed information about a checkpoint-enabled backup.

### About this task

#### Procedure

1. From the **Administration** window, select **Monitoring > Groups**.
2. Right-click the group to which the checkpoint enabled client belongs, then select **Show Details**.
3. View the detailed information related to the group backups:
  - If the partial save set is in the work list for the group, the save set appears in the **Waiting to Run** section.
  - If the partial save set is running, the save set appears in the **Currently Running** section.

- If the entire partial save sets sequence of the savegroup is complete, the save set appears in the **Completed Successfully** section.
  - If the entire partial save sets sequence of the savegroup is not complete, the save set appears in the **Failed** section.
-  **NOTICE** If any messages are generated, the Show Messages button is enabled. Click Show Messages to view the messages.

4. Click **OK** to close the **Group Details** window.

## Query the media database for partial save sets

The savegrp completion report does not provide detailed information about partial save sets that might be necessary to perform a recovery.

### Querying partial save sets from the Console

You can view information about the partial save sets by using the NetWorker Console.

#### About this task

-  **NOTICE** If no partial save sets are found that match the query, ensure that the backup of the partial save sets was started within the Save Time period. To change the values for the Save Time attribute, open the Save Set Query tab and select a date and time from the Save Time calendar.

#### Procedure

1. In the **Administration** window, click **Media**. Media-related topics appear in the navigation tree.
2. Select **Save Sets**. The following tabs appear in the **Save Sets** window:
  - Query Save Set
  - Save Set List
3. Select the **Query Save Set** tab, to query:
  - All partial save sets, select **Checkpoint Enabled**.
  - All partial save sets with the same Checkpoint ID, in the **Checkpoint ID** field, type the **Checkpoint ID** of the partial save set on which you want to perform the query.
4. Select the **Save Set List** tab to view the result of the save set query:
  - The **Checkpoint ID** column displays the partial save set **Checkpoint ID** and its Sequence ID. The **Checkpoint ID** is listed first followed by the **Sequence ID**, which is encased within brackets.
  - Sort the **Checkpoint ID** column to view the complete sequence of partial save sets.
  - The **Status** column displays the status of the partial save sets:
    - A Checkpoint browsable status indicates that the save sets can be browsed for recover.
    - A Checkpoint aborted status indicates that the backup of the partial save set was stopped or aborted. A save set recover is used to recover the partial save set. Consider the following:
      - When a checkpoint-enabled backup completes successfully, the status of the last partial save set is Checkpoint browsable.
      - When a checkpoint-enabled backup completes successfully, on the first backup attempt, the save set status is Checkpoint browsable. Only one Sequence id is

associated with the Checkpoint ID. The Sequence id is 1. If the Sequence id is 2, the first partial save set in the checkpoint-enabled backup is missing.

## Querying partial save sets by using the `mminfo` command

By default, the `mminfo` command output only displays the browsable save sets. The first and intermediate partial save sets are not displayed. Only complete checkpoint-enabled save sets or final partial save sets are displayed.

Use the `mminfo` command with specific queries to display more information about checkpoint-enabled save sets.

The following table lists the new media attributes support the Checkpoint Restart feature.

**Table 45** New Checkpoint Restart media attributes

| Media attribute    | Description                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------|
| checkpoint_id      | Displays the checkpoint restart id of the partial save set in the <code>chkpt_id</code> column. |
| checkpoint_seq     | Displays the partial save set sequence id in the <code>chkpt_seq</code> column.                 |
| checkpoint-restart | This flag attribute is used to only display checkpoint restart enabled save sets.               |

Also, several media sumflags are used with the Checkpoint Restart feature:

- `k`—Indicates this is a checkpoint enabled save set.
- `a`—The first and all intermediate partial save sets of a checkpoint sequence has aborted status.
- `b`—The last partial or complete save set of a checkpoint sequence is marked browsable.

## Displaying checkpoint enabled save sets

Display all checkpoint enabled save sets by using the following `mminfo` command:

```
mminfo -q 'checkpoint-restart' -r 'client,nsavetime,ssid(11),
sumflags(3),name,checkpoint_id,checkpoint_seq'
```

**Table 46** Checkpoint enabled save sets

| client | save time  | ssid           | ssflags | filename | chkpt_id   | chkpt_seq |
|--------|------------|----------------|---------|----------|------------|-----------|
| plapew | 1251910303 | 420470031<br>9 | cak     | /space   | 1251910303 | 1         |
| plapew | 1251910327 | 4187923127     | cbk     | /space   | 1251910303 | 2         |
| plapew | 1251910710 | 408726021<br>4 | cak     | /space   | 1251910710 | 1         |
| plapew | 1251910725 | 407048301<br>3 | cbk     | /space   | 1251910710 | 2         |

## Displaying all partial save sets for the checkpoint id

Display all partial savesets for the checkpoint id by using the following `mminfo` command:

```
mminfo -q "checkpoint_id=1251910303"
```

**Table 47** Partial save sets for the checkpoint id

| volume     | client | date     | size   | level | name   |
|------------|--------|----------|--------|-------|--------|
| plapew.001 | plapew | 08/02/15 | 17 MB  | full  | /space |
| plapew.001 | plapew | 08/02/15 | 799 MB | full  | /space |

## Reporting NDMP Data

The NetWorker software reports information about NDMP clients, data, and volumes in two ways:

- The NetWorker Management Console (NMC) reporting feature—Reports NDMP data in the same manner as non-NDMP data. The *NetWorker Administration Guide* provides more information.
- The `mminfo` command. Use the `mminfo` program to query the media database for NDMP volume or save set information.

### Querying the NDMP volumes by backup type with the `mminfo` command

You can query save sets by backup format (NDMP or DSA) to display volume information.

For example:

- To query NDMP volumes, type `mminfo -q ndmp`. Output similar to the following appears:

```
volume client date size level name
005D0000 simlcifs1 6/22/2011 1036 MB full /fs1
005D0001 simlcifs1 6/22/2011 173 MB full /fs1
005D0001 simlcifs1 6/22/2011 862 MB full /fs1
005D0002 simlcifs1 6/22/2011 348 MB full /fs1
```

- To query NDMP -DSA volumes, type `mminfo -q dsa`. Output similar to the following appears:

```
volume client date size level name
NDMP.001 10.8.67.219 12/13/2011 644 MB full /vol/vol0
NDMP.001 10.8.67.219 12/13/2011 402 MB full /vol/vol1
NDMP.001 10.8.67.219 12/13/2011 402 MB full /vol/vol1
NDMP.001 10.8.67.219 12/13/2011 36 MB full /vol/vol2
```

### Querying the NDMP save sets with the `mminfo` command

To determine which save sets are Network Data Management Protocol (NDMP) save sets and the status of an NDMP save set in the media database, query the media database. NDMP save set status information is important when performing NDMP recoveries:

- To perform a browsable NDMP recover, the `ssflags (f1)` field for an NDMP save set must contain a (b). The b value denotes a browsable save set.
- To perform a save set recover from the NetWorker User program, the `ssflags (f1)` field for an NDMP save set must contain either (r) or (b).
- An NDMP save set contains an N attribute in the `ssflags (f1)` field.
- An NDMP-DSA save set contains an s attribute in the `ssflags (f1)` field.

In the following example, the NDMP save set status is recoverable (r). To recover the data, you can only perform a save set recovery from a command line.

```
mminfo -av
```

```
volume type client date time size ssid fl lvl name
vol1 dlt clnt 6/22/2011 3:15:12 1036MB 3842140553 hrN full /fs1
```

In the following example, the NDMP-DSA save set status is browsable (b). Recover the data by using the NetWorker User program, or from the command line. A browsable NDMP-DSA save set supports browsable and save set recoveries.

```
mminfo -av
```

```
volume type client date time size ssid fl lvl name
vol1 dlt clnt 6/22/2011 3:15:12 36MB 4259813785 cbNs full /fs1
```

## Performing NDMP recoveries

NetWorker uses the `nsrndmp_recover` program to coordinate recover operations between the NetWorker software and the Network Data Management Protocol (NDMP) client. The `nsrndmp_recover` program does not move data to the NDMP client. When the `nsrndmp_recover` program identifies an NDMP-DSA save set, `nsrndmp_recover` automatically runs the `nsrdsa_recover` program on the same host that runs the `nsrndmp_recover` command.

To recover NDMP data, you can run the `nsrndmp_recover` program from a command prompt, or use one of following programs, which automatically starts `nsrndmp_recover`:

- `recover`—The command line program on Windows and UNIX.
- `winworkr`—The NetWorker User GUI on Windows.
- The NMC Recovery wizard.

During the recovery process, the `nsrndmp_recover` program passes `nlist` information to the NDMP client. There are three methods to recover NDMP backups:

- Index-based file-by-file recover—The `nlist` includes file offset and ACL information. When you recover many files, the recover process uses a significant amount of system resources on both the NetWorker server and the NDMP client to build and process the `nlist` information.
- Full save set recovery—The `nlist` only includes the path to the recovery directory, down to and including the mount point. When you recover many files, the recover process uses less system resources than an index-based NDMP recover to build and process the `nlist` information.
- NDMP directory restore—A partial save set recovery of a single file or single directory.

For example, when the NetWorker software writes NDMP data to a remote storage node, start the `recover` program on the NetWorker storage node to prevent the data from traversing the network.

**Note:** When you start the `recover` program on the NetWorker server, the data flows from the storage node to the NetWorker server and from the NetWorker server to the NDMP client, over the network.

## NDMP recovery requirements for Isilon

The following list summarizes the requirements:

### scanner

You cannot use the `scanner` command with the `-i`, `-f` and `-r` options on an NDMP volume. You cannot use the `scanner` command on a volume that contains NDMP and non-NDMP save sets when you load the volume in an NDMP device. The *Scanner command usage*

technical note provides more information about using the `scanner` command with NDMP data.

### Cross platform recoveries

You can recover NDMP data to different NDMP client however, you cannot perform a cross platform recover. Recover NDMP data to an NDMP client that is the same brand, a compatible model, and the same operating system as the original NDMP client.

### Devices

Recover Direct-NDMP and Three-party backups performed to an NDMP device from an NDMP device. To improve recover performance from an NDMP tape device, configure the tape device to support variable length records. Recover NDMP-DSA backups from a non-NDMP device.

### Localized environments

When recovering data in a localized NDMP environment, the Index Recover status window shows the process in English and not the localized language.

### NDMP-DSA

For better recovery performance, start the recover process on the NetWorker host where the backup volume resides.

### Immediate recoveries

Run the `nsrndmp_recover` program on the storage node with the locally attached backup device to perform an immediate recovery of NDMP-DSA data.

## DAR and DDAR

By default, the Network Data Management Protocol (NDMP) recover process reads an entire tape from start to finish. The recover process extracts the data as it encounters the data on the tape. For large backup images, recovery is slow.

The Direct Access Recovery (DAR) and Directory DAR (DDAR) recovery process:

- Provides the ability to recover a file or directory from the exact location on a tape.
- DDAR only passes the directory path to the NAS filer. DAR passes the paths of each file individually.
- Reduces the size of the nlist information that the recover process stores in memory. During the recover process, the NAS filer (DDAR) assumes that the directory path includes all cataloged files and directories. However, DAR mentions each file that it wants recovered.
- Does not sequentially read the file or record numbers on the tape to locate the data, which reduces the amount of time that you require to recover specific files from a backup.

**Note:** [Creating and configuring the NDMP client resource](#) on page 67 describes how to configure the DAR and DDAR Application Information attributes for NDMP clients.

**Note:** DDAR is not applicable when multiple levels of backups are present and for isilon multi-stream backups. DDAR is applicable only for FULL (level 0 ) backups.

### When not to use DAR or DDAR

DAR and DDAR recoveries send multiple pathnames across the network to the NDMP Data Server and, in three-party configurations, to the NetWorker server. The recover process stores the pathnames in memory on the NDMP Data Server. Recoveries of a large amount of data from a large save set can negatively affect the network and the NDMP Data Server resources.

Do not use DAR and DDAR to recover the following objects:

- Several thousands of files in a single index-based recover operation.

- A specific directory structure containing several thousand or millions of files.

To perform a non-DAR-based recovery of a save set when you set the *DIRECT=y* at the time of backup, first define the *NSR\_NDMP\_RECOVER\_NO\_DAR=y* variable in the Application Information attribute of the NDMP client.

## Recovering data from partial save sets

If there is a complete sequence of partial save sets that span the original save set, then you can browse to and recover individual files and directories. If the sequence of partial save sets is incomplete and does not make up the original save set, then you must perform a save set recovery to recover the data from the partial save set.

To recover data from partial save sets that span the original save sets, perform a query for all partial save sets, and then use either the NetWorker User program on Windows or the `recover` program on UNIX to restore the data.

The steps to recover data from a single partial save set are the same as save set recovery from a complete save set. The partial save set contains only files that were successfully backed up. You cannot browse partial save sets.

When you perform a save set recovery of a partial NDMP save set, the recovery process recovers all partial save sets in the checkpoint sequence. You cannot recover data in a partial save set separately from other partial save sets in the checkpoint sequence.

Use the `nsrinfo` command to display the contents of a partial save set.

## Recover window

The **Recover** window displays information about recover configurations that are created with the NetWorker Management Console (NMC) Recovery wizard.

You can use this window to:

- Start the NMC Recovery wizard to create recover configurations or modify saved recover configurations.
- Identify the status of a recover configuration that is created with the NMC Recovery wizard.
- Start and stop a recover job.

The **Recover** window is divided into five sections:

- **Toolbar**—The toolbar is hidden by default. To display the recovery toolbar, select **View > Show toolbar**
- **Summary**
- **Configured Recovers**
- **Currently Running**

A splitter separates the **Configured Recovers** section from **Currently running** window. You can click and move the splitter to resize these two windows.



## Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. The following table summarizes the function of each toolbar button.

**Table 48** Recovery toolbar options

| Button                                                                              | Function                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|    | Starts the NMC <b>Recover</b> wizard to create recover configurations.                                                                                                                        |
|    | Displays the <b>Properties</b> window for the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                          |
|    | Deletes the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                                                            |
|    | Displays online help for the <b>Recover</b> window.                                                                                                                                           |
|    | Displays the <b>Find</b> window at the bottom of the <b>Recover</b> window. The <b>Find</b> window allows you to perform keyword searches for messages that appear in the <b>Logs</b> window. |
|   | Start the recover operation for a selected saved recover configuration. This option is only available for a recover configuration that has a Never run, or Failed status.                     |
|  | Stop in-progress recover operation that you selected in the <b>Currently Running</b> window.                                                                                                  |

 **Note:** The **Recover** toolbar does not appear by default. To display the **Recover** toolbar, select **View > Show toolbar**.

## Recover Summary

The Recover Summary section displays a high-level overview of recover jobs.

This section includes the following information:

- Total Recovers—The total number of successful recover jobs.
- Since—The number of successful recover jobs since this date.

## Configured Recovers

The **Configured Recovers** window displays a list of saved recover configurations in a tabular format. You can sort the information by column. The **Configured Recovers** table displays the following information for each saved recover configuration:

- Status—The job status of a saved recover configuration.
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB.

- Comment
- OS—The operating system of the source host.
- Recover requestor—The Windows or UNIX account used to create the recover configuration.
- Start Time
- End Time
- Start date

**Table 49** Save recover configuration job status

| Icon                                                                              | Description                                        |
|-----------------------------------------------------------------------------------|----------------------------------------------------|
|  | The last recover attempt failed.                   |
|  | The last recover attempt completed successfully.   |
|  | The recover job has never run.                     |
|  | The recover job is scheduled to run in the future. |
|  | The recover job has expired.                       |

## Currently running

The **Currently Running** window displays a list of in progress recover jobs in a tabular format. You can sort the information by column. The **Currently Running** table displays the following information for each job:

- Status
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB
- Volume
- Comment
- Device
- Size
- Total size
- % complete
- Rate (KB/s)
- Start time
- Duration
- Currently running

## Find

The **Find** section appears along the bottom of the **Recover** window, after you select the **Find** button on the **Recover** toolbar. **Find** allows you to search for keywords in the **Configured Recovers** window. The following table summarizes the available find options.

**Table 50** Find options

| Find option          | Description                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Find</b>          | Highlight the first saved recover configuration that contains the specified keyword.                                                         |
| <b>Prev</b>          | Highlight the previous saved recover configuration that contains the specified keyword.                                                      |
| <b>Highlight All</b> | Highlights each saved recover configuration that contains the specified keyword.                                                             |
| <b>Sort Selected</b> | Sorts each highlighted recover configuration in the Configured Recover table so that they appear at the top of the Configured Recover table. |
| <b>Match case</b>    | Make the keyword search case sensitive.                                                                                                      |

## Performing an NDMP index-based file-by-file data recovery

Perform an NDMP index based file-by-file recover in the same manner as a non-NDMP data recover. You can restore the data to the original NDMP client or perform a directed recovery to a different NDMP client.

Before you perform an index-based file-by-file recover, review the following information:

- Set the `HIST=y` in the Application Information attribute of the NDMP client at the time of the backup.
- The NDMP save set must be browsable. You cannot perform a browsable recover of a recoverable or recyclable save set. [Reporting NDMP Data](#) on page 93 describes how to determine the status of an NDMP save set.
- Do not use an index-based recovery to recover a large numbers of files or directories. For better recovery performance, use a save set recover. [Performing a Full or Directory Restore of NDMP data by using a save set recovery](#) on page 103 provides more information.
- To perform an index-based file-by-file recover:
  - Use the NetWorker User program on a Windows host. [Performing an NDMP index-based file-by-file recover using the NetWorker User program](#) on page 100 provides detailed information.
  - Use the `recover` program. [Performing an NDMP index-based file-by-file recover from a command prompt](#) on page 102 provides detailed information.

### Performing an NDMP index-based file-by-file recover using the NetWorker User program

On Windows, to recover data to the original NDMP client or to a different NDMP client, perform the following steps.

#### Procedure

1. Open the NetWorker User program and connect to the NetWorker server.

**NOTICE** If you receive the error:

No file indexes were found for client `client_name` on server `server_name`

Try connecting to a different NetWorker server and you selected the correct NetWorker server, then ensure that you selected a browsable save set. Alternatively, perform a save set recover.

2. Select **Recover** to open the **Source Client** window.
3. Select source NDMP client and click **OK**. The local client is the default selection.
4. Select the destination client for the recovered data and click **OK**. If the destination client is not the source client, ensure the NAS filer is the same brand, a compatible model and the same operating system as the source NDMP client.
5. (Optional) Recover the data from an earlier backup time. The **Recover** window appears with the latest version of the backup files. To recover data from an earlier backup, change the date and time of backup using one of the following methods:
  - a. Change the browse time for all files in the recover window:
    - From the **View** menu, select **Change Browse Time**.
    - In the **Change Browse Time** window, select a new day within the calendar. Select **Previous Month** or **Next Month** to change from the current month.
    - In the **Time** field, change the time of day by typing an hour, a minute, and the letter a for A.M. or p for P.M. Use the 12-hour format.
    - Click **OK**.
  - b. View all versions of the selected file system object:
    - Highlight the file or directory for review.
    - From the **View** menu select **Versions**.
    - Once you locate the version to recover, change the browse time. To change the browse time, highlight the volume, directory, or file and click **Change Browse Time**. The **Version** window closes and the **Recover** window reflects the new browse time.
6. (Optional) Search for the files. To search for and recover the most recently backed-up version of a file or directory:
  - a. From the **File** menu, select **Find**.
  - b. Type the name of the file or directory. Use wildcards to expand the search; without wildcards, partial filenames do not provide any results.
7. Mark the data to recover. To select file system objects to recover:
  - a. In the left pane of the **Recover** window, click the appropriate directory folder.
  - b. Mark each directory or file to recover by selecting the checkbox next to each directory or file.
8. (Optional) Relocate the data to a different location. By default, the recover process recovers the selected files to the original location.

**Note:** The NDMP protocol does not support name conflict resolutions. NetWorker will always overwrite existing files that have the same name as the recovered file. It is recommended that you recover the NDMP data to a different location, to avoid data loss.

To relocate the files to a different location:

- a. Select **Recover Options** from the **Options** menu.

NDMP recoveries do not support the following options:

- Rename recovered file
- Discard recovered file
- Prompt for every file conflict

NDMP recoveries will always overwrite existing files. It is recommended that you relocate the NDMP data to a different location, to avoid data loss.

- b. In the **Relocate Recovered Data To** field, type the full path name of the target directory, click **OK**.

The target directory is a literal string and must match the path as seen by the NAS filer in its native OS, exactly. Otherwise, the recover process uses the original location and overwrites existing files with the same name.

9. (Optional) To view the volumes required to recover the marked file system objects, from the **View** menu, select **Required Volumes**.
10. Click **Start** to begin the recovery. If any required volume is not available to the NetWorker server, a volume status warning appears.

When this warning appears:

- a. Click **No**.
- b. From the **View** menu, select **Required Volumes**.
- c. Ensure that the NetWorker software can mount each listed volume into an available device.
- d. Attempt the recover operation again.

The NetWorker server takes a few moments to recover the files, depending on file size, network traffic, server load, and tape positioning. During this time, messages appear so that you can monitor the progress of the recovery.

When the recovery completes successfully, a message similar to the following appears:

```
Received 1 file(S) from NSR server
server Recover completion time: Tue Jan 21 08:33:04 2009
```

## Performing an NDMP index-based file-by-file recover from a command prompt

This section applies to command line recoveries from a Windows and UNIX client.

### About this task

To avoid using the Windows version of `recover.exe` on Windows operating systems, perform one of the following actions:

- Specify the full path to the recover program. For example: `C:\Program Files\EMC NetWorker\nsr\bin\recover.exe`
- Ensure that the `$PATH` environment variable contains the `NetWorker_install_path\bin` directory before `%SystemRoot%\System32`

To recover Network Data Management Protocol (NDMP) data from a command prompt on a UNIX or Windows NetWorker host, perform the following steps.

## Procedure

1. From the command prompt, type:

```
recover -s NetWorker_servername -c client_name
```

where:

- *-s NetWorker\_servername* specifies a particular NetWorker server on the network to use when recovering data.

When you do not use the *-s* option, the `recover` program tries to connect to the first computer listed in the servers file. When the servers file does not contain any servers, or lists more than one server, the **Change Server** window appears, and you can select the server.

- *-c client\_name* specifies the source NDMP client.

2. When prompted, type the directory to browse, for example:

```
cd /mydirectory
```

3. Use the `add` command to add the required files or folders to the recover list. The *NetWorker Command Reference Guide* provides a complete list of options for the `recover` command.
4. When restoring NDMP data, it is recommended that you relocate the NDMP data to a different location.

**Note:** The NDMP protocol does not support name conflict resolutions. NetWorker will always overwrite existing files that have the same name as the recovered file. It is recommended that you recover the NDMP data to a different location, to avoid data loss.

- To relocate the data to a different directory, type:

```
relocate destination_directory_name
```

The target pathname for *destination\_directory\_name* is a literal string and must match the path as seen by the NAS filer in its native OS, exactly. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

- To recover the data to a different host, type:

```
relocate target_hostname::/mount_point
```

Data ONTAP may require you to add a backslash (\) after the mount point. For example, *target\_hostname::\mount\_point\*.

5. After you add all of the required files, type:

```
recover
```

## Performing a Full or Directory Restore of NDMP data by using a save set recovery

You perform a Network Data Management Protocol (NDMP) save set recover in the same manner as a non-NDMP save set recovery. You can recover data to the original NDMP client or perform a directed recovery of the data to a different NDMP client of the same platform.

Before you perform a full save set recover, review the following information:

- Use a full save set recovery to recover all files and folders in an NDMP data save set, or to recover an entire directory within an NDMP save set. You cannot use the NetWorker User program to perform an NDMP Directory Restore.
- To use the NetWorker User program on Windows, a client file index entry for the save set must exist. When the index entry for the save set does not exist, the recover fails with an `index not found` error. When the client file index entries do not exist for the save set, use the `nsrndmp_recover` program with the `-v off` option.
- You cannot perform a save set recover from the NetWorker User program when the save set status is eligible for recycling (E). The recover process requires a recoverable (r) or browsable (b) save set status. The *NetWorker Administration Guide* provides information on how to change the status of a save set. A save set recover reads the entire tape, from beginning to end, to find and recover the requested files. The recovery process completes when the recover operations reads all required tapes in their entirety.
- As each file recovers, the file name appears on the target share but the file size is 0 KB. The actual file size update occurs after the recovery completes.

## Performing an NDMP save set recover by using the NetWorker User in Windows

### About this task

**NOTICE** When the recover operations fails with the error:

```
Failed to propagate handle <number> to child process: Access is denied
```

The save set is not in the client file index of the NDMP client. Perform a save set recover from a command prompt. [Performing an NDMP save set recovery from the command prompt](#) on page 105 provides more information.

### Procedure

1. Start the NetWorker User program.
2. On the **Change Server** window, select the NetWorker server and click **OK**.
3. Select **Options > Recover Save Sets**.
4. On the **Source Client** window, select the appropriate NDMP client, and then click **OK**.
5. On the **Save Sets** window, select the name of the save set.
6. Select the version of the save set, if there are multiple versions. You can also select the cloned version of a save set, if applicable.
7. To recover specific files and directories instead of the entire save set:
  - a. Click **Files**.
  - b. Specify the files and directories, one per line.
  - c. Click **OK**.

**NOTICE** Do not use this method to mark tens of thousands of files. Instead, perform an NDMP Directory Restore. Marking many files and directories generates a large nlist and requires intensive resources on both the NetWorker server and the NAS filer.

8. Click **Recover Options**.

An NDMP data recovery does not support the following options:

- Rename recovered file

- Discard recovered file
- Prompt for every file conflict

 **NOTICE** It is recommended that you relocate the NDMP data to a different location. NDMP recoveries always overwrite existing files.

9. To recover the data to a pathname that is different from the original backup location, in the **Relocate Recovered Data To** field, type the full pathname of the destination directory, then click **Ok**.

For NDMP data recoveries, the target pathname is a literal string and must exactly match the path as seen by the native OS on the NAS filer. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

10. To recover the data to a different NDMP client, specify the name of the client to receive the NDMP data in the **Destination Client** field.
11. To view the volumes that are required to perform the recover, select **View > Required Volumes**
12. Click **OK** to begin the recovery. The recovery status appears in the **Recover Status** window.

## Performing an NDMP save set recovery from the command prompt

To perform a save set recovery to the original Network Data Management Protocol (NDMP) client or to a different NDMP client, use the `nsrndmp_recover` command.

### About this task

For example:

```
nsrndmp_recover -s NetWorker_server -c source_ndmp_client -S ssid/cloneid -v off
-m target_ndmp_client::/target_path /source_path
```

where:

- `source_ndmp_client` is the hostname of the source NDMP client.
- `target_ndmp_client` is the hostname of the destination NDMP client.
- `/source_path` is the original location of the data.
- `/target_path` is the location to recover the data.

 **NOTICE** It is recommended that you relocate the NDMP data to a different location. NDMP recoveries always overwrite existing files. The `/target_path` is a literal string and must exactly match the path as seen by the native OS on the NAS filer. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

- `-v off` allows you to restore data when client file index of the NDMP client does not contain information about the NDMP save set.

In the following examples, the NetWorker server is mars and the backup client is venus:

- To recover a mount point `/mnt` from a backup of NDMP host venus to a directory `/newmnt` on NDMP host jupiter, type:

```
nsrndmp_recover -s mars -c venus -S 123456789 -v off -m jupiter::/newmnt
```

- To recover a mount point `/mnt` from a backup of NDMP host venus to NDMP host pluto, type:



```
nsrndmp_recover -s mars -c venus -R pluto -S 123456789 -v off -m /mnt
```

Data ONTAP may require that you to add a slash (/) after the mount point. For example, *target\_hostname:/mount\_point/*.

## Troubleshooting NDMP recover

This section provides a list of the possible causes and the possible resolutions for NDMP recovery issues.

### RESTORE: could not create path *pathname*

This error message appears when restoring NetApp data. This error, when encountered, appears in the `daemon.raw` file of the NetWorker server and the recovery output.

To resolve this issue:

- Ensure that you specify a source and a target path during the recover that exists on the target filer.
- If you set the *UTF8=Y* application information variable during an NDMP client backup and the backup contains path names with non-ASCII characters, then perform a save set recover. Index-based recoveries will fail with this error message.

### These files were not restored (Restore failed with error, or file/directory specified but not found in backup)

This error message appears in the `daemon.raw` file of the NetWorker server and the in the recovery output.

To resolve this issue:

- Ensure that the file or directory specified during the recover, exists in the save set.
- Ensure that the pathname specified to relocate the data exists on the destination filer. For NDMP data recoveries, the target pathname is a literal string and must exactly match the path as seen by the native OS on the NAS filer.



# CHAPTER 4

## NetApp

This chapter includes the following topics:

|                                                                                             |     |
|---------------------------------------------------------------------------------------------|-----|
| • <a href="#">Choosing a device type</a> .....                                              | 188 |
| • <a href="#">Configuring devices for NDMP operations</a> .....                             | 188 |
| • <a href="#">Configure NetWorker for NDMP backup and clone operations</a> .....            | 199 |
| • <a href="#">Monitoring NetWorker Server activities in the Administration window</a> ..... | 238 |
| • <a href="#">Reporting NDMP Data</a> .....                                                 | 255 |
| • <a href="#">Performing NDMP recoveries</a> .....                                          | 256 |

## Choosing a device type

Network Data Management Protocol (NDMP) backups can be written to either an NDMP device, or if using NDMP-DSA, to a non-NDMP device.

Perform either of the following tasks:

- Configure devices for NDMP operations.
- Configure non-NDMP devices. If you are using NDMP-DSA, refer to the *NetWorker Administration Guide* for device configuration.

For a description of each configuration, refer to [Configurations in a NetWorker NDMP environment](#) on page 16.



**Note:**

PowerStore does not support local or direct NDMP backup.

## Configuring devices for NDMP operations

Review this section for information about how to configure the NetWorker environment for Network Data Management Protocol (NDMP) data operations.

The *NetWorker Hardware Compatibility Guide* on the Support website provides a list of NDMP devices that the NetWorker software supports.

### NDMP device limitations

Review these limitations before you configure Network Data Management Protocol (NDMP) devices:

- The timeout of the NetWorker server `nsrmmmd` resource attribute does not apply to NDMP devices, but it does apply to storage nodes devices.
- You cannot use the `jbexercise` utility with an NDMP autochanger.
- You cannot configure NDMP devices on a dedicated storage node.
- You must use a non-rewind device handle for the NDMP media device handle.
- You cannot configure advanced file type devices and file type devices as NDMP devices.
- You cannot configure an NDMP autochanger when the NDMP protocol is earlier than version 3. You must determine the NDMP device handles, then use the `jbconfig` command to configure the autochanger.

### Determining NDMP device pathnames

To configure an NDMP stand-alone device or an NDMP jukebox, you must first determine the path names of the media devices. If the NAS filer does not support the NDMP\_CONFIG interface or uses NDMP version 3, you must also determine the library device handle.

To determine the NDMP device path names and the library handle, use the `inquire` command or vendor-specific commands.

### Determining the NDMP device path names using the `inquire` command

#### About this task

Use the `inquire` command to determine the path names and library handle.

## Procedure

1. From a command prompt on the NetWorker server, type:

```
inquire -N NAS_hostname -T
```

2. When prompted, specify the NAS username and password.

**NOTICE** Use the `inquire` command with caution. When you run `inquire`, the command sends the SCSI `inquiry` command to all devices that are detected on the SCSI bus. If you use the `inquire` command during normal operations, unforeseen errors can occur, which might result in data loss.

## Determining the NDMP device path name for NetApp

Before you configure an NDMP autochanger you must determine the device path names of NDMP devices and the robotic arm.

### Procedure

- Log in to the appliance as root or as a Windows Administrator and type:

```
sysconfig -t
```

The host responds with a list of media device names, for example:

```
Tape drive (DS_300B:3.126L10) Hewlett-Packard LTO-4 nrst0l - no
rewind device, format is: LTO-2(ro)/3 2/400GB nrst0m - no rewind
device, format is: LTO-2(ro)/3 4/800GB cmp nrst0h - no rewind device,
format is: LTO-4 800GB nrst0a - no rewind device, format is: LTO-4
1600GB cmp
```

where:

- (DS\_300B:3.126L10) indicates the switch (DS\_3--6 ), the port number (3) and the LUN number( 10). This information must match the output in the `sysconfig -v` command.
- nrst0l is the media device name.  
When the filer uses NDMP v2, or does not support the NDMP\_CONFIG interface, to determine the autochanger handle, type:

```
sysconfig -m
```

The host responds with the devices on the host, for example:

```
Medium changer (DS_300B:3.126L9) ADIC Scalar i2000 mc5 - medium
changer device
```

where mc5 is the autochanger handle.

## Dynamic drive sharing

Dynamic Drive Sharing (DDS) is a feature that provides NetWorker software with the ability to recognize shared physical tape drives. DDS enables NetWorker software to perform the following operations:

- Skip the shared tape drives that are in use.
- Route the backups or recoveries to other available shared tape drives.

## Introduction to DDS

DDS controls application requests for tape media and allows the NetWorker server and all storage nodes to access and share all attached devices.

A system administrator can configure DDS by setting a sharing policy for devices that are accessible from multiple storage nodes.

There are two terms that are central to the use of DDS are drive and device. Within the context of DDS, these terms are defined as follows:

- Drive—The physical backup object, such as a tape drive, disk, or file.
- Device—The access path to the physical drive.

 **Note:** NetWorker only supports DDS in a storage area network (SAN) Fibre Channel environment and not in a direct-connect SCSI environment.

## Benefits of DDS


Enabling DDS on a NetWorker system provides these benefits:

- Reduces storage costs—You can share a single tape drive among several storage nodes. In fact, since NetWorker software uses the same open tape format for UNIX, Windows, NetWare and Linux, you can share the same tape between different platforms (assuming that respective save sets belong to the same pool).
- Reduces LAN traffic—You can configure clients as SAN storage nodes that can send save sets over the SAN to shared drives.
- Provides fault tolerance—Within a SAN environment, you can configure hardware to eliminate a single point of failure.
- Provides configuration over a greater distance—You can configure a system over a greater distance than with SCSI connections.

## DDS configuration overview

The following figure illustrates the DDS process and potential device sharing configurations. This basic configuration consists of a server, two storage nodes, and a library with two tape drives.

**Figure 17** Dynamic Drive Sharing



In this figure:

- Storage nodes sn\_1 and sn\_2 are attached to the library.
- Each storage node, on its own, has access to drive\_1 and drive\_2.
- With DDS enabled, both storage nodes have access to both drives and can recognize when a shared drive is in use.

This configuration requires two DDS licenses, one for each drive.

**Note:** Ensure that all applicable devices can be seen from each storage node by running the `inquire -l` command locally on each storage node.

## DDS block-size compatibility between UNIX and Windows

With DDS enabled, drives can be shared between storage nodes on different platforms, such as UNIX and Microsoft Windows. For NetWorker software operations (such as backups and recoveries) to take place successfully, ensure that the block size is compatible between different platforms or hardware.

To ensure compatibility, make sure one of the following conditions is met:

- The various storage nodes sharing a drive support the same block sizes.
- When a tape is labeled on a drive, it is labeled with the block size defined on the storage nodes.

## Block-size incompatibility between UNIX and Windows

Incompatible block-size settings between UNIX and Microsoft Windows storage nodes could result in any of these error scenarios:

- A backup taken on a UNIX node might not be recoverable on a Microsoft Windows node if the Windows node does not support large block sizes.
- A UNIX process labels and saves data to a tape and leaves the tape mounted. A Microsoft Windows process subsequently attempts to verify the label on this tape and fails because the label verification is done by reading a header from the data portion.
- A tape on a UNIX node is labeled with a large block size. The backup is started on a Microsoft Windows node and the Windows node attempts to write the backup by using the default block size. Internally, the backup on Windows is written by breaking down the big buffer of data into smaller segments of writable block sizes.  
Attempting to recover a specific file on Windows in this situation fails due to positioning errors on the tape. The data is still recoverable from the Windows side, since the NetWorker software will switch from using file and block positioning to reading the tape from the beginning to reach the correct position. The data might not, however, be recoverable from the UNIX side.

## Unintended Access to DDS device prevention

The Reserve/Release attribute has been added to the Device resource for tape devices to support Reserve/Release, including the Persistent Reserve commands.

Reserve/Release is a mechanism that uses SCSI commands to attempt to prevent unintended access to tape drives that are connected by using a shared-access technology, such as Fibre Channel, iSCSI, or SCSI multiplexers. It is a “cooperative” and host-based mechanism, which means that all applications should respect the reservations and not purposely break them. Access is granted based on the host system that reserved the device. Other applications that run on that host cannot be prevented from accessing a reserved device.

Reserve/Release cannot prevent a malicious or badly behaved application from accessing a reserved device. It also cannot prevent all problems caused by hardware issues (such as SCSI resets or FC LIPs) from interrupting data access.

The basic sequence requires that a host reserve a tape drive (using specific SCSI commands) before attempting to access the tape drive. If this “reservation” succeeds, then the host can use the drive. If the reservation fails (usually because the device is reserved by someone else), then the host attempting the reservation should not attempt to use the drive. When a host has finished using a reserved drive, that host must release the drive by using the appropriate SCSI commands.

The reservation is maintained by the drive itself. With older (called “Simple” in NetWorker software) Reserve/Release, the reservation is based on the SCSI ID of the system that issued the reserve command. For tape drives connected to Fibre Channel (FC) using FC-SCSI bridges, the mapping between FC host and reservation is done inside the bridge, since the initiator on the SCSI side is always the bridge itself, regardless which host actually issued the reserve command.

For Persistent Reserve, the reservation is associated with a 64-bit “key” that is registered by the host. Several keys can be registered with a given drive at any given time, but only one may hold the active reservation. NetWorker software uses the “exclusive” reservation method for Persistent Reserve. Only the host that holds the active reservation is allowed to access the drive.

The Reserve/Release attribute does not support file type or advanced file type devices.

The settings that relate to Reserve/Release and Persistent Reserve are found in a device’s **Properties** window, on the **Advanced** tab. They are visible only when diagnostic mode is turned on.

The default setting for Reserve/Release is None. Once any other Reserve/Release setting is selected, it works automatically, without further user intervention. The Reserve/Release attribute



is supported only on Common Device Interface (CDI) platforms, so if the CDI attribute in a device's **Properties** is set to Not Used, then Reserve/Release settings are ignored.

For newer hardware, once a Reserve/Release setting (other than None) has been selected, the appropriate Persistent Reserve commands are automatically issued before a device is opened for reading or writing, and before the device is closed. With older hardware, a SCSI-2 Reserve command is issued before opening the device, and a SCSI-2 Release command is issued after the device is closed.

Reserve/Release has these possible settings:

- None (the default)
- Simple
- Persistent Reserve
- Persistent Reserve + APTPL (Activate Persist Through Power Loss)

The Persistent Reserve Key attribute has also been added. It is used with Persistent Reservation calls.

### Restrictions for use of the SCSI Reserve/Release setting

There are restrictions for using the SCSI Reserve or Release setting.

Consider the following:

- It is available on CDI platforms only. Consequently, since CDI is not supported within an NDMP environment, Reserve/Release is not supported with NDMP.
- Not all drives support persistent Reserve/Release. (All drives support at least simple reserve release. The code automatically drops back from Persistent +APTPL or Persistent to Simple on drives that do not support Persistent.)
- SCSI resets can clear Simple reservations at the device.
- Even with Reserve/Release, there is no guarantee against data loss.
- If the operating system has its own Reserve/Release feature, that feature must be disabled in order for the NetWorker Reserve/Release feature to work.
- Even if all of the enterprise's NetWorker storage nodes have this feature enabled, then it is possible that, on the storage node where a backup operation is run, data loss can be caused by the operating system's utilities or by third-party programs.

## DDS on NDMP nodes in a SAN environment

You can configure shared drives between NDMP nodes in a SAN environment.

Ensure that:

- All the components of a SAN configuration are compatible when DDS is enabled with the NetWorker NDMP feature.
- The Fibre Channel switches are compatible with any NDMP hosts within a SAN.
- NDMP hosts and libraries in the SAN are compatible with each other.
- The NDMP nodes that will share the drives are homogeneous.

**Note:** The current NDMP implementation does not allow the sharing of drives between non-homogeneous NDMP nodes. There is, however, no inherent limitation within DDS that would prevent this.

## NetApp zoning requirements for DDS in a SAN environment

To configure DDS with NetApp filers, a zoned SAN configuration is required. Zoning is a feature of the Fibre Channel switch.


Consider the following when configuring DDS with NetApp filers:

- The NetApp zone, which contains only the NetApp filers and tape devices, must be configured on the Fibre Channel switch. This NetApp zone may also include the robotic arm and must also be configured in an arbitrated loop.
- All non-NetApp servers that are attached to the same Fibre Channel switch must be excluded from the NetApp zone. A separate zone must be configured for the non-NetApp servers, in which an arbitrated loop may, or may not be a requirement.
- The NetApp zone and all other zones can overlap on the tape devices within the SAN, so that the tape devices are visible to both zones.

This figure illustrates a basic DDS configuration with NDMP:

- Both the NDMP Filer (sn\_1) node and the NDMP Filer (sn\_2) node must be the same kind for DDS to be enabled.
- The hardware id for Drive 1 is drive\_1
- The hardware id for Drive 2 is drive\_2.

**Figure 18** DDS with NDMP



## DDS attributes in the device properties

Configure the attributes that DDS uses, in the **Properties** window for a device.

The attributes include:

- Hardware ID
- Shared Devices

### Hardware ID attribute

The Hardware ID attribute tracks the drives that are shared between multiple hosts. Device instances that share the same physical drive across multiple hosts have the same hardware ID. The device autoconfiguration process automatically assigns the Hardware ID to a device, or it is added when manually configuring a device. Users cannot edit the Hardware ID.

You can view the Hardware ID in the **Properties** window for a device, on the **General** tab, in the **Device Sharing** area.

NetWorker generates the Hardware ID when a device is scanned or configured. The Hardware ID consists of the following components:

- Hardware serial number
- Device type
- Worldwide part number (WWPN)
- Worldwide name (WWN)

## Shared Devices attribute

The Shared Devices attribute appears on the **Operations** tab of a device's **Properties** window when in diagnostic mode. It features values that can be used to manipulate all shared instances of a drive simultaneously. This attribute enables or disables all devices that share the same Hardware ID with a single action. The following table lists allowed values and descriptions for the attribute.

**Table 51** Shared Devices attributes

| Value       | Description                                                                                                                                              |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable All  | When selected, enables all devices with the same Hardware ID.                                                                                            |
| Disable All | When selected, disables all the devices with the same Hardware ID.                                                                                       |
| Done        | This value is the default setting. After the server has enabled or disabled all devices with the same Hardware ID, the attribute value is reset to Done. |

You cannot configure the Shared Devices attribute with the `jbconfig` program.

## Idle Device Timeout attribute and DDS

A tape might remain mounted in a drive after a backup completes. Other requests for the drive from another device path must wait during this timeout period. Use the Idle Device Timeout attribute to adjust the timeout value.

The Idle Device Timeout attribute is not specifically a DDS attribute, but is useful in configuring shared drives. This attribute appears on the device **Properties** window on the **Advanced** tab when displayed in Diagnostic Mode. The default value is 0 (zero) minutes, which means that the device never times out and you must manually eject the tape.

If the device belongs to a library, you can also specify the Idle Device Timeout value for all devices in the library. However, the library value will take effect only on those devices whose **Idle Device Timeout** value is 0. The Idle Device Timeout value for a library is located on the **Timer** tab of the library **Properties** window.

## Max active devices

In a DDS environment, use the Max active devices attribute, on the **General** tab of the Storage Node resource to define the maximum number of active devices for a storage node.

This attribute sets the maximum number of devices that NetWorker may use from the storage node in a DDS configuration. In large environments with media libraries that have a large number of devices, storage nodes might not have the ability to optimize all the drives in the library. The Max active devices attribute allows you to limit the number of devices that the storage node uses at a specified time, which allows the storage node to have access to all the devices in the library, but does not limit the storage node to the number of devices it can fully optimize.

## Configuring NDMP devices

You can back up NDMP data to an NDMP or non-NDMP device in a standalone or library configuration. You can also back up NDMP data to ACSLS controlled silos.

### Configuring NetApp in a cluster mode

For NetApp in a cluster mode, tape drivers are only attached to physical nodes. The tape drivers are not accessible from the Vserver. To enable backups to the tape drivers, you must configure a NAS client with a physical node on the NetWorker server. Meaning you must configure Device and storage pool with the physical node. The devices are also accessible from the cluster. Another option is to create a NAS client of the cluster, and configure Device and Storage pool with the cluster.

### Configuring a standalone NDMP device

Use the NetWorker Management Console (NMC) to configure a standalone Network Data Management Protocol (NDMP) tape device for Direct NDMP backups.

#### Procedure

1. In the **Administration** window, click **Devices**.
2. In the navigation tree, right-click **Devices**, and then select **New**.
3. In the **Name** attribute, specify the NDMP device in the format:

`rd=NAS_hostname:NAS_device_handle (NDMP)`  
where:

- *NAS\_hostname* is the hostname of the NAS that has the NDMP device attached.
- *NAS\_device\_handle* is the path of the device.

 **Note:** Configure the NDMP device as a remote device and add **(NDMP)** after the pathname. Otherwise, you receive a message similar to the following:

NDMP device name shall be in `rd=snode:devname (NDMP)` format

4. In the **Media Type** attribute, specify the device type.
5. Specify a valid NAS administrator account in the **Remote User** attribute.
6. Specify the password for the NAS administrator account in the **Password** attribute.
7. On the **Configuration** tab:
  - a. Select the **NDMP** checkbox. You can only set this attribute when you create the device. You cannot change the NDMP attribute after you create the device. To change the device configuration, delete and re-create the device.
  - b. Set the **Target Sessions** attribute to 1. NDMP devices do not support multiplexing.
  - c. The **Dedicated Storage Node** attribute must remain at the default value: **No**.
8. Under the **Advanced** tab, the **CDI** attribute must remain at the default value: **Not used**.
9. (Optional) Change the block size that is used by the NDMP device.  
By default, NDMP devices use a block size of 60 KB. If required, select a different block size in the **Device block size** field. When you configure the NDMP client, set the *NDMP\_AUTO\_BLOCK\_SIZE* environment variable in the **Application Information** attribute.
10. Click **OK**.

## Configuring an NDMP autochanger

You can use an NDMP autochanger to manage Direct NDMP or Three-party backups with NDMP devices. To configure an NDMP autochanger, use NMC or the `jbconfig` command.

### Configuring an NDMP autochanger with NMC

When you configure an NDMP autochanger in NMC, the NetWorker software first detects the NDMP devices and then configures the library.

#### Procedure

1. In the **NetWorker Administration** window, click **Devices**.
2. Right-click the NetWorker Server, and then select **Configure All Libraries**.
3. On the **Provide General Configuration Information** window, accept the default library type, **SCSI/NDMP**, and then click **Next**.
4. On the **Select Target Storage Nodes** window, click **Create a new Storage Node**.
5. On the **Storage Node Name** field, specify the hostname of the NAS.
6. In the **Device Scan Type** attribute, select **NDMP**.
7. In the **NDMP User Name** and **NDMP Password** fields, specify the NAS administrator account. If DinoStor TapeServer manages the autochanger, specify the DinoStor username and password.
8. Click **Start Configuration**.
9. Click **Finish**.
10. Monitor the **Log** window for the status of the device scan.

When you specify an incorrect username and password combination:

- The Log status window reports:

```
No configured libraries detected on storage node
storage_node_name
```

- The `daemon.raw` file on the NetWorker server reports:

```
NDMP Service Debug: The process id for NDMP service is 0xb6c0b7b0
42597:dvdetect: connect auth: connection has not been authorized
42610:dvdetect: The NDMP connection is not successfully
authorized on host 'storage_node_name'
```

To resolve this issue, relaunch the **Configure All Libraries** wizard and correct the NDMP username and password combination.

**Note:** If the **Log** window reports that NetWorker cannot detect the serial numbers for the library, see [Configuring an NDMP autochanger by using the `jbconfig` command](#) on page 39 for detailed instructions.

### Configuring an NDMP autochanger by using the `jbconfig` command

The NMC interface is the preferred method to configure an NDMP autochanger. Use the `jbconfig` command when you cannot configure the autochanger by using the NMC Configure Library wizard.

#### About this task

The *NetWorker Command Reference Guide* or the UNIX man page provides more information about the `jbconfig` command.

## Procedure

1. Log in to the NetWorker server as root on UNIX, or Administrator on Windows.
2. At the command prompt, type `jbconfig`
3. At the **What kind of jukebox are you configuring** prompt, type 2 to configure an autodetected NDMP SCSI jukebox.
4. When prompted for an NDMP username, specify the NAS administrator account.
5. When prompted for an NDMP password, specify the NAS administrator password.
6. When prompted for the NDMP Tape Server Name, specify the NAS filer hostname.
7. At the **What name do you want to assign to this jukebox device** prompt, provide a name to identify the autochanger.
8. To enable auto-cleaning, accept the default value of **Yes**, otherwise type **no**.
9. At the **Is (any path of) any drive intended for NDMP use? (yes / no) [no]** prompt, type **Yes**.
10. At the **Is any drive going to have more than one path defined? (yes / no) [no]** prompt, type **no** if you will not configure shared devices. Type **yes** to configure shared drives.
11. When prompted, for the first pathname for the NDMP devices in the jukebox, perform the following steps:
  - a. Specify the pathname in the following format:
 

```
NDMP_tape_server_name:device_path
```

 where:
    - *NDMP\_tape\_server\_name* is the hostname of the NDMP Server.
    - *device\_path* is the first device path.

Do not type a slash before the device name. Although the `jbconfig` command completes without errors, the NetApp filer will not recognize the tape device or autochanger.
12. Complete the prompts for the second device.
13. In the **Enter the drive type of drive 1** prompt, specify the number that corresponds to the NDMP device type.
14. If each drive in the autochanger is the same model, then type **yes**. Otherwise, type **no**, and then specify the appropriate device types for each additional autochanger device.
15. When prompted to configure another autochanger, type **no**.

## Changing the block size of an NDMP device

By default, the block size that is used to write data to an NDMP backup is 60KB. With the exception of Celerra, when you specify the `NDMP_AUTO_BLOCK_SIZE=Y` variable for an NDMP client, an NDMP device can use the value that is defined in its Device block size attribute.

### About this task

To determine the block sizes that are supported by the NDMP filer before setting the block size for an NDMP device, consult the applicable vendor documentation.

To change the block size that is defined for the NDMP device, perform the following steps:

## Procedure

1. From the **View** menu, select **Diagnostic Mode**.
2. In the **Devices** window, right-click the NDMP device, and then select **Properties**.

3. On the **Advanced** tab, select a value in the **Device block size** field.

 **Note:** The selected block size must not exceed the block size that is configured on the NAS filer.

4. Click **Ok**.

### Message displayed when CDI enabled on NDMP or file type device

If you enable the CDI feature for an NDMP tape device or file type device (FTD), a message similar to the following appears:

```
nsrd: media notice: The CDI attribute for device "/dev/rmt/3cbn" has
been changed to "Not used".
```

To avoid this message, do not enable the CDI attribute for these device types.

## Configuring NDMP-DSA devices

When you use DSA, NetWorker sends the NDMP data to a NDMP-DSA device, which includes tape, virtual tape, AFTD, and Data Domain devices. The steps to configure a NDMP-DSA device for a specified device type is the same as configuring a non-NDMP device. The *NetWorker Administration Guide* provides detailed information.

## Configuring the Clone Storage Node

When cloning NDMP data, specify the destination storage node, called the clone “write source” (the device that receives the clone data), in the Clone storage nodes attribute. The *NetWorker Administration Guide* provides details.

## Pools requirements for NDMP

When you create a pool for non-NDMP devices, select only the devices that are required by the NDMP clients.

NetWorker cannot send bootstrap and index backups to an NDMP device. When you do not configure a non-NDMP devices or a non-NDMP device is not available to receive the index and bootstrap backups, the NDMP client backup appears to hang. Configure a separate pool to direct the index and bootstrap to a non-NDMP device.

Auto media verification in the Pool resource does not support NDMP.

When an NDMP client backup is a member of a clone-enabled group, configure a clone pool with non-NDMP devices that are local to the NetWorker server to receive the clone bootstrap and index.

# Configure NetWorker for NDMP backup and clone operations

This section explains how to configure NetWorker for NDMP backup and clone operations.

## Creating and configuring the NDMP client resource

Use the NMC Client Configuration wizard to create the NDMP client or create the client manually. It is recommended that you use the NMC Client Configuration wizard to create NDMP clients.

## Using the Client Configuration wizard

### About this task

Use the NMC Client Configuration wizard to create the NDMP client.

### Procedure

1. From the **Administration** window in NMC, click **Protection**.
2. In the expanded left pane, select **Clients**, and then select **Protection > New Client Wizard**.
3. On the **Specify Client Information** window:
  - a. In the **Client Name** field, specify the hostname of the filer.
  - b. (Optional) Add comments in the **Comment** field.
  - c. (Optional) In the **Tag** field, specify the name of the tag for the dynamic group in which you want to add this client.
  - d. (Optional) In the **Groups** area, select an existing group, in which to add the client.
  - e. In the **Type** area, select **NDMP**, and then click **Next**.
4. On the **Specify the NDMP Client Credentials** window:
  - a. In the **NDMP User Name** field, specify a valid NAS administrator account.
  - b. In the **NDMP Password** attribute, specify the password for the NAS administrator account, and then click **Next**.
5. In the **Specify the NDMP Client Backup Options** window:
  - a. In the **NDMP backup type** attribute, select or specify the backup type:
    - `dump`— Traverses a file tree in mixed width first and depth-first order.
    - `smtape`—Performs a block-level backup of a SnapMirror volume.
  - b. In the **NDMP Array Name** field, specify the logical name that is assigned to the NDMP NAS array.  
  
 The **NDMP Array Name** field enables you to configure the same NAS device with multiple NDMP clients that have different host IDs.  
 **Note:** NDMP clients that use the same NAS device must have the same NDMP array name.
  - c. Review the **App Info** options and disable options, as required. It is recommended that the default options remain enabled.

**Table 52** Application information variable types

| App Info Type | Description                                                                                                                                                                                                                  |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HIST          | Enables the backup of index data. If you do not select this option, you can only perform full recoveries of the backup data.                                                                                                 |
| UPDATE        | Enables the backup process to update the last backup dates in database on the NDMP client, after the backup completes. This applies to NetApp when using backup type SMTape. It has no effect when using other backup types. |



**Table 52** Application information variable types (continued)

| App Info Type          | Description                                                                                                                                                                                                                                                                                            |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DIRECT                 | Enables DAR or DDAR support. <a href="#">DAR and DDAR</a> on page 257 provides more information.                                                                                                                                                                                                       |
| Use Token-Based Backup | Enables the NDMP backup to use last backup time tokens to decide what files to backup. Not all NDMP clients support token based backup. When you select this option and the NDMP data server on the client does not support token based backups, NetWorker performs the backup by using backup levels. |

- d. In the **Advanced App Info** field, specify additional NAS specific environments variables, one per line. The following table provides a list of the available **Application Information** environment variables for each NAS.

 **NOTICE** Environment variables are case-sensitive. Use an equal (=) sign to separate the environment variable name from its value.

**Table 53** Vendor-specific Application Information variables

| Variables                    | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>FILESYSYTEM=path</i>      | Optional. Use this variable to define the file system to back up and override the value in the Save set attribute of the client.                                                                                                                                                                                                                                                                                                                                                                                                       |
| <i>DIRECT=y/n</i>            | Optional. When you use DAR or DDAR, you must set this value to y.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <i>EXCLUDE=string</i>        | Optional. This string specifies the files to exclude from backup. The following rules apply: <ul style="list-style-type: none"> <li>• The string must be a file name. Use file names, not absolute paths.</li> <li>• You can use the asterisk (*) as a wildcard, only when * is the first or last character in the string, or both.</li> <li>• To list multiple files, separate each name with a comma. A comma cannot appear as part of the file name. You cannot use spaces.</li> <li>• You can specify up to 32 strings.</li> </ul> |
| <i>EXTRACT_ACL=y</i>         | Optional. Specify this variable to recover ACLs when you use DAR.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <i>SMTAPE_BREAK_MIRROR=Y</i> | Optional when you use SMTape. During the recovery of the mirror, setting <i>SMTAPE_BREAK_MIRROR=Y</i> ensures that the mirror breaks and the volume becomes available for reuse. If you do not set the variable or you specify <i>SMTAPE_BREAK_MIRROR=N</i> , the mirror                                                                                                                                                                                                                                                               |

**Table 53** Vendor-specific Application Information variables (continued)

| Variables                       | Definition                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                 | remains in the same state as at the time of backup.                                                                                                                                                                                                                                                                                                            |
| <i>SMTAPE_DELETE_SNAPSHOT=Y</i> | Optional when you use SMTape. When backing up the filer volumes, setting <i>SMTAPE_DELETE_SNAPSHOT=Y</i> ensures the removal of the mirror created during the backup, at the end of backup. If you do not set the variable or you specify <i>SMTAPE_BREAK_MIRROR=N</i> , each backup attempt creates a new snap mirror image.                                  |
| <i>RECURSIVE=y</i>              | Optional for DAR and DDAR recoveries. <i>RECURSIVE=y</i> ensures the correct recovery of ACLs, permissions, and ownerships for all interim directories selected in the recover operation.                                                                                                                                                                      |
| <i>USE_TBB_IF_AVAILABLE=n</i>   | Optional. The NetWorker software enables TBB automatically. Specify this variable to disable TBB support for incremental backups. This value reverts the backup to the native level-based backup of the NAS.                                                                                                                                                   |
| <i>UTF8=n</i>                   | Optional. Provides support for UTF-8 formatted data. When you do not define this variable, the default value is n.<br><br>When you set <i>UTF8=Y</i> during an NDMP client backup and the backup contains path names with non-ASCII characters, an index-based recovery of this backup fails with the error:<br><br>RESTORE: could not create path<br>pathname |
| <i>NDMP_AUTO_BLOCK_SIZE=Y</i>   | Optional. Specify this variable to override the default block size of 60 KB when writing NDMP backups to an NDMP device. Uses the block size value defined in the Device block size attribute when you labeled the NDMP volume.                                                                                                                                |
| <i>NDR_DSA_NODE</i>             | When this <i>NDR_DSA_NODE</i> is used, <i>nsrdsa_save</i> binary is started on the specified sn-server during the backup. During recover, <i>nsrdsa_recover</i> binary is started on the specified sn-server.                                                                                                                                                  |
| <i>NSR_DSA_REC_SERVER</i>       | During recover, <i>nsrdsa_recover</i> binary is started on the specified remote storage node.                                                                                                                                                                                                                                                                  |

6. Click **Next**.
7. On the **Select the NetWorker Client Properties** window:

- a. In the **Priority** field, specify the order in which the NetWorker server contacts clients in a protection group for backup. The attribute can contain a value between 1 and 1,000. The lower the value, the higher the priority.
  - b. In the **Parallelism** attribute:
    - For Direct-NDMP, set the **Parallelism** attribute to 1.
    - For NDMP-DSA, the parallelism value depends on the NAS capabilities and set parallelism to a value that is appropriate for the NAS. Parallelism values of 4 to 8 are common. In general, the best parallelism setting depends on filer configuration and the amount of installed RAM.
  - c. In the **Remote Access** attribute:
 

Specify the root account on Linux/UNIX, and/or the administrator account on Windows, of any computer that you use to browse backups of the NAS. Specify each account on a separate line. For example:

```
administrator@windows_hostname

root@linux_hostname
```
  - d. Select the **Data Domain Interface**. This option specifies the protocol to use if you send the backup data to a Data Domain Device. Available selections are Any, Fibre Channel, or IP.
  - e. Do not select the **Block Based Backup** or **Client Direct** options, as they do not apply to NDMP backups.
  8. Click **Next**.
  9. On the **Client Configuration Summary** window, review the attributes, and then click **Create**.
  10. On the **Client Configuration Results** window, review the results, and then click **Finish** to exit the wizard.
- [Troubleshooting NDMP configuration and backup failures for PowerStore, Unity, Celerra, VNX, and VNXe](#) describes how to resolve errors that you may experience when you configure the NDMP client.

## NetApp Environment Variables

Restart the NetWorker services after setting the following environment variables:

**Table 54** NetApp Environment Variables

| Variables                   | Definitions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>NDMP_RECOVER_LIMIT=1</i> | When you select to restore multiple files, if the files are on different save sets, NetWorker sends multiple NDMP recover requests to the NAS client. With NetApp 8.2 and later releases, the ability to run parallel recoveries if the files are restored to the same destination volume is disabled, and the restore fails. You can use the <i>NDMP_RECOVER_LIMIT</i> parameter to configure the number of sessions on the NetWorker server. For NetApp 8.2 and later releases, set the parameter as follows: |

**Table 54** NetApp Environment Variables (continued)

| Variables                                                                                                                                              | Definitions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                                                                                                        | NDMP_RECOVER_LIMIT=1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <i>RECOVER_FULL_PATHS=y</i>                                                                                                                            | Optional for DAR and DDAR recoveries. <i>RECOVER_FULL_PATHS=y</i> ensures that NetWorker recovers the ACLs, permissions, and ownerships for each interim directories selected in the recover operation.                                                                                                                                                                                                                                                                                                                                                                             |
| <i>NSR_NDMP_TMP_DIR=&lt;path to tmp&gt;</i>                                                                                                            | Use this variable to redirect the tmp files going to /nsr/tmp (optional).<br><b>Note:</b> Path to tmp should be created manually.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <i>NSR_NDMP_RESTORE_LIMIT</i>                                                                                                                          | Use this variable while recovering large NDMP savesets using DAR recovery. The <i>NSR_NDMP_RESTORE_LIMIT</i> environment variable is used to limit memory consumption during recoveries involving a large number of index entries (millions). This is specifically useful if there is not enough swap space or memory for the number of index entries selected for the recovery. If the variable is not set, the recovery might fail with an out of memory error message.<br><b>Note:</b> <i>NSR_NDMP_RESTORE_LIMIT</i> can be set to 50,000 or 1,000,000, but less than 3,000,000. |
| <i>NSR_NDMP_DDAR=y</i><br>This environment variable must be set in the operating system before invoking either the recover or winworkr program.        | Use this variable to recover the permissions of the parent directories of selected files. Specify this variable to perform a DDAR recovery when (optional): <ul style="list-style-type: none"> <li>You set the <i>DIRECT=y</i> variable during the backup</li> <li>The DART version is 5.5 and later</li> </ul> <b>Note:</b> Do not specify <i>NSR_NDMP_DDAR</i> when you also use <i>NSR_NDMP_RECOVER_DIR</i> .                                                                                                                                                                    |
| <i>NSR_NDMP_RECOVER_NO_DAR=y</i>                                                                                                                       | Define this variable to perform a non-DAR recovery when you set the <i>DIRECT=y</i> variable during the backup (optional).                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <i>NSR_NDMP_RECOVER_DIR=y</i><br>This environment variable must be set in the operating system before invoking either the recover or winworkr program. | Specify this variable to perform a DAR recovery when (optional): <ul style="list-style-type: none"> <li>When you set the <i>DIRECT=y</i> variable during the backup</li> <li>The DART version is 5.5 and later</li> </ul>                                                                                                                                                                                                                                                                                                                                                           |

**Table 54** NetApp Environment Variables (continued)

| Variables              | Definitions                                                                                                                                                  |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                        | <p><b>Note:</b> Do not use <i>NSR_NDMP_RECOVER_DIR</i> when you also use <i>NSR_NDMP_DDAR</i>.</p>                                                           |
| <i>NSR_NDMP_NDAR=y</i> | <p>It can be used instead of <i>NSR_NDMP_DDAR</i> when both full and incremental backups are present. <i>NSR_NDMP_DDAR</i> works only with full backups.</p> |

**Note:** To set environment variable on Linux NetWorker server edit `/opt/nsr/admin/networkerrc` and add the variable as shown below: ,

- `NSR_NDMP_DDAR =y`
- `export NSR_NDMP_DDAR`

To set environment variable on Windows NetWorker server navigate to Control Panel > System and Security > System > Advanced system settings > Environment Variables > System variables > New.

- Variable Name = `NSR_NDMP_TMP_DIR`
- Variable Value = `C:\NW\nsr\tmp`

## Performing post Client Configuration Wizard steps

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

### Modifying the Storage Node

#### About this task

On the **Globals (2 of 2)** tab, specify the storage node in the **Storage Nodes** attribute.

The attribute value depends on the type of backup:

- When you perform Direct-NDMP backups with NDMP devices, specify the hostname of the NAS that manages the tape device or autochanger.
- For three-party backups, specify the destination host first.
- For NDMP-DSA backups, specify the hostname of the storage node that manages the tape device or autochanger. If the NetWorker server is the storage node, specify *nsrserverhost*.

**NOTICE** For NDMP-DSA backups, the NetWorker software uses the **Storage Node** attribute field of the NDMP client to determine which host receives the backup data. The `nsrndmp_save` command does not require the `-M` and `-P` options. If you specify the `-M` and `-P` options, they override the **Storage Node** attribute value.

### Enabling Checkpoint Restart

To allow a failed backup for a client to restart from a known good point, you must enable Checkpoint Restart for the NetWorker client resource and configure the number of automatic retries for the backup action in the data protection policy.

#### Procedure

1. In the **Administration** window, click **Protection**.

2. In the expanded left pane, select **Clients**.
3. Right-click the client resource and select **Properties**.  
The **Client Properties** dialog box appears.
4. On the **General** tab, select the **Checkpoint enabled** checkbox.  

**Note:** When you enable **Diagnostic Mode** the **Checkpoint granularity** attribute appears. This attribute does not apply to NDMP.
5. Define the interval at which checkpoints are written during the backup:
  - a. Select the **Apps & Modules** tab on the **Client Properties** dialog box.
  - b. In the **Application information** attribute, specify the `CHECKPOINT_INTERVAL_IN_BYTES` variable.  
 For example, to write a checkpoint after every 1,000,000 bytes, type:  

```
CHECKPOINT_INTERVAL_IN_BYTES=1000000
```

 To define the value by using a different multiplier, specify the multiplier with the numeric value. Supported multipliers include KB, MB, GB, TB, kb, mb, gb, and tb. For example, to write a checkpoint after every 1 GB, type:  

```
CHECKPOINT_INTERVAL_IN_BYTES=1GB
```

**Note:** The checkpoint interval value is automatically rounded up to a multiple of the tape block size.
6. Click **OK** on the **Client Properties** dialog box.

## Adding NDMP Client Properties

### Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Clients**.
3. Right-click a **Client**, and then select **New Client Properties**. The **Client Properties** screen appears.
4. Select the **Apps & Modules** tab. Select the **NDMP** attributes of the **Client**. The **NDMP** options are as follows:
  - **NDMP**—Select this box to indicate whether this client is an NDMP client.
  - **NDMP multistreams enabled**—Do not select this box. Only Isilon supports multistreaming.
  - **NDMP log successful file recovery**—By default, NetWorker does not print each successfully recovered file name in the log messages, because this logging impacts performance and takes up space. To enable the logging of successful recoveries for each file, select this checkbox.
  - **Disable IPv6**—Check this box to disable IPv6 on NDMP backup and recovery.
  - **NDMP array name**—This name is the logical name that is assigned to the array in NDMP NAS array configurations.
  - **NDMP vendor information**—This attribute contains NDMP client vendor information.
5. Click **OK**.

## Configuring the NDMP client manually

It is recommended that you create Network Data Management Protocol (NDMP) clients by using the Client Configuration wizard. If you create the NDMP client manually, then the configuration

details for each attribute in the Client Configuration wizard apply when you create the client manually.

Review this information before you configure an NDMP client manually:

- For an NDMP configuration that includes Storage Node resources, configure a Client resource for each storage node that you define for an NDMP backup and clone operation.
- For NDMP three-party storage nodes that use NDMP devices, repeat these steps for each NDMP storage node.
- For NDMP-DSA storage nodes, create the NetWorker Client resources in the same manner as non-NDMP clients. The *NetWorker Administration Guide* provides details on how to create a non-NDMP Client resource.
- NDMP does not support the use of directives including AES encryption. The NetWorker software ignores any value that you define in the **Directives** attribute for an NDMP client.
- When you select **Checkpoint enabled** on the **General** tab, do not modify the **Checkpoint granularity** attribute. NDMP backups do not support checkpoint granularity and the NetWorker software ignores any value that you define for this attribute.
- If the NAS supports the NDMP snapshot management extension, then you can browse and mark individual file systems for backup instead of specifying the save sets in the **Save set** attribute. You cannot use the **Save set browse** icon to browse the NDMP file system until you:
  - Select the **NDMP** checkbox, on the **Apps & Modules** tab.
  - Specify the NDMP username and password in the **Remote user and password** fields on the **Apps and Modules** tab.

 **Note:**

- Celerra, VNX, VNXe, and NetApp C-mode do not support Snapshot Management Extension. Only NetApp 7-Mode supports Snapshot Management Extension.
- Isilon, Celerra, VNX, VNXe, and NetApp C-mode filers do not allow you to browse. Only NetApp 7-Mode allows you to browse.

## Performing schedule backup and clone operations

Data Protection Policies provide you with the ability to schedule backup and clone operations, to protect NDMP data.

You can use the NDMP protocol to protect data on NAS devices.

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

## Overview of protection policies

A protection policy allows you to design a protection solution for your environment at the data level instead of at the host level. With a data protection policy, each client in the environment is a backup object and not simply a host.

Data protection policies enable you to back up and manage data in a variety of environments, as well as to perform system maintenance tasks on the NetWorker server. You can use either the **NetWorker Management Web UI** or the NMC **NetWorker Administration** window to create your data protection policy solution.

A data protection policy solution encompasses the configuration of the following key NetWorker resources:

## Policies

Policies provide you with a service-catalog approach to the configuration of a NetWorker datazone. Policies enable you to manage all data protection tasks and the data protection lifecycle from a central location.

Policies provide an organizational container for the workflows, actions, and groups that support and define the backup, clone, management, and system maintenance actions that you want to perform.

## Workflows

The policy workflow defines a list of actions to perform sequentially or concurrently, a schedule window during which the workflow can run, and the protection group to which the workflow applies. You can create a workflow when you create a new policy, or you can create a workflow for an existing policy.

A workflow can be as simple as a single action that applies to a finite list of Client resources, or a complex chain of actions that apply to a dynamically changing list of resources. In a workflow, some actions can be set to occur sequentially, and others can occur concurrently.

You can create multiple workflows in a single policy. However, each workflow can belong to only one policy. When you add multiple workflows to the same policy, you can logically group data protection activities with similar service level provisions together, to provide easier configuration, access, and task execution.

## Protection groups

Protection groups define a set of static or dynamic Client resources or save sets to which a workflow applies. There are also dedicated protection groups for backups in a VMware environment or for snapshot backups on a NAS device. Review the following information about protection groups:

- Create one protection group for each workflow. Each group can be assigned to only one workflow.
- You can add the same Client resources and save sets to more than one group at a time.
- You can create the group before you create the workflow, or you can create the group after you create the workflow and then assign the group to the workflow later.

## Actions

Actions are the key resources in a workflow for a data protection policy and define a specific task (for example, a backup or clone) that occurs on the client resources in the group assigned to the workflow. NetWorker uses a work list to define the task. A work list is composed of one or several work items. Work items include client resources, virtual machines, save sets, or tags. You can chain multiple actions together to occur sequentially or concurrently in a workflow. All chained actions use the same work list.

When you configure an action, you define the days on which to perform the action, as well as other settings specific to the action. For example, you can specify a destination pool, a retention period, and a target storage node for the backup action, which can differ from the subsequent action that clones the data.

When you create an action for a policy that is associated with the virtual machine backup, you can select one of the following data protection action types:


- **Backup** — Performs a backup of virtual machines in vCenter to a Data Domain system. You can only perform one VMware backup action per workflow. The VMware backup action must occur before clone actions.
- **Clone** — Performs a clone of the VMware backup on a Data Domain system to any clone device that NetWorker supports (including Data Domain system or tape targets). You can specify multiple clone actions. Clone actions must occur after the Backup action.



You can create multiple actions for a single workflow. However, each action applies to a single workflow and policy.

The following figure provides a high level overview of the components that make up a data protection policy in a datazone.

**Figure 19** Data Protection Policy



## 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 20** 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.

## Overview of configuring a new data protection policy

### About this task

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

**Example 3** Example of a data protection policy with 2 workflows (continued)

**Figure 21** 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*.

## Strategies for traditional backups

The primary considerations for a traditional backup strategy are the groups of Client resources, the workflows that define the series of actions that are associated with the backup, and the schedule for the backup.

### 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 traditional backups

A protection groups for traditional backups identifies the client resources to back up.

Traditional backups support the following types of protection groups:

- 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.

Create multiple groups to perform different types of backups for different Client resources, or to perform backups on different schedules. For example:

- Create one group for backups of clients in the Accounting department, and another group for backups of clients in the Marketing department.

- Create one group for file system backups and one group for backups of Microsoft Exchange data with the NetWorker Module for Microsoft.
- Create one group for a workflow with backups actions that start at 11 p.m., and another group for a workflow with backup actions that start at 2 a.m.

 **Note:** A Client resource can belong to more than one group.

### 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 Datanones** 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.

#### About this task

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](#).

## Actions sequences in traditional backup workflows

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


A workflow for a traditional backup can optionally include a probe or check connectivity action before the backup, and a clone action either concurrently with or after the backup.

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

### Workflow path from a traditional backup action

The only action that can follow a traditional backup is a clone action.

**Figure 22** Workflow path from a traditional backup action



## Creating a check connectivity action

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

### Before you begin

Create the policy and 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, monthly, or reference schedule for the action:
    - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
    - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
    - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
  9. Click the icon on each day to specify whether to check connectivity with the client.
- The following table provides details on the icons.

**Table 55** Schedule icons

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

To check connectivity every day, select **Execute** from the list, and then click **Make All**.

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 successful connectivity is achieved with all clients in the assigned group, select the **Succeed only after all clients succeed** checkbox.
  - To specify that the connectivity check is successful if connectivity is achieved with one or more clients in the assigned group, clear the checkbox.

12. Click **Next**.

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

13. Optionally, 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

For NetApp clients that use checkpoint restart, on the **Advanced** tab, set the value of the **Client retries** attribute to a number greater than 0.

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) 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.
- If a schedule is associated to an action, then override option is disabled.

19. Click **Next**.

The **Action Configuration Summary** page appears.

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

**After you finish**

Optionally, 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 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. In-built commands from NetWorker should be avoided as probe command.

**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, monthly, or reference schedule for the action:
  - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
  - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
  - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
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 56** 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 action 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 action if any of the probes 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

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.
- If a schedule is associated to an action, then override option is disabled.

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 traditional backup action

A traditional backup is a scheduled backup of the save sets defined for the Client resources in the assigned group for the workflow.

#### Before you begin

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

#### 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 **Backup**.
6. From the secondary action list, select the backup type, for example, **Traditional**.
7. (Optional) From the **Force Backup Level** list select a backup level.

For workflows that have more than one scheduled backup within a 24-hour period, use the **Force Backup Level** attribute to allow more than one backup to occur at two different backup levels in a 24-hour period. When you select a backup level in the **Force Backup Level** attribute, the first backup is performed at the scheduled backup level. Each subsequent occurrence of the backup action in the next 24 hours occurs at the level defined in the **Force Backup Level** attribute. For example, if the level defined by the schedule is Full and the **Force Backup Level** attribute is set to Incr, the first backup started by the action occurs at a level full and subsequent backups, within 24 hours of the start of the full backup are incremental. By default this option is cleared, which means that if the action runs multiple backup operations in a 24 period, all the backups occur at the scheduled backup level.

8. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
9. 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.
10. Specify a weekly, monthly, or reference schedule for the action:
  - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
  - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
  - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
11. To specify the backup level to perform when **Define** option under **Select Schedule** is selected, click the icon on each day. If it is **Select** option under **Select Schedule**, choose the

customized NSR schedule using the drop down, after specifying the backup level within the NSR schedule resource.

The following table provides details about the backup level that each icon represents.

**Table 57** Schedule icons

| Icon                                                                                | Label                                                                                                                                                         | Description                                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|    | <b>Full</b>                                                                                                                                                   | Perform a full backup on this day. Full backups include all files, regardless of whether the files changed.                                                                                                               |
|    | <b>Incr</b>                                                                                                                                                   | Perform an incremental backup on this day. Incremental backups include files that have changed since the last backup of any type (full or incremental).                                                                   |
|    | <b>Cumulative Incr</b>                                                                                                                                        | Perform a cumulative incremental backup. Cumulative incremental backups include files that have changed since the last full backup.                                                                                       |
|  | <b>Logs Only</b>                                                                                                                                              | Perform a backup of only database transaction logs.                                                                                                                                                                       |
|  | <b>Incremental Synthetic Full</b><br> <b>Note:</b> Not supported for NDMP. | Perform an incremental synthetic backup on this day. An incremental synthetic full backup includes all data that changed since the last full backup and subsequent incremental backups to create a synthetic full backup. |
|  | <b>Skip</b>                                                                                                                                                   | Do not perform a backup on this day.                                                                                                                                                                                      |

To perform the same type of backup on each day, select the backup type from the list and click **Make All**.

NetWorker does not support the use of synthetic full backup levels for NDMP data.

Celerra, Isilon, VNX, Unity, and NetApp filers with NDMP version 4 or later support token-based backups (TBB) to perform NDMP full and incremental backups. NetWorker supports the same number of incremental levels that the NAS vendor supports. Celerra, Isilon, and NetApp documentation provide the maximum number of incremental levels that the TBB incremental backup can support.

When you configure TBB after you update the NetWorker server from 7.6 SP1 or earlier, the first incremental backup does not occur until after one complete full backup.

Filers that do not support TBB, do not support incremental backups. If you select the level Incr, the NetWorker server performs a full backup.



Verify that the NAS storage vendor supports NDMP incremental backups before you use this feature.

12. Click **Next**.

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

13. From the **Destination Storage Node** box, select the storage node with the devices on which to store the backup data.
14. From the **Destination Pool** box, select the media pool in which to store the backup data.
15. From the **Retention** boxes, specify the amount of time to retain the backup data.

After the retention period expires, the save set is removed from the client file index and marked as recyclable in the media database during an expiration server maintenance task.

When you define the retention policy an NDMP client, consider the amount of disk space that is required for the client file index. NDMP clients with several thousands of small files have significantly larger client file indexes on the NetWorker server than a non-NDMP client. A long retention policy for an NDMP client increases disk space requirements on the file system that contains the client file indexes.

16. From the **Client Override Behavior** box, specify how NetWorker uses certain client configuration attributes that perform the same function as attributes in the Action resource:

- **Client Can Override**—The values in the Client resource for **Schedule**, **Pool**, **Retention policy**, and the **Storage Node** attributes take precedence over the values that are defined in the equivalent Action resource attributes.

 **Note:** If the NetWorker policy action schedule is set to the **Skip** backup level, the **Client can Override** option is not honored. For NetWorker to consider the **Client can Override** option, change the action schedule to a different level.

- **Client Can Not Override**—The values in the Action resource for the **Schedule**, **Destination Pool**, **Destination Storage Node**, and the **Retention** attributes take precedence over the values that are defined in the equivalent Client resource attributes.
- **Legacy Backup Rules**—This value only appears in actions that are created by the migration process. The updating process sets the **Client Override Behavior** for the migrated backup actions to **Legacy Backup Rules**.

17. Click **Next**.

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

18. 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.

19. 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.

20. 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.

21. 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

22. 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**.
- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

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

**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.

23. 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.

24. 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.

25. (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.

26. (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.
- If a schedule is associated to an action, then override option is disabled.

27. 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**.
28. 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:
- On Linux, to send an email notification, type the following command:
- For NetWorker Virtual Edition (NVE), to send an email notification, type the following command:

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

```
mail -s subject recipient
```

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

- 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.

29. Click **Next**.

The **Action Configuration Summary** page appears.

30. 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 after the backup. A clone action is the only supported action after a backup action in a workflow.

### Cloning NDMP save sets

You can clone Direct-NDMP and NDMP-DSA save sets by using the same methods used to clone non-NDMP save sets.

Before you clone NDMP save sets, review these requirements:

- To clone Direct-NDMP or Three-party backup data:
  - The source NAS must run NDMP version 3 or later.
  - The destination NAS can run any version of NDMP, but you cannot clone a volume cloned with NDMP earlier than version 3 to another volume.
  - You cannot clone NDMP save sets to a non-NDMP device.
  - You can clone NDMP tapes from one NDMP host to another NDMP host of the same type. For example, you can clone tapes from a NetApp filer with an attached library to another NetApp filer or to the same filer.
- You require two NDMP devices to clone the NDMP save sets, one device to perform the read operation and one device to perform the write operation.
- You must clone NDMP-DSA backups to non-NDMP devices. You can however, clone NDMP-DSA save from one type of tape device to another. For example you can clone save sets on a DLT device to an AIT device.
- Use the nsrclone program to clone NDMP save sets from a command prompt. The *NetWorker Command Reference Guide* or the UNIX man pages provide more information on nsrclone usage.

#### 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

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 **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, monthly, or reference schedule for the action:
  - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
  - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
  - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
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 58** 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

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 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.

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.
- If a schedule is associated to an action, then override option is disabled.

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 traditional backup workflows

**Figure 23** Traditional backup workflow



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.

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**.

## Performing manual NDMP backups

After you configure the NetWorker server for NDMP backup data operations, you can perform manual NDMP backups.

On Windows, you can manually back up NDMP data by using the NetWorker User program. The method to backup NDMP data is the same as a non-NDMP local backup. You cannot perform a three-party backup with the NetWorker User program.

On Windows and UNIX, you can perform a manual backup from a command prompt by using the `nsrndmp_save` command.

Before performing a manual backup by using the `nsrndmp_save` command or the NetWorker User program, review these requirements:

- You can only perform manual Direct-NDMP backups from a NetWorker server.
- You can start a manual NDMP-DSA backup from a NetWorker server, storage node, or client. When you do not start the NDMP-DSA backup from the NetWorker server, the `servers` file on the NetWorker server and storage node, must contain the hostname of the host that initiates the backup.
- Before you perform a manual backup, you must configure the NDMP client on the NetWorker server. Manual backups use client configuration information for example, the variables that are defined in the **Application Information** attribute of an NDMP client.
- Direct-NDMP and three-party NDMP backups support manual DAR backups when the NDMP client contains the `DIRECT=Y` and `HIST=Y` environment variables in the **Application Information** attribute for the NDMP client.

**NOTICE** To use DAR, the NAS filer must use NDMP version 4. The *NetWorker E-LAB Navigator* describes how to determine if a particular NDMP vendor supports DAR.

## Performing an NDMP backup from the command line

### About this task

Use the `nsrndmp_save` command to perform a manual command line NDMP backup.

The `nsrndmp_save` command does not back up the bootstrap. Without the bootstrap, you cannot perform a disaster recovery of the NetWorker server. To back up the bootstrap, run the `nsrpolicy -G nsrpolicy policy_name start` command from the NetWorker server. The `nsrpolicy` command uses the attribute values specified for the policy. For example, the pool and schedule values.

To perform an NDMP backup from the command prompt, use the following syntax:

```
nsrndmp_save -T backup_type -s NetWorker_servername -c clientname -l
backup_level -t date_time -g nsrpolicy path
```

where:

- `backup_type` is a supported backup type for the NAS filer: NetApp supports the `dump` and `smtape` backup types.
- `backup_level` is a full for a full backup, `incr` for an incremental backup. Each NAS supports full backups.
- `date_time` is the date and time of the last backup, which is enclosed in double quotes. Specify this value for `incr` level backups. When you do not specify the date and time, the backup is a native NDMP level-based backup.

**NOTE:** NetWorker only supports or allows Full, Incremental, and Differential level 1 NDMP backups of all NAS.

**NOTICE** During a NetWorker scheduled policy backup, the NetWorker software supplies the date and the time information, and incremental and level backups work as expected.

Use one of these methods to determine the date and time of the last NDMP backup:

- Review the `daemon.raw` file on the NetWorker server or the savegroup completion report for a line similar to the following:

```
42920:nsrndmp_save: browsable savetime=1296694621
```

Use the value after `savetime=` with the `-t` option.

- Specify the date and time of the last backup reported by the `mminfo` command for the NDMP save set.

**Note:** By default the backup operation uses IPv6 addressing, if enabled. NetApp always uses IPV6 if it is enabled, regardless if IPV6 is configured on the NetApp system or not. If the NetApp system is not configured with IPV6, NDMP backup and recovery displays the following error message:

```
42597:nsrndmp_save: mover listen: communication failure
```

```
42801:nsrndmp_save: Failed to send the MOVER_LISTEN message to the
NDMP tape server.
```

You can force the NDMP backup and recovery to ignore IPv6 and instead use IPv4 in one of two ways:

1. Add the `-f` option to the `nsrndmp_save` or the `nsrndmp_recover` commands as required.
2. Select the **Disable IPv6** checkbox on the Apps & Modules tab in the **Client Properties** window.

## Example of an NDMP backup

To perform an incremental backup of a NetApp client that is named `mynetapp`, perform the following steps:

1. Determine the time of the last full backup:

```
mminfo -v -c mynetapp
```

**Table 59** NDMP backup

| client   | date     | time     | size   | ssid       | fl   | lvl  | name      |
|----------|----------|----------|--------|------------|------|------|-----------|
| mynetapp | 08/16/15 | 15:23:58 | 1853MB | 3864812701 | cbNs | full | /.../set1 |
| mynetapp | 08/17/15 | 15:39:58 | 815MB  | 3848036430 | cbNs | incr | /.../set2 |

2. Specify the last backup time in `nsrndmp_save` command:

```
nsrndmp_save -T dump -s my_nwserver -c mynetapp -l incr -t "02/16/11
15:23:58" -g mygroup path
```

For NDMP-DSA backups, the NetWorker software uses the Storage Node attribute field of the NDMP client to determine which host receives the backup data. The `nsrndmp_save` command does not require the `-M` and `-P` options. If you specify the `-M` and `-P` options, they override the **Storage Node** attribute value. The *NetWorker Command Reference Guide* and the `nsrndmp_save` man page on UNIX provide more information.

## Troubleshooting NDMP configuration and backup failures for NetApp

This section provides a list of the possible causes and the resolutions for NDMP backup failures.

### Unable to connect to NDMP host *hostname*

This message appears when the NetWorker server cannot create or modify an NDMP client.

To resolve this issue ensure that the environment meets the following requirements:

- Username and password specified for the client is correct and has sufficient permissions to perform NDMP operations.
- NDMP service is running on the filer.

## NetWorker features not supported on NetApp NDMP v3 and earlier

Features such as Checkpoint restart require NDMP v4.

To verify the NDMP version, perform the following steps.

1. Log in to the NetApp host as root or as a Windows Administrator.
2. Display the NDMP version:

```
ndmpd version
```

To change the NDMP version:

1. Log in to the NetApp host as root or as Windows Administrator.
2. Stop the **NDMP** process:

```
ndmpd off
```

3. Change the **NDMP** version:

```
ndmpd version 4
```

4. Restart the **NDMP** process:

```
ndmpd on
```

## Cannot perform NDMP backup after the NetWorker server licenses expire

If a NetWorker sever running in evaluation mode expires before you authorize the server, NDMP devices remain disabled after the addition of the required licenses and authorization of the NetWorker server.

To re-enable NDMP devices, perform the following steps:

1. To connect to the NetWorker server, use NMC, and then click the **Devices** button.
2. In the **Devices** windows, right-click the NDMP device, and then select **Properties**.
3. Click the **Configuration** tab, and then set the **Target Sessions** attribute to **1**.
4. Click the **General** tab, and then in the **Enabled** section, select **Yes**.
5. Click **Ok**.

## Failed to store index entries

This error message occurs in the `daemon.raw` file when an index backups fails due to an insufficient amount of swap space.

To resolve this issue, increase the amount of swap space available to the NetWorker server.

 **NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

## IO\_WritePage write failed - No space left on device (28): No space left on device

This error message appears in the `daemon.raw` file when the index backup fails. There is insufficient temporary space to store the index entries before the NetWorker software commits the information into the client file index.

To resolve this issue, specify a new the temp directory with sufficient disk space in one of the following ways:

- Define the `NSR_NDMP_TMP_DIR` environment variable in the Application Information attribute of the client.
- Define the `NSR_NDMP_TMP_DIR` as an operating system environment variable on the NetWorker server.

[Memory and space requirements for NDMP FH updates](#) on page 23 describes how to determine the amount of disk space the NetWorker software requires to temporarily store client files index entries.

**NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

## Error reading the FH entries from save through stdin

This error message appears in the `daemon.raw` file of the NetWorker server when there is a communication error between the `nsrndmp_save` and `nsrndmp_2fh` processes.

Resolve any communication or connection issues, then retry the backup.

**NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

## Cannot find file history info for file name...You may still be able to recover this file with a save set recovery

This error message appears in the `daemon.raw` file of the NetWorker server when file history (FH) information is missing or corrupted for the file that is specified in the error message. For example, NetWorker cannot update the client file index (CFI) with FH information when a backup process interruption occurs during the failover of a clustered NetWorker environment.

You cannot perform an Network Data Management Protocol (NDMP) file-by-file recover or a save set recover when the CFI does not contain the associated FH information.

To recover this file, perform a save set recover from the command prompt. [Performing an NDMP save set recovery from the command prompt](#) on page 105 provides for further information.

**NOTICE** The NetWorker server does not delete the FH files that are stored in the `tmp` directory when the CFI updates fail.

## nsrndmp\_save: data connect: failed to establish connection

This error message appears in the `daemon.raw` file of the NetWorker server for several reasons:

- Network connectivity or name resolution issues exist between the NetWorker server and the NDMP client.
- You specified an incorrect NDMP username or password specified for the NDMP client.
- The NDMP service is not started on the NAS filer.
- The NetWorker server cannot communicate with the NAS filer over port 10000.
- A free port in the NetWorker server's default port range (7937-9936) is not available during an NDMP-DSA backup.  
The *NetWorker Security Configuration Guide* provides more information about NDMP port requirements and configuration.
- A misconfigured loop router. For a Celerra filer, the server route command utility configures the loop router. For NetApp, the route utility configures loop back router. The value of this setup is network-specific and depends on the number of switches and hubs between the NAS filer, NetWorker server, and NetWorker storage node.
- On the host where DSA is running, if the hostname is present in the hosts file, the `nsrdsa_save` process uses this name during backup. The DSA host passes the loopback

entry to the NDMP data server and the connection fails. To resolve this issue, remove the hostname from the localhost list.

Knowledge base articles on the Support website provides detailed troubleshooting information for this error message and other failed to establish connection failures that you might encounter during an NDMP backup.

### nsrndmp\_save: get extension list: communication failure

This message appears during a NDMP local backup when NetWorker cannot determine the filer name.

To resolve this issue, perform the following steps:

1. From a command prompt on the NetWorker server, type:

```
ndmpsup -c NDMP_hostname -o output_filename
```

For example:

```
ndmpsup -c myfiler.mnd.com -o ndmpsup.txt
```

2. Edit the output file that the `ndmpsup` command generates and search for the string **Vendor Name**. Make note of the reported Vendor Name.  
For example:

```
Vendor Name = BlueArc Corp
```

3. Change to the `/nsr/debug` directory on UNIX or the `NetWorker_installation_dir\nsr\debug` directory on Windows.

4. Create new empty file and name it with the following format:

```
ndmpgettextlist_disable_VENDOR_NAME
```

where you replace `VENDOR_NAME` with the vendor name of the filer reported in the `ndmpsup` output file.

For example, to create this file for a BlueArc filer on UNIX, type:

```
touch "ndmpgettextlist_disable_BlueArc Corp"
```

## Monitoring NetWorker Server activities in the Administration window

The **Monitoring** window in the NetWorker **Administration** application enables you to monitor the activities of an individual NetWorker Server.

### About this task

The **Monitoring** window provides the following types of activity and status information:

- Data protection policies, workflows, and individual actions.
- Cloning, recovering, synthetic full backups, and browsing of client file indexes.
- Operations that are related to devices and jukeboxes.
- Alerts and log messages.

You can also perform some management operations from the **Monitoring** window, for example, starting, stopping, or restarting a data protection policy.

## Procedure

1. From the **NMC Console** window, click **Enterprise**.
2. In the **Enterprise** view, right-click the NetWorker Server, and then select **Launch Application**.

The **Administration** window appears.

3. To view the **Monitoring** window, click **Monitoring**.

## About the Monitoring window

On the **Administration** window taskbar, select **Monitoring** to view the details of current NetWorker server activities and status, such as:

- Policies and actions.
- Cloning, recovering, synthetic backups, checkpoint restart backups, and browsing of client file indexes.
- Alerts and log messages, and operations that are related to devices and jukeboxes.

While the **Monitoring** window is used primarily to monitor NetWorker server activities, it can also be used to perform certain operations. These operations include starting, stopping, or restarting a workflow.

The **Monitoring** window includes a docking panel that displays specific types of information. Select the types of information you want to view from the docking panel.

A portion of the **Monitoring** window, which is known as the task monitoring area, is always visible across all windows. A splitter separates the task monitoring area from the rest of the window. You can click and move the splitter to resize the task monitoring area. The arrow icon in the upper right corner of the **Monitoring** window allows you to select which tasks you want to appear in this view.

Smaller windows appear within the **Monitoring** window for each window. Each smaller window, once undocked, is a floating window and can be moved around the page to customize the view. You can select multiple types from the panel to create multiple floating windows that can be viewed simultaneously. The following table describes the various types of information available in the docking panel, and the details each one provides.

**Table 60** Monitoring window panel

| Window                  | Information provided                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Policies/Actions</b> | The <b>Policies</b> tab provides you with status information about all configure policies and the associated workflows and actions. The <b>Actions</b> tab provides you with status information for all actions. <a href="#">Policies/Actions pane</a> on page 81 provides more information.                                                                                                                                   |
| <b>Sessions</b>         | Allows you to customize whether to display all session types, or only certain session types. The information that is provided depends on which session type you select. For example, if you select <b>Save Sessions</b> , the window lists clients, save sets, groups, backup level, backup start time, duration of the backup, devices, rate, and size. <a href="#">Sessions window</a> on page 85 provides more information. |
| <b>Alerts</b>           | Lists the priority, category, time, and message of any alerts. <a href="#">Alerts pane</a> on page 86 provides more information.                                                                                                                                                                                                                                                                                               |
| <b>Devices</b>          | Lists devices, device status, storage nodes, libraries, volumes, pools, and related messages. <a href="#">Devices pane</a> on page 87 provides more information.                                                                                                                                                                                                                                                               |

**Table 60** Monitoring window panel (continued)

| Window            | Information provided                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Operations</b> | <p>Lists the status of all library and silo operations, including <code>nsrjb</code> operations that are run from the command prompt. Also lists user input, libraries, origin, operation data, operation start time, duration of the operation, progress messages, and error messages.</p> <p>When displaying Show Details from the <b>Operations</b> window, the length of time that the window is displayed depends on the value that is typed in the <b>Operation Lifespan</b> attribute on the <b>Timers</b> tab of the <b>Properties</b> dialog box for the corresponding library. To access library properties, click <b>Devices</b> in the taskbar. By default, this pane is hidden.</p> |
| <b>Log</b>        | <p>Lists messages that are generated by the NetWorker server, including the priority of each message, the time the message was generated, the source of the message, and the category. <a href="#">Log window</a> on page 90 provides more information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Customizing the Monitoring window

This section describes how to customize the **Monitoring** window in the **Administration** interface.

### Customizing tables

You can customize the organization and display of tabular information in the **Monitoring** window.

#### Sorting tables

You can change the display of tabular information that appears in the window. You can sort Table grids by column heading, and then by alphabetic or numeric order within those columns.

1. Drag and drop the column heading to its new position.
2. Click the column heading to sort the items into alphabetic and numeric order. An arrow appears in the column heading to indicate the sort order.

#### Sorting selected rows in a table

Selected rows are sorted to the top of the table. This sorting is particularly useful when you select **Highlight All** from the Find panel to select all rows matching the Find criteria and then moving all selected rows to the top of the table to view the results.

1. From the **Edit** menu, select **Find**, or press **Ctrl + F** to view the **Find** panel.
2. To select the rows, click each row or use the Find criteria.
3. Select **Sort Selected**.

#### Sorting multiple columns in a table

You can select the column that you want to use as the tertiary sort key, the secondary sort key, and the primary sort key.

1. Click the column that you want to use as the last sort key.
2. Click the column that you want to use as the next-to-last sort key, and so on, until you select the primary column.



### Displaying columns in a table

You can select which columns to display in a table.

1. From the **View** menu, select **Choose Table Columns**.
2. Click a column name to select or clear the column and then click **OK**. You can also select the columns to display by right-clicking a table header and selecting **Add Column** from the drop-down.

## Displaying panes

You can choose to show or hide panes in the **Monitoring** window.

### About this task

Perform the following steps to hide or show a pane in the **Monitoring** window.

### Procedure

1. From the **View** menu, select **Show**. A check mark appears beside the panes that appear in the **Monitoring** window.
2. To hide a pane, select a marked pane.  
A check mark does not appear beside the pane.
3. To show a pane, select an unmarked pane.  
A check mark appears beside the pane.

## Policies/Actions pane

The **Policies/Actions** pane provides you with the ability to review status information about policies and actions.

This pane has two tabs:

- **Policies**—Provides a navigation tree that displays all configured policies on the NetWorker Server. Expand each policy to display the workflows that are associated with each policy. Expand each workflow to display each action that is contained in the workflow.
- **Actions**—Provides a list of all Action resources.

### Policies pane

The **Monitoring** window in the **NetWorker Administration** window enables you to monitor activities for specific policies, workflows, and actions.

The **Policies/Actions** pane at the top of the **Monitoring** window lists the policies on the NetWorker Server by default. Click the + (plus) sign next to a policy in the list to view the workflows in the policy, and the + (plus) sign next to a workflow to view the actions for a workflow.

The **Policies** pane provides the following information for each item (where applicable):

- Overall status

The following table provides details on the status icons that may appear in the **Policies** pane.

**Table 61** Policy status icons

| Icon                                                                                | Status    |
|-------------------------------------------------------------------------------------|-----------|
|  | Never run |

**Table 61** Policy status icons (continued)

| Icon                                                                              | Status                                                |
|-----------------------------------------------------------------------------------|-------------------------------------------------------|
|  | Running                                               |
|  | Succeeded                                             |
|  | Failed                                                |
|  | Probing                                               |
|  | Interrupted                                           |
|  | Queued                                                |
|  | Cloning                                               |
|  | Consolidating (NetWorker Server 8.2.x and lower only) |

 **Note:** When the schedule for an action is skip, the status of the action appears as Never Run and the status of the Workflow is Succeeded.

- Most recent start time.
- Duration of the most recent run.
- Next scheduled runtime.
- Name of the assigned save set.
- Device on which the save set is stored.
- Backup level.
- Data transfer rate.
- Size of the save set.
- Messages that resulted from an action.

Right-click an action in the **Policies** pane and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

When you sort the items on the **Policies/Actions** pane by using the **Status** column, NetWorker sorts the items in alphabetical order that is based on the label of the icon.

Consider the following when a policy/action is in a probing state:

- A message is sent when the group starts and finishes the probe operation.
- The results of the probe operation (run backup/do not run backup) are also logged.
- Probes do not affect the final status of the group, and the group status does not indicate the results of the probe.
- If probing indicates that a backup should not run, then the group status reverts to its state before the group running.
- Check the results of the probe in the **Log** window to ensure that the probe indicates that the backup can be taken.

## Actions pane

To view a list of all actions, click the **Actions** tab at the bottom of the **Policies** pane. The **Policies** pane becomes the **Actions** pane.

The **Actions** pane provides the following information for each action:

- Overall status

 **Note:** The **Actions** pane displays the same status icons as the **Policies** pane.

- Name
- Assigned policy
- Assigned workflow
- Type
- Date and time of the most recent run
- Duration of the most recent run
- Percent complete, for actions that are in progress
- Next scheduled runtime

Right-click an action in the **Actions** pane and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

## Workflow operations

This section describes how to use the **Monitoring** window to start, stop, and restart workflows.

### Starting, stopping, and restarting policies

The workflows in a policy can run automatically, based on a schedule. You can also manually start, stop, and restart specific workflows by using the the NMC **NetWorker Administration Monitoring** window.

#### About this task

You can restart any failed or canceled workflow. Note, however, that the restart must occur within the restart window that you specified for the workflow. Additionally, for a VMware backup, if you cancel a workflow from **NetWorker Administration** and then want to restart the backup, ensure that you restart the workflow from the **NetWorker Administration** window. If a workflow that was started from **NetWorker Administration** is restarted from the **vSphere Web Client**, the backup fails.

#### Procedure

1. In the **Monitoring** window, select the workflow or actions.
2. Right-click and then select **Start**, **Stop**, or **Restart**.

A confirmation message appears.

 **Note:** You cannot stop, restart, or start individual actions.

3. Click **Yes**.

### Viewing workflow backup details

Perform the following steps to view backup details for workflows.

#### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Policies** in the docking panel, and expand the Policy that you want to monitor.

3. Right-click the workflow, and then select **Show Details**. The **Workflow Summary** window appears.
4. In the **Workflow runs** pane of the **Workflow Summary** window, select the workflow.
5. Click **Show Messages**. In the **Show Messages** window, select one of the following options:
  - Get Full Log—To display all messages.
  - Print—To print the log.
  - Save—To save the log to a local file.
  - OK—To close the **Show Messages** window.
6. Click **OK** to close the **Workflow Summary** window.

## Viewing action backup details

Perform the following steps to view backup details for actions.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Actions** in the docking panel.
3. In the **Actions** pane, right-click the action, and then select **Show Details**. The details window for the action appears.
4. Review the information in the **Actions Messages** pane. To display detailed information from the action log file, click **Show Action Logs**, and then select one of the following options:
  - Get Full Log—To display all messages.
  - Print—To print the log.
  - Save—To save the log to a local file.
  - OK—To close the **Show Messages** window.
5. In one of the Actions detail panes, for example, the **Completed successfully** pane, select the action that you want to review.
6. Click **Show Messages**. In the **Show Messages** window, select one of the following options:
  - Get Full Log—To display all messages.
  - Print—To print the log.
  - Save—To save the log to a local file.
  - OK—To close the **Show Messages** window.
7. Click **OK** to close the **Details** window.

## Sessions window

Use the **Sessions** window to view the sessions that are running on a NetWorker server. You can change the view of this window to display these sessions:

The **Sessions** pane below the **Policies/Actions** pane provides details on individual save, recover, clone, and synthetic full sessions by client.

To view all sessions or to limit the list of sessions by the session type, click the tabs at the bottom of the **Sessions** pane. Session types include:

- Save
- Recover

- Clone
- Browse
- Synthetic Full/Rehydrated Sessions
- All

To change the displayed session types go to **View > Show**, and select the type of sessions to display. To display all sessions currently running on the NetWorker Server, regardless of type, select **All Sessions**.

You can stop a session (backup, synthetic full backup, clone, and recovery sessions) from the **Monitoring** window, even if the session was started by running the `nsrworkflow` command.

To stop a session, right-click the workflow in the pane, and select **Stop** from the list box.

## Changing displayed session types

The column headings that are displayed on this window differ depending on the type of sessions you chose to display.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Sessions** in the docking panel.
3. Select **View > Show** and then select the type of sessions to display. To display all sessions currently running on the NetWorker server, regardless of type, select **All Sessions**.

## Stopping a session

You can stop a session (backup, synthetic full backup, clone, and recovery sessions) from the Monitoring window, even if the session was started by running `savegrp`.

To stop a session, right-click the session in the window and select Stop from the drop-down.

The following table provides a list of actions that can be stopped from NMC.

**Table 62** Sessions that can be stopped from NMC

| Session type                             | Stop from NMC? |
|------------------------------------------|----------------|
| Save by Savegroup                        | Yes            |
| Synthetic Full by Savegroup              | Yes            |
| Clone by Savegroup                       | Yes            |
| Schedule Clone                           | Yes            |
| Manual Save                              | No             |
| Manual Clone via NMC                     | No             |
| Manual Clone via CLI                     | No             |
| Winworker and CLI Recovery               | No             |
| Recovery started from Recover wizard     | Yes            |
| VMware Backup Appliance Save and Recover | No             |

 **NOTICE** Stopping a session from NMC does not affect any other group operations running.

## Alerts pane

The **Alerts** pane displays alerts that are generated by a particular NetWorker server or Data Domain system that has devices that are configured on the NetWorker server. The **Alerts** pane includes priority, category, time, and message information.

An icon represents the priority of the alert. The following table lists and describes each icon.

**Table 63** Alerts window icons

| Icon                                                                                | Label        | Description                                                                                                                           |
|-------------------------------------------------------------------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
|    | Alert        | Error condition detected by the NetWorker server that should be fixed by a qualified operator.                                        |
|    | Critical     | Severe error condition that demands immediate attention.                                                                              |
|    | Emergency    | Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority. |
|    | Information  | Information about the current state of the server. This icon represents the lowest priority.                                          |
|    | Notification | Important information.                                                                                                                |
|   | Waiting      | The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.                                           |
|  | Warning      | A non-fatal error has occurred.                                                                                                       |

When items on the **Alerts** pane are sorted by the **Priority** column, they are sorted in alphabetical order based on the label of the icon.

## Removing alerts

Remove individual alert messages from the **Events** tables by removing them from the **Events** table. To delete a message in the **Events** table, right-click the message, and select **Dismiss**.

### About this task

 **Note:** The alert message remains in the **Log** window in the NetWorker **Administration** program.

## Devices pane

The **Devices** pane allows you to monitor the status of all devices, including NDMP devices. If the NetWorker server uses shared and logical devices, the window is adjusted dynamically to present a set of columns appropriate for the current configuration.

The **Devices** pane provides the following information:

- Status of the operation.
- Name of the device.
- Name of the storage node that contains the device.
- For tape devices, the name of the library that contains the device.

- Name of the volume in the device.
- Name of the pool that is associated with the volume.
- Last message generated for the device.
- Whether the operation requires user input.

For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape.

[Entering user input](#) on page 89 provides instructions on how to deal with a user input notification.

If the current server configuration includes a shared device, a **Shared Device Name** column appears on the **Devices** pane. The name of the shared device appears in the **Shared Device Name** column. If other devices for that configuration are not shared devices, then the **Shared Device Name** column is blank for those devices. Only a single device per hardware ID can be active at any particular moment. The information for inactive shared devices is filtered out, and as a result, only one device per hardware ID is presented on the window at any time.

An icon represents the device status. The following table lists and describes each icon.

**Table 64** Devices status icons

| Icon                                                                                | Label                       | Description                         |
|-------------------------------------------------------------------------------------|-----------------------------|-------------------------------------|
|    | Library device active       | The library device is active.       |
|   | Library device disabled     | The library device is disabled.     |
|  | Library device idle         | The library device is idle.         |
|  | Stand-alone device active   | The stand-alone device is active.   |
|  | Stand-alone device disabled | The stand-alone device is disabled. |
|  | Stand-alone device idle     | The stand-alone device is idle.     |

When you sort items in the **Devices** pane by the **Status** column, NetWorker sorts the devices in alphabetical order based on the label name of the icon.

## Operations window

The **Operations** window displays information about device operations. It provides the following information:

- Status of the operation.
- Name of the library.
- Whether the operation requires user input.  
For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape. [Entering user input](#) on page 89 provides instructions on how to deal with a user input notification.
- The origin, or source, of the operation.  
For example, the interface, nsrjb or the NetWorker server.
- Time the operation started.
- Type of operation.

- Duration of the operation.
- Status messages from the operation.
- Any error messages.

 **NOTICE** Only the last error message of the operation appears in the **Error Messages** column. Move the mouse pointer over the cell containing the last error message to display the entire list of error messages.

The operation status is represented by an icon. The following table lists and describes each of the icons.

**Table 65** Operations window icons

| Icon                                                                               | Label      | Description                                          |
|------------------------------------------------------------------------------------|------------|------------------------------------------------------|
|   | Failed     | The operation failed.                                |
|   | Queued     | The operation is waiting in the queue to run.        |
|   | Retry      | The operation failed, but may work if you try again. |
|   | Running    | The operation is running.                            |
|   | Successful | The operation completed successfully.                |
|  | User Input | The operation requires user input.                   |

When items on the **Operations** window are sorted by the Status column, they are sorted in alphabetical order based on the label of the icon.

## Viewing operation details

The **Operation Details** dialog box opens, providing information about the completion of the operation. The **Completion Time** displays the time that the operation finished. The time that it took to complete the operation is the difference between the completion and start times of the operation.

### About this task

To save operation details to a file, click **Save** in the **Operation Details** dialog box. When prompted, identify a name and location for the file.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation, then select **Show Details**.

## Stopping an operation

Certain operations can be stopped from the **Operations** window.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation to stop, then select **Stop**.



- Click **Yes** to confirm the stop.

 **Note:** Operations that were started from a command line program, such as the `nsrjb` command, cannot be stopped from the **Operations** window. To stop these operations, press `Ctrl-C` from the window where the command was started.

## Entering user input

If the system requires user input, select the labeling operation in slow/verbose mode and the **Supply User Input** icon appears.

### Procedure

- Right-click the operation, then select **Supply Input**.
  - Confirm the requirement to supply input.
    - If **Yes**, and input is supplied, the icon in the **User Input** column disappears.
-  **Note:** If two users try to respond to the same user input prompt, the input of the first user takes precedence, and the second user receives an error message.
- If **No**, and input is not supplied, the operation will time out and fail.

## Log window

To view the most recent notification logs, click the **Log** window from the docking panel in the **Monitoring** window. The **Log** window provides the priority, time, source, category, and message for each log.

 **Note:** If a particular log file is no longer available, check the log file on the NetWorker server. The log files are located in `NetWorker_install_path\logs` directory.

An icon represents the priority of the log entry. The following table lists and describes each icon.

**Table 66** Icons in the Log pane

| Icon                                                                                | Label        | Description                                                                                                                           |
|-------------------------------------------------------------------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
|  | Alert        | Error condition that is detected by the NetWorker server that should be fixed by a qualified operator.                                |
|  | Critical     | Severe error condition that demands immediate attention.                                                                              |
|  | Emergency    | Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority. |
|  | Information  | Information about the current state of the server. This icon represents the lowest priority.                                          |
|  | Notification | Important information.                                                                                                                |
|  | Waiting      | The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.                                           |
|  | Warning      | Non-fatal error has occurred.                                                                                                         |

When you sort items on the **Log** pane by using the **Priority** column, NetWorker sorts the icons in alphabetical order based on the name of the label.

## Recover window

The **Recover** window displays information about recover configurations that are created with the NetWorker Management Console (NMC) Recovery wizard.

You can use this window to:

- Start the NMC Recovery wizard to create recover configurations or modify saved recover configurations.
- Identify the status of a recover configuration that is created with the NMC Recovery wizard.
- Start and stop a recover job.

The **Recover** window is divided into five sections:

- **Toolbar**—The toolbar is hidden by default. To display the recovery toolbar, select **View > Show toolbar**
- **Summary**
- **Configured Recovers**
- **Currently Running**

A splitter separates the **Configured Recovers** section from **Currently running** window. You can click and move the splitter to resize these two windows.

## Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. The following table summarizes the function of each toolbar button.

**Table 67** Recovery toolbar options

| Button                                                                              | Function                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Starts the NMC <b>Recover</b> wizard to create recover configurations.                                                                                                                        |
|  | Displays the <b>Properties</b> window for the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                          |
|  | Deletes the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                                                            |
|  | Displays online help for the <b>Recover</b> window.                                                                                                                                           |
|  | Displays the <b>Find</b> window at the bottom of the <b>Recover</b> window. The <b>Find</b> window allows you to perform keyword searches for messages that appear in the <b>Logs</b> window. |
|  | Start the recover operation for a selected saved recover configuration. This option is only available for a recover configuration that has a Never run, or Failed status.                     |
|  | Stop in-progress recover operation that you selected in the <b>Currently Running</b> window.                                                                                                  |

 **Note:** The **Recover** toolbar does not appear by default. To display the **Recover** toolbar, select **View > Show toolbar**.

## Recover Summary

The Recover Summary section displays a high-level overview of recover jobs.

This section includes the following information:

- Total Recovers—The total number of successful recover jobs.
- Since—The number of successful recover jobs since this date.

## Configured Recovers

The **Configured Recovers** window displays a list of saved recover configurations in a tabular format. You can sort the information by column. The **Configured Recovers** table displays the following information for each saved recover configuration:

- Status—The job status of a saved recover configuration.
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB.
- Comment
- OS—The operating system of the source host.
- Recover requestor—The Windows or UNIX account used to create the recover configuration.
- Start Time
- End Time
- Start date

**Table 68** Save recover configuration job status

| Icon                                                                                | Description                                        |
|-------------------------------------------------------------------------------------|----------------------------------------------------|
|  | The last recover attempt failed.                   |
|  | The last recover attempt completed successfully.   |
|  | The recover job has never run.                     |
|  | The recover job is scheduled to run in the future. |
|  | The recover job has expired.                       |

## Currently running

The **Currently Running** window displays a list of in progress recover jobs in a tabular format. You can sort the information by column. The **Currently Running** table displays the following information for each job:

- Status

- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB
- Volume
- Comment
- Device
- Size
- Total size
- % complete
- Rate (KB/s)
- Start time
- Duration
- Currently running

## Find

The **Find** section appears along the bottom of the **Recover** window, after you select the **Find** button on the **Recover** toolbar. **Find** allows you to search for keywords in the **Configured Recovers** window. The following table summarizes the available find options.

**Table 69** Find options

| Find option          | Description                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Find</b>          | Highlight the first saved recover configuration that contains the specified keyword.                                                         |
| <b>Prev</b>          | Highlight the previous saved recover configuration that contains the specified keyword.                                                      |
| <b>Highlight All</b> | Highlights each saved recover configuration that contains the specified keyword.                                                             |
| <b>Sort Selected</b> | Sorts each highlighted recover configuration in the Configured Recover table so that they appear at the top of the Configured Recover table. |
| <b>Match case</b>    | Make the keyword search case sensitive.                                                                                                      |

## Monitoring checkpoint-enabled backups

You can view detailed information about a checkpoint-enabled backup.

### About this task

#### Procedure

1. From the **Administration** window, select **Monitoring > Groups**.
2. Right-click the group to which the checkpoint enabled client belongs, then select **Show Details**.

3. View the detailed information related to the group backups:
  - If the partial save set is in the work list for the group, the save set appears in the **Waiting to Run** section.
  - If the partial save set is running, the save set appears in the **Currently Running** section.
  - If the entire partial save sets sequence of the savegroup is complete, the save set appears in the **Completed Successfully** section.
  - If the entire partial save sets sequence of the savegroup is not complete, the save set appears in the **Failed** section.
-  **NOTICE** If any messages are generated, the Show Messages button is enabled. Click Show Messages to view the messages.
4. Click **OK** to close the **Group Details** window.

## Query the media database for partial save sets

The savegrp completion report does not provide detailed information about partial save sets that might be necessary to perform a recovery.

### Querying partial save sets from the Console

You can view information about the partial save sets by using the NetWorker Console.

#### About this task

-  **NOTICE** If no partial save sets are found that match the query, ensure that the backup of the partial save sets was started within the Save Time period. To change the values for the Save Time attribute, open the Save Set Query tab and select a date and time from the Save Time calendar.

#### Procedure

1. In the **Administration** window, click **Media**. Media-related topics appear in the navigation tree.
2. Select **Save Sets**. The following tabs appear in the **Save Sets** window:
  - Query Save Set
  - Save Set List
3. Select the **Query Save Set** tab, to query:
  - All partial save sets, select **Checkpoint Enabled**.
  - All partial save sets with the same Checkpoint ID, in the **Checkpoint ID** field, type the **Checkpoint ID** of the partial save set on which you want to perform the query.
4. Select the **Save Set List** tab to view the result of the save set query:
  - The **Checkpoint ID** column displays the partial save set **Checkpoint ID** and its Sequence ID. The **Checkpoint ID** is listed first followed by the **Sequence ID**, which is encased within brackets.
  - Sort the **Checkpoint ID** column to view the complete sequence of partial save sets.
  - The **Status** column displays the status of the partial save sets:
    - A Checkpoint browsable status indicates that the save sets can be browsed for recover.
    - A Checkpoint aborted status indicates that the backup of the partial save set was stopped or aborted. A save set recover is used to recover the partial save set.

Consider the following:

- When a checkpoint-enabled backup completes successfully, the status of the last partial save set is Checkpoint browsable.
- When a checkpoint-enabled backup completes successfully, on the first backup attempt, the save set status is Checkpoint browsable. Only one Sequence id is associated with the Checkpoint ID. The Sequence id is 1. If the Sequence id is 2, the first partial save set in the checkpoint-enabled backup is missing.

## Querying partial save sets by using the `mminfo` command

By default, the `mminfo` command output only displays the browsable save sets. The first and intermediate partial save sets are not displayed. Only complete checkpoint-enabled save sets or final partial save sets are displayed.

Use the `mminfo` command with specific queries to display more information about checkpoint-enabled save sets.

The following table lists the new media attributes support the Checkpoint Restart feature.

**Table 70** New Checkpoint Restart media attributes

| Media attribute    | Description                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------|
| checkpoint_id      | Displays the checkpoint restart id of the partial save set in the <code>chkpt_id</code> column. |
| checkpoint_seq     | Displays the partial save set sequence id in the <code>chkpt_seq</code> column.                 |
| checkpoint-restart | This flag attribute is used to only display checkpoint restart enabled save sets.               |

Also, several media sumflags are used with the Checkpoint Restart feature:

- `k`—Indicates this is a checkpoint enabled save set.
- `a`—The first and all intermediate partial save sets of a checkpoint sequence has aborted status.
- `b`—The last partial or complete save set of a checkpoint sequence is marked browsable.

## Displaying checkpoint enabled save sets

Display all checkpoint enabled save sets by using the following `mminfo` command:

```
mminfo -q 'checkpoint-restart' -r 'client,nsavetime,ssid(11),
sumflags(3),name,checkpoint_id,checkpoint_seq'
```

**Table 71** Checkpoint enabled save sets

| client | save time  | ssid           | ssflags | filename | chkpt_id   | chkpt_seq |
|--------|------------|----------------|---------|----------|------------|-----------|
| plapew | 1251910303 | 420470031<br>9 | cak     | /space   | 1251910303 | 1         |
| plapew | 1251910327 | 4187923127     | cbk     | /space   | 1251910303 | 2         |
| plapew | 1251910710 | 408726021<br>4 | cak     | /space   | 1251910710 | 1         |
| plapew | 1251910725 | 407048301<br>3 | cbk     | /space   | 1251910710 | 2         |

## Displaying all partial save sets for the checkpoint id

Display all partial savesets for the checkpoint id by using the following `mminfo` command:

```
mminfo -q "checkpoint_id=1251910303"
```

**Table 72** Partial save sets for the checkpoint id

| volume     | client | date     | size   | level | name   |
|------------|--------|----------|--------|-------|--------|
| plapew.001 | plapew | 08/02/15 | 17 MB  | full  | /space |
| plapew.001 | plapew | 08/02/15 | 799 MB | full  | /space |

## Reporting NDMP Data

The NetWorker software reports information about NDMP clients, data, and volumes in two ways:

- The NetWorker Management Console (NMC) reporting feature—Reports NDMP data in the same manner as non-NDMP data. The *NetWorker Administration Guide* provides more information.
- The `mminfo` command. Use the `mminfo` program to query the media database for NDMP volume or save set information.

## Querying the NDMP volumes by backup type with the mminfo command

You can query save sets by backup format (NDMP or DSA) to display volume information.

For example:

- To query NDMP volumes, type `mminfo -q ndmp`. Output similar to the following appears:

```
volume client date size level name
005D0000 simlcifs1 6/22/2011 1036 MB full /fs1
005D0001 simlcifs1 6/22/2011 173 MB full /fs1
005D0001 simlcifs1 6/22/2011 862 MB full /fs1
005D0002 simlcifs1 6/22/2011 348 MB full /fs1
```

- To query NDMP -DSA volumes, type `mminfo -q dsa`. Output similar to the following appears:

```
volume client date size level name
NDMP.001 10.8.67.219 12/13/2011 644 MB full /vol/vol0
NDMP.001 10.8.67.219 12/13/2011 402 MB full /vol/vol1
NDMP.001 10.8.67.219 12/13/2011 402 MB full /vol/vol1
NDMP.001 10.8.67.219 12/13/2011 36 MB full /vol/vol2
```

## Querying the NDMP save sets with the mminfo command

To determine which save sets are Network Data Management Protocol (NDMP) save sets and the status of an NDMP save set in the media database, query the media database. NDMP save set status information is important when performing NDMP recoveries:

- To perform a browsable NDMP recover, the `ssflags (f1)` field for an NDMP save set must contain a (b). The b value denotes a browsable save set.
- To perform a save set recover from the NetWorker User program, the `ssflags (f1)` field for an NDMP save set must contain either (r) or (b).

- An NDMP save set contains an N attribute in the ssflags (f1) field.
- An NDMP-DSA save set contains an s attribute in the ssflags (f1) field.

In the following example, the NDMP save set status is recoverable (r). To recover the data, you can only perform a save set recovery from a command line.

```
mminfo -av
```

```
volume type client date time size ssid fl lvl name
vol1 dlt clnt 6/22/2011 3:15:12 1036MB 3842140553 hrN full /fs1
```

In the following example, the NDMP-DSA save set status is browsable (b). Recover the data by using the NetWorker User program, or from the command line. A browsable NDMP-DSA save set supports browsable and save set recoveries.

```
mminfo -av
```

```
volume type client date time size ssid fl lvl name
vol1 dlt clnt 6/22/2011 3:15:12 36MB 4259813785 cbNs full /fs1
```

## Performing NDMP recoveries

NetWorker uses the `nsrndmp_recover` program to coordinate recover operations between the NetWorker software and the Network Data Management Protocol (NDMP) client. The `nsrndmp_recover` program does not move data to the NDMP client. When the `nsrndmp_recover` program identifies an NDMP-DSA save set, `nsrndmp_recover` automatically runs the `nsrdsa_recover` program on the same host that runs the `nsrndmp_recover` command.

To recover NDMP data, you can run the `nsrndmp_recover` program from a command prompt, or use one of following programs, which automatically starts `nsrndmp_recover`:

- `recover`—The command line program on Windows and UNIX.
- `winworkr`—The NetWorker User GUI on Windows.
- The NMC Recovery wizard.

During the recovery process, the `nsrndmp_recover` program passes nlist information to the NDMP client. There are three methods to recover NDMP backups:

- Index-based file-by-file recover—The nlist includes file offset and ACL information. When you recover many files, the recover process uses a significant amount of system resources on both the NetWorker server and the NDMP client to build and process the nlist information.
- Full save set recovery—The nlist only includes the path to the recovery directory, down to and including the mount point. When you recover many files, the recover process uses less system resources than an index-based NDMP recover to build and process the nlist information.
- NDMP directory restore—A partial save set recovery of a single file or single directory.

For example, when the NetWorker software writes NDMP data to a remote storage node, start the `recover` program on the NetWorker storage node to prevent the data from traversing the network.

 **Note:** When you start the `recover` program on the NetWorker server, the data flows from the storage node to the NetWorker server and from the NetWorker server to the NDMP client, over the network.

## NDMP recovery requirements for NetApp

The following list summarizes the requirements:



### scanner

You cannot use the `scanner` command with the `-i`, `-f` and `-r` options on an NDMP volume. You cannot use the `scanner` command on a volume that contains NDMP and non-NDMP save sets when you load the volume in an NDMP device. The *Scanner command usage* technical note provides more information about using the `scanner` command with NDMP data.

### Cross platform recoveries

You can recover NDMP data to different NDMP client however, you cannot perform a cross platform recover. Recover NDMP data to an NDMP client that is the same brand, a compatible model, and the same operating system as the original NDMP client.

### Devices

Recover Direct-NDMP and Three-party backups performed to an NDMP device from an NDMP device. To improve recover performance from an NDMP tape device, configure the tape device to support variable length records. Recover NDMP-DSA backups from a non-NDMP device.

### Localized environments

When recovering data in a localized NDMP environment, the Index Recover status window shows the process in English and not the localized language.

### NDMP-DSA

For better recovery performance, start the recover process on the NetWorker host where the backup volume resides.

### Immediate recoveries

Run the `nsrndmp_recover` program on the storage node with the locally attached backup device to perform an immediate recovery of NDMP-DSA data.

**Note:** When you select to restore multiple files, if the files are on different save sets, NetWorker sends multiple NDMP recover requests to the NAS client. With NetApp 8.2 and later releases, the ability to run parallel recoveries if the files are restored to the same destination volume is disabled, and the restore fails. You can use the `NDMP_RECOVER_LIMIT` parameter to configure the number of sessions on the NetWorker server. For NetApp 8.2 and later releases, set the parameter as follows:

```
NDMP_RECOVER_LIMIT=1
```

## DAR and DDAR

By default, the Network Data Management Protocol (NDMP) recover process reads an entire tape from start to finish. The recover process extracts the data as it encounters the data on the tape. For large backup images, recovery is slow.

The Direct Access Recovery (DAR) and Directory DAR (DDAR) recovery process:

- Provides the ability to recover a file or directory from the exact location on a tape.
- DDAR only passes the directory path to the NAS filer. DAR passes the paths of each file individually.
- Reduces the size of the nlist information that the recover process stores in memory. During the recover process, the NAS filer (DDAR) assumes that the directory path includes all cataloged files and directories. However, DAR mentions each file that it wants recovered.
- Does not sequentially read the file or record numbers on the tape to locate the data, which reduces the amount of time that you require to recover specific files from a backup.

 **Note:** [Creating and configuring the NDMP client resource](#) on page 67 describes how to configure the DAR and DDAR Application Information attributes for NDMP clients.

### When not to use DAR or DDAR

DAR and DDAR recoveries send multiple pathnames across the network to the NDMP Data Server and, in three-party configurations, to the NetWorker server. The recover process stores the pathnames in memory on the NDMP Data Server. Recoveries of a large amount of data from a large save set can negatively affect the network and the NDMP Data Server resources.

Do not use DAR and DDAR to recover the following objects:

- Several thousands of files in a single index-based recover operation.
- A specific directory structure containing several thousand or millions of files.

To perform a non-DAR-based recovery of a save set when you set the *DIRECT=y* at the time of backup, first define the *NSR\_NDMP\_RECOVER\_NO\_DAR=y* variable in the Application Information attribute of the NDMP client.

## Recovering data from partial save sets

If there is a complete sequence of partial save sets that span the original save set, then you can browse to and recover individual files and directories. If the sequence of partial save sets is incomplete and does not make up the original save set, then you must perform a save set recovery to recover the data from the partial save set.

To recover data from partial save sets that span the original save sets, perform a query for all partial save sets, and then use either the NetWorker User program on Windows or the `recover` program on UNIX to restore the data.

The steps to recover data from a single partial save set are the same as save set recovery from a complete save set. The partial save set contains only files that were successfully backed up. You cannot browse partial save sets.

When you perform a save set recovery of a partial NDMP save set, the recovery process recovers all partial save sets in the checkpoint sequence. You cannot recover data in a partial save set separately from other partial save sets in the checkpoint sequence.

Use the `nsrinfo` command to display the contents of a partial save set.

## Recover window

The **Recover** window displays information about recover configurations that are created with the NetWorker Management Console (NMC) Recovery wizard.

You can use this window to:

- Start the NMC Recovery wizard to create recover configurations or modify saved recover configurations.
- Identify the status of a recover configuration that is created with the NMC Recovery wizard.
- Start and stop a recover job.

The **Recover** window is divided into five sections:

- **Toolbar**—The toolbar is hidden by default. To display the recovery toolbar, select **View > Show toolbar**
- **Summary**
- **Configured Recovers**
- **Currently Running**

A splitter separates the **Configured Recovers** section from **Currently running** window. You can click and move the splitter to resize these two windows.

## Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. The following table summarizes the function of each toolbar button.

**Table 73** Recovery toolbar options

| Button                                                                              | Function                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|    | Starts the NMC <b>Recover</b> wizard to create recover configurations.                                                                                                                        |
|    | Displays the <b>Properties</b> window for the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                          |
|    | Deletes the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                                                            |
|    | Displays online help for the <b>Recover</b> window.                                                                                                                                           |
|   | Displays the <b>Find</b> window at the bottom of the <b>Recover</b> window. The <b>Find</b> window allows you to perform keyword searches for messages that appear in the <b>Logs</b> window. |
|  | Start the recover operation for a selected saved recover configuration. This option is only available for a recover configuration that has a Never run, or Failed status.                     |
|  | Stop in-progress recover operation that you selected in the <b>Currently Running</b> window.                                                                                                  |

 **Note:** The **Recover** toolbar does not appear by default. To display the **Recover** toolbar, select **View > Show toolbar**.

## Recover Summary

The Recover Summary section displays a high-level overview of recover jobs.

This section includes the following information:

- **Total Recovers**—The total number of successful recover jobs.
- **Since**—The number of successful recover jobs since this date.

## Configured Recovers

The **Configured Recovers** window displays a list of saved recover configurations in a tabular format. You can sort the information by column. The **Configured Recovers** table displays the following information for each saved recover configuration:

- **Status**—The job status of a saved recover configuration.
- **Name**
- **Source client**
- **Destination client**

- Recovery list
- Recover type—For example, file system or BBB.
- Comment
- OS—The operating system of the source host.
- Recover requestor—The Windows or UNIX account used to create the recover configuration.
- Start Time
- End Time
- Start date

**Table 74** Save recover configuration job status

| Icon                                                                               | Description                                        |
|------------------------------------------------------------------------------------|----------------------------------------------------|
|   | The last recover attempt failed.                   |
|   | The last recover attempt completed successfully.   |
|   | The recover job has never run.                     |
|   | The recover job is scheduled to run in the future. |
|  | The recover job has expired.                       |

## Currently running

The **Currently Running** window displays a list of in progress recover jobs in a tabular format. You can sort the information by column. The **Currently Running** table displays the following information for each job:

- Status
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB
- Volume
- Comment
- Device
- Size
- Total size
- % complete
- Rate (KB/s)
- Start time
- Duration

- Currently running

## Find

The **Find** section appears along the bottom of the **Recover** window, after you select the **Find** button on the **Recover** toolbar. **Find** allows you to search for keywords in the **Configured Recovers** window. The following table summarizes the available find options.

**Table 75** Find options

| Find option          | Description                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Find</b>          | Highlight the first saved recover configuration that contains the specified keyword.                                                         |
| <b>Prev</b>          | Highlight the previous saved recover configuration that contains the specified keyword.                                                      |
| <b>Highlight All</b> | Highlights each saved recover configuration that contains the specified keyword.                                                             |
| <b>Sort Selected</b> | Sorts each highlighted recover configuration in the Configured Recover table so that they appear at the top of the Configured Recover table. |
| <b>Match case</b>    | Make the keyword search case sensitive.                                                                                                      |

## Performing an NDMP index-based file-by-file data recovery

Perform an NDMP index based file-by-file recover in the same manner as a non-NDMP data recover. You can restore the data to the original NDMP client or perform a directed recovery to a different NDMP client.

Before you perform an index-based file-by-file recover, review the following information:

- Set the *HIST=y* in the Application Information attribute of the NDMP client at the time of the backup.
- The NDMP save set must be browsable. You cannot perform a browsable recover of a recoverable or recyclable save set. [Reporting NDMP Data](#) on page 93 describes how to determine the status of an NDMP save set.
- Do not use an index-based recovery to recover a large numbers of files or directories. For better recovery performance, use a save set recover. [Performing a Full or Directory Restore of NDMP data by using a save set recovery](#) on page 103 provides more information.
- To perform an index-based file-by-file recover:
  - Use the NetWorker User program on a Windows host. [Performing an NDMP index-based file-by-file recover using the NetWorker User program](#) on page 100 provides detailed information.
  - Use the `recover` program. [Performing an NDMP index-based file-by-file recover from a command prompt](#) on page 102 provides detailed information.

### Performing an NDMP index-based file-by-file recover using the NetWorker User program

On Windows, to recover data to the original NDMP client or to a different NDMP client, perform the following steps.

#### Procedure

1. Open the NetWorker User program and connect to the NetWorker server.

**NOTICE** If you receive the error:

No file indexes were found for client `client_name` on server `server_name`

Try connecting to a different NetWorker server and you selected the correct NetWorker server, then ensure that you selected a browsable save set. Alternatively, perform a save set recover.

2. Select **Recover** to open the **Source Client** window.
3. Select source NDMP client and click **OK**. The local client is the default selection.
4. Select the destination client for the recovered data and click **OK**. If the destination client is not the source client, ensure the NAS filer is the same brand, a compatible model and the same operating system as the source NDMP client.
5. (Optional) Recover the data from an earlier backup time. The **Recover** window appears with the latest version of the backup files. To recover data from an earlier backup, change the date and time of backup using one of the following methods:
  - a. Change the browse time for all files in the recover window:
    - From the **View** menu, select **Change Browse Time**.
    - In the **Change Browse Time** window, select a new day within the calendar. Select **Previous Month** or **Next Month** to change from the current month.
    - In the **Time** field, change the time of day by typing an hour, a minute, and the letter a for A.M. or p for P.M. Use the 12-hour format.
    - Click **OK**.
  - b. View all versions of the selected file system object:
    - Highlight the file or directory for review.
    - From the **View** menu select **Versions**.
    - Once you locate the version to recover, change the browse time. To change the browse time, highlight the volume, directory, or file and click **Change Browse Time**. The **Version** window closes and the **Recover** window reflects the new browse time.
6. (Optional) Search for the files. To search for and recover the most recently backed-up version of a file or directory:
  - a. From the **File** menu, select **Find**.
  - b. Type the name of the file or directory. Use wildcards to expand the search; without wildcards, partial filenames do not provide any results.
7. Mark the data to recover. To select file system objects to recover:
  - a. In the left pane of the **Recover** window, click the appropriate directory folder.
  - b. Mark each directory or file to recover by selecting the checkbox next to each directory or file.
8. (Optional) Relocate the data to a different location. By default, the recover process recovers the selected files to the original location.

**Note:** The NDMP protocol does not support name conflict resolutions. NetWorker will always overwrite existing files that have the same name as the recovered file. It is recommended that you recover the NDMP data to a different location, to avoid data loss.

To relocate the files to a different location:

- a. Select **Recover Options** from the **Options** menu.

NDMP recoveries do not support the following options:

- Rename recovered file
- Discard recovered file
- Prompt for every file conflict

NDMP recoveries will always overwrite existing files. It is recommended that you relocate the NDMP data to a different location, to avoid data loss.

- b. In the **Relocate Recovered Data To** field, type the full path name of the target directory, click **OK**.

The target directory is a literal string and must match the path as seen by the NAS filer in its native OS, exactly. Otherwise, the recover process uses the original location and overwrites existing files with the same name.

9. (Optional) To view the volumes required to recover the marked file system objects, from the **View** menu, select **Required Volumes**.
10. Click **Start** to begin the recovery. If any required volume is not available to the NetWorker server, a volume status warning appears.

When this warning appears:

- a. Click **No**.
- b. From the **View** menu, select **Required Volumes**.
- c. Ensure that the NetWorker software can mount each listed volume into an available device.
- d. Attempt the recover operation again.

The NetWorker server takes a few moments to recover the files, depending on file size, network traffic, server load, and tape positioning. During this time, messages appear so that you can monitor the progress of the recovery.

When the recovery completes successfully, a message similar to the following appears:

```
Received 1 file(S) from NSR server
server Recover completion time: Tue Jan 21 08:33:04 2009
```

## Performing an NDMP index-based file-by-file recover from a command prompt

This section applies to command line recoveries from a Windows and UNIX client.

### About this task

To avoid using the Windows version of `recover.exe` on Windows operating systems, perform one of the following actions:

- Specify the full path to the recover program. For example: `C:\Program Files\EMC NetWorker\nsr\bin\recover.exe`
- Ensure that the `$PATH` environment variable contains the `NetWorker_install_path\bin` directory before `%SystemRoot%\System32`

To recover Network Data Management Protocol (NDMP) data from a command prompt on a UNIX or Windows NetWorker host, perform the following steps.

## Procedure

1. From the command prompt, type:

```
recover -s NetWorker_servername -c client_name
```

where:

- *-s NetWorker\_servername* specifies a particular NetWorker server on the network to use when recovering data.

When you do not use the *-s* option, the `recover` program tries to connect to the first computer listed in the servers file. When the servers file does not contain any servers, or lists more than one server, the **Change Server** window appears, and you can select the server.

- *-c client\_name* specifies the source NDMP client.

2. When prompted, type the directory to browse, for example:

```
cd /mydirectory
```

3. Use the `add` command to add the required files or folders to the recover list. The *NetWorker Command Reference Guide* provides a complete list of options for the `recover` command.
4. When restoring NDMP data, it is recommended that you relocate the NDMP data to a different location.

**Note:** The NDMP protocol does not support name conflict resolutions. NetWorker will always overwrite existing files that have the same name as the recovered file. It is recommended that you recover the NDMP data to a different location, to avoid data loss.

- To relocate the data to a different directory, type:

```
relocate destination_directory_name
```

The target pathname for *destination\_directory\_name* is a literal string and must match the path as seen by the NAS filer in its native OS, exactly. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

- To recover the data to a different host, type:

```
relocate target_hostname::/mount_point
```

Data ONTAP may require you to add a backslash (\) after the mount point. For example, *target\_hostname::\mount\_point\*.

5. After you add all of the required files, type:

```
recover
```

## Performing a Full or Directory Restore of NDMP data by using a save set recovery

You perform a Network Data Management Protocol (NDMP) save set recover in the same manner as a non-NDMP save set recovery. You can recover data to the original NDMP client or perform a directed recovery of the data to a different NDMP client of the same platform.

Before you perform a full save set recover, review the following information:



- Use a full save set recovery to recover all files and folders in an NDMP data save set, or to recover an entire directory within an NDMP save set. You cannot use the NetWorker User program to perform an NDMP Directory Restore.
- To use the NetWorker User program on Windows, a client file index entry for the save set must exist. When the index entry for the save set does not exist, the recover fails with an `index not found` error. When the client file index entries do not exist for the save set, use the `nsrndmp_recover` program with the `-v off` option.
- You cannot perform a save set recover from the NetWorker User program when the save set status is eligible for recycling (E). The recover process requires a recoverable (r) or browsable (b) save set status. The *NetWorker Administration Guide* provides information on how to change the status of a save set. A save set recover reads the entire tape, from beginning to end, to find and recover the requested files. The recovery process completes when the recover operations reads all required tapes in their entirety.
- As each file recovers, the file name appears on the target share but the file size is 0 KB. The actual file size update occurs after the recovery completes.

## Performing an NDMP save set recover by using the NetWorker User in Windows

### About this task

**NOTICE** When the recover operations fails with the error:

```
Failed to propagate handle <number> to child process: Access is denied
```

The save set is not in the client file index of the NDMP client. Perform a save set recover from a command prompt. [Performing an NDMP save set recovery from the command prompt](#) on page 105 provides more information.

### Procedure

1. Start the NetWorker User program.
2. On the **Change Server** window, select the NetWorker server and click **OK**.
3. Select **Options > Recover Save Sets**.
4. On the **Source Client** window, select the appropriate NDMP client, and then click **OK**.
5. On the **Save Sets** window, select the name of the save set.
6. Select the version of the save set, if there are multiple versions. You can also select the cloned version of a save set, if applicable.
7. To recover specific files and directories instead of the entire save set:
  - a. Click **Files**.
  - b. Specify the files and directories, one per line.
  - c. Click **OK**.

**NOTICE** Do not use this method to mark tens of thousands of files. Instead, perform an NDMP Directory Restore. Marking many files and directories generates a large nlist and requires intensive resources on both the NetWorker server and the NAS filer.

8. Click **Recover Options**.

An NDMP data recovery does not support the following options:

- Rename recovered file

- Discard recovered file
- Prompt for every file conflict

 **NOTICE** It is recommended that you relocate the NDMP data to a different location. NDMP recoveries always overwrite existing files.

9. To recover the data to a pathname that is different from the original backup location, in the **Relocate Recovered Data To** field, type the full pathname of the destination directory, then click **Ok**.

For NDMP data recoveries, the target pathname is a literal string and must exactly match the path as seen by the native OS on the NAS filer. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

10. To recover the data to a different NDMP client, specify the name of the client to receive the NDMP data in the **Destination Client** field.
11. To view the volumes that are required to perform the recover, select **View > Required Volumes**
12. Click **OK** to begin the recovery. The recovery status appears in the **Recover Status** window.

## Performing an NDMP save set recovery from the command prompt

To perform a save set recovery to the original NDMP client or to a different NDMP client, use the `nsrndmp_recover` command.

### About this task

 **Note:** By default the backup operation uses IPv6 addressing, if enabled. NetApp always uses IPV6 if it is enabled, regardless if IPV6 is configured on the NetApp system or not. If the NetApp system is not configured with IPV6, NDMP backup and recovery displays the following error message:

```
42597:nsrndmp_recover: mover listen:communication failure
42801:nsrndmp_recover: Failed to send the MOVER_LISTEN message to the
NDMP tape server.
```

You can force the NDMP backup and recovery to ignore IPv6 and instead use IPv4 in one of two ways:

1. Add the `-f` option to the `nsrndmp_save` or the `nsrndmp_recover` commands as required.
2. Select the **Disable IPv6** checkbox on the Apps & Modules tab in the **Client Properties** window.

For example:

```
nsrndmp_recover -s NetWorker_server -c source_ndmp_client -S ssid/cloneid -v
off -m target_ndmp_client::/target_path /source_path
```

where:

- `source_ndmp_client` is the hostname of the source NDMP client.
- `target_ndmp_client` is the hostname of the destination NDMP client.
- `/source_path` is the original location of the data.

- */target\_path* is the location to recover the data.
- **NOTICE** It is recommended that you relocate the NDMP data to a different location. NDMP recoveries always overwrite existing files. The */target\_path* is a literal string and must exactly match the path as seen native OS on the NAS filer. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.
- **-v off** allows you to restore data when client file index of the NDMP client does not contain information about the NDMP save set.  
In the following examples, the NetWorker server is mars and the backup client is venus.
  - To recover a mount point */mnt* from a backup of NDMP host venus to a directory */newmnt* on NDMP host jupiter, type: `nsrndmp_recover -s mars -c venus -S 123456789 -v off -m jupiter:./newmnt`
  - To recover a mount point */mnt* from a backup of NDMP host venus to NDMP host pluto, type:  
`nsrndmp_recover -s mars -c venus -R pluto -S 123456789 -v off -m /mnt`  
Data ONTAP may require you to add a slash (/) after the mount point. For example, *target\_hostname:./mount\_point/*.

## Troubleshooting NDMP recover

This section provides a list of the possible causes and the possible resolutions for NDMP recovery issues.

### RESTORE: could not create path *pathname*

This error message appears when restoring NetApp data. This error, when encountered, appears in the `daemon.raw` file of the NetWorker server and the recovery output.

To resolve this issue:

- Ensure that you specify a source and a target path during the recover that exists on the target filer.
- If you set the `UTF8=Y` application information variable during an NDMP client backup and the backup contains path names with non-ASCII characters, then perform a save set recover. Index-based recoveries will fail with this error message.

### These files were not restored (Restore failed with error, or file/directory specified but not found in backup)

This error message appears in the `daemon.raw` file of the NetWorker server and the in the recovery output.

To resolve this issue:

- Ensure that the file or directory specified during the recover, exists in the save set.
- Ensure that the pathname specified to relocate the data exists on the destination filer. For NDMP data recoveries, the target pathname is a literal string and must exactly match the path as seen by the native OS on the NAS filer.



# CHAPTER 5

## Other filers

This chapter includes the following topics:

|                                                                                             |     |
|---------------------------------------------------------------------------------------------|-----|
| • <a href="#">Choosing a device type</a> .....                                              | 270 |
| • <a href="#">Configuring devices for NDMP operations</a> .....                             | 270 |
| • <a href="#">Configure NetWorker for NDMP backup and clone operations</a> .....            | 281 |
| • <a href="#">Monitoring NetWorker Server activities in the Administration window</a> ..... | 316 |
| • <a href="#">Reporting NDMP Data</a> .....                                                 | 333 |
| • <a href="#">Performing NDMP recoveries</a> .....                                          | 334 |

## Choosing a device type

Network Data Management Protocol (NDMP) backups can be written to either an NDMP device, or if using NDMP-DSA, to a non-NDMP device.

Perform either of the following tasks:

- Configure devices for NDMP operations.
- Configure non-NDMP devices. If you are using NDMP-DSA, refer to the *NetWorker Administration Guide* for device configuration.

For a description of each configuration, refer to [Configurations in a NetWorker NDMP environment](#) on page 16.



**Note:**

PowerStore does not support local or direct NDMP backup.

## Configuring devices for NDMP operations

Review this section for information about how to configure the NetWorker environment for Network Data Management Protocol (NDMP) data operations.

The *NetWorker Hardware Compatibility Guide* on the Support website provides a list of NDMP devices that the NetWorker software supports.

### NDMP device limitations

Review these limitations before you configure Network Data Management Protocol (NDMP) devices:

- The timeout of the NetWorker server `nsrmmmd` resource attribute does not apply to NDMP devices, but it does apply to storage nodes devices.
- You cannot use the `jbexercise` utility with an NDMP autochanger.
- You cannot configure NDMP devices on a dedicated storage node.
- You must use a non-rewind device handle for the NDMP media device handle.
- You cannot configure advanced file type devices and file type devices as NDMP devices.
- You cannot configure an NDMP autochanger when the NDMP protocol is earlier than version 3. You must determine the NDMP device handles, then use the `jbconfig` command to configure the autochanger.

### DinoStor-managed jukeboxes

The DinoStor software provides a web-based interface for administering and controlling the TapeServer settings.

Review this information before you use a DinoStor-managed jukebox with NetWorker:

- When you configure the DinoStor TapeServer, set the port number to 10000 on the NDMP page of the **Configure** tab.
- NetWorker supports:
  - SCSI tape devices.
  - GigE and 10/100 Base-T networks.

- You cannot use DDS because the DinoStor TapeServer is not fibre-equipped.

## Determining NDMP device pathnames

To configure an NDMP stand-alone device or an NDMP jukebox, you must first determine the path names of the media devices. If the NAS filer does not support the NDMP\_CONFIG interface or uses NDMP version 3, you must also determine the library device handle.

To determine the NDMP device path names and the library handle, use the `inquire` command or vendor-specific commands.

### Determining the NDMP device path names using the `inquire` command

#### About this task

Use the `inquire` command to determine the path names and library handle.

#### Procedure

1. From a command prompt on the NetWorker server, type:

```
inquire -N NAS_hostname -T
```

2. When prompted, specify the NAS username and password.

**NOTICE** Use the `inquire` command with caution. When you run `inquire`, the command sends the SCSI `inquiry` command to all devices that are detected on the SCSI bus. If you use the `inquire` command during normal operations, unforeseen errors can occur, which might result in data loss.

### Determining the NDMP device pathname for the DinoStor-managed filer

#### About this task

Before you configure an NDMP autochanger, determine the device pathnames of NDMP devices and of the robotic arm.

#### Procedure

1. Access the DinoStor TapeServer interface.
2. Click the **Configure** page.
3. Click the **SCSI** tab.
4. Make note of the device names and device handles.

### Determining the NDMP device pathname for the MiraPoint filer

#### About this task

Before you configure an NDMP autochanger, determine the device pathnames of NDMP devices and of the robotic arm.

#### Procedure

1. On the MiraPoint filer, type:

```
diag tape inquiry ** ** **
```

The device pathname is in the format `/dev/nrst n`, where `n` starts at 0 and increases one number for each tape drive. This value is constant.

When the filer uses NDMP v2 or does not support the NDMP\_CONFIG interface, you must specify the autochanger handle, `/dev/ch0`, when running the `jbconfig` command.

2. To determine the autochanger handle, type:

```
diag changer inquiry ** ** **
```

## Determining the NDMP device pathname for the Procom NetFORCE filer

### About this task

Before you configure an NDMP autochanger, determine the device pathnames of NDMP devices and of the robotic arm.

### Procedure

1. Log in as root and type:

```
status dm
```

2. If the filer uses NDMP v2 or does not support the NDMP\_CONFIG interface, you must determine the autochanger handle.

On a Procom NetFORCE filer, the SCSI device name format is `isp1tSSL[L]`, where `isp1` is the autochanger handle. The Fibre Channel device format is `ffx1tSSL[L]`, where `ffx1` is the autochanger handle.

## Dynamic drive sharing

Dynamic Drive Sharing (DDS) is a feature that provides NetWorker software with the ability to recognize shared physical tape drives. DDS enables NetWorker software to perform the following operations:

- Skip the shared tape drives that are in use.
- Route the backups or recoveries to other available shared tape drives.

## Introduction to DDS

DDS controls application requests for tape media and allows the NetWorker server and all storage nodes to access and share all attached devices.

A system administrator can configure DDS by setting a sharing policy for devices that are accessible from multiple storage nodes.

There are two terms that are central to the use of DDS are drive and device. Within the context of DDS, these terms are defined as follows:

- Drive—The physical backup object, such as a tape drive, disk, or file.
  - Device—The access path to the physical drive.
-  **Note:** NetWorker only supports DDS in a storage area network (SAN) Fibre Channel environment and not in a direct-connect SCSI environment.

### Benefits of DDS

Enabling DDS on a NetWorker system provides these benefits:

- Reduces storage costs—You can share a single tape drive among several storage nodes. In fact, since NetWorker software uses the same open tape format for UNIX, Windows, NetWare and Linux, you can share the same tape between different platforms (assuming that respective save sets belong to the same pool).
- Reduces LAN traffic—You can configure clients as SAN storage nodes that can send save sets over the SAN to shared drives.




- Provides fault tolerance—Within a SAN environment, you can configure hardware to eliminate a single point of failure.
- Provides configuration over a greater distance—You can configure a system over a greater distance than with SCSI connections.

## DDS configuration overview

The following figure illustrates the DDS process and potential device sharing configurations. This basic configuration consists of a server, two storage nodes, and a library with two tape drives.

**Figure 24** Dynamic Drive Sharing



In this figure:

- Storage nodes sn\_1 and sn\_2 are attached to the library.
- Each storage node, on its own, has access to drive\_1 and drive\_2.
- With DDS enabled, both storage nodes have access to both drives and can recognize when a shared drive is in use.

This configuration requires two DDS licenses, one for each drive.

**Note:** Ensure that all applicable devices can be seen from each storage node by running the `inquire -l` command locally on each storage node.

## DDS block-size compatibility between UNIX and Windows

With DDS enabled, drives can be shared between storage nodes on different platforms, such as UNIX and Microsoft Windows. For NetWorker software operations (such as backups and recoveries) to take place successfully, ensure that the block size is compatible between different platforms or hardware.

To ensure compatibility, make sure one of the following conditions is met:

- The various storage nodes sharing a drive support the same block sizes.
- When a tape is labeled on a drive, it is labeled with the block size defined on the storage nodes.

## Block-size incompatibility between UNIX and Windows

Incompatible block-size settings between UNIX and Microsoft Windows storage nodes could result in any of these error scenarios:

- A backup taken on a UNIX node might not be recoverable on a Microsoft Windows node if the Windows node does not support large block sizes.
- A UNIX process labels and saves data to a tape and leaves the tape mounted. A Microsoft Windows process subsequently attempts to verify the label on this tape and fails because the label verification is done by reading a header from the data portion.
- A tape on a UNIX node is labeled with a large block size. The backup is started on a Microsoft Windows node and the Windows node attempts to write the backup by using the default block size. Internally, the backup on Windows is written by breaking down the big buffer of data into smaller segments of writable block sizes.  
Attempting to recover a specific file on Windows in this situation fails due to positioning errors on the tape. The data is still recoverable from the Windows side, since the NetWorker software will switch from using file and block positioning to reading the tape from the beginning to reach the correct position. The data might not, however, be recoverable from the UNIX side.

## Unintended Access to DDS device prevention

The Reserve/Release attribute has been added to the Device resource for tape devices to support Reserve/Release, including the Persistent Reserve commands.

Reserve/Release is a mechanism that uses SCSI commands to attempt to prevent unintended access to tape drives that are connected by using a shared-access technology, such as Fibre Channel, iSCSI, or SCSI multiplexers. It is a “cooperative” and host-based mechanism, which means that all applications should respect the reservations and not purposely break them. Access is granted based on the host system that reserved the device. Other applications that run on that host cannot be prevented from accessing a reserved device.

Reserve/Release cannot prevent a malicious or badly behaved application from accessing a reserved device. It also cannot prevent all problems caused by hardware issues (such as SCSI resets or FC LIPs) from interrupting data access.

The basic sequence requires that a host reserve a tape drive (using specific SCSI commands) before attempting to access the tape drive. If this “reservation” succeeds, then the host can use the drive. If the reservation fails (usually because the device is reserved by someone else), then the host attempting the reservation should not attempt to use the drive. When a host has finished using a reserved drive, that host must release the drive by using the appropriate SCSI commands.

The reservation is maintained by the drive itself. With older (called “Simple” in NetWorker software) Reserve/Release, the reservation is based on the SCSI ID of the system that issued the reserve command. For tape drives connected to Fibre Channel (FC) using FC-SCSI bridges, the mapping between FC host and reservation is done inside the bridge, since the initiator on the SCSI side is always the bridge itself, regardless which host actually issued the reserve command.

For Persistent Reserve, the reservation is associated with a 64-bit “key” that is registered by the host. Several keys can be registered with a given drive at any given time, but only one may hold the active reservation. NetWorker software uses the “exclusive” reservation method for Persistent Reserve. Only the host that holds the active reservation is allowed to access the drive.

The Reserve/Release attribute does not support file type or advanced file type devices.

The settings that relate to Reserve/Release and Persistent Reserve are found in a device’s **Properties** window, on the **Advanced** tab. They are visible only when diagnostic mode is turned on.

The default setting for Reserve/Release is None. Once any other Reserve/Release setting is selected, it works automatically, without further user intervention. The Reserve/Release attribute

is supported only on Common Device Interface (CDI) platforms, so if the CDI attribute in a device's **Properties** is set to Not Used, then Reserve/Release settings are ignored.

For newer hardware, once a Reserve/Release setting (other than None) has been selected, the appropriate Persistent Reserve commands are automatically issued before a device is opened for reading or writing, and before the device is closed. With older hardware, a SCSI-2 Reserve command is issued before opening the device, and a SCSI-2 Release command is issued after the device is closed.

Reserve/Release has these possible settings:

- None (the default)
- Simple
- Persistent Reserve
- Persistent Reserve + APTPL (Activate Persist Through Power Loss)

The Persistent Reserve Key attribute has also been added. It is used with Persistent Reservation calls.

### Restrictions for use of the SCSI Reserve/Release setting

There are restrictions for using the SCSI Reserve or Release setting.

Consider the following:

- It is available on CDI platforms only. Consequently, since CDI is not supported within an NDMP environment, Reserve/Release is not supported with NDMP.
- Not all drives support persistent Reserve/Release. (All drives support at least simple reserve release. The code automatically drops back from Persistent +APTPL or Persistent to Simple on drives that do not support Persistent.)
- SCSI resets can clear Simple reservations at the device.
- Even with Reserve/Release, there is no guarantee against data loss.
- If the operating system has its own Reserve/Release feature, that feature must be disabled in order for the NetWorker Reserve/Release feature to work.
- Even if all of the enterprise's NetWorker storage nodes have this feature enabled, then it is possible that, on the storage node where a backup operation is run, data loss can be caused by the operating system's utilities or by third-party programs.

## DDS on NDMP nodes in a SAN environment

You can configure shared drives between NDMP nodes in a SAN environment.

Ensure that:

- All the components of a SAN configuration are compatible when DDS is enabled with the NetWorker NDMP feature.
- The Fibre Channel switches are compatible with any NDMP hosts within a SAN.
- NDMP hosts and libraries in the SAN are compatible with each other.
- The NDMP nodes that will share the drives are homogeneous.

 **Note:** The current NDMP implementation does not allow the sharing of drives between non-homogeneous NDMP nodes. There is, however, no inherent limitation within DDS that would prevent this.

## DDS attributes in the device properties

Configure the attributes that DDS uses, in the **Properties** window for a device.

The attributes include:

- Hardware ID
- Shared Devices

### Hardware ID attribute

The Hardware ID attribute tracks the drives that are shared between multiple hosts. Device instances that share the same physical drive across multiple hosts have the same hardware ID. The device autoconfiguration process automatically assigns the Hardware ID to a device, or it is added when manually configuring a device. Users cannot edit the Hardware ID.

You can view the Hardware ID in the **Properties** window for a device, on the **General** tab, in the **Device Sharing** area.

NetWorker generates the Hardware ID when a device is scanned or configured. The Hardware ID consists of the following components:

- Hardware serial number
- Device type
- Worldwide part number (WWPN)
- Worldwide name (WWN)

### Shared Devices attribute

The Shared Devices attribute appears on the **Operations** tab of a device's **Properties** window when in diagnostic mode. It features values that can be used to manipulate all shared instances of a drive simultaneously. This attribute enables or disables all devices that share the same Hardware ID with a single action. The following table lists allowed values and descriptions for the attribute.

**Table 76** Shared Devices attributes

| Value       | Description                                                                                                                                              |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable All  | When selected, enables all devices with the same Hardware ID.                                                                                            |
| Disable All | When selected, disables all the devices with the same Hardware ID.                                                                                       |
| Done        | This value is the default setting. After the server has enabled or disabled all devices with the same Hardware ID, the attribute value is reset to Done. |

You cannot configure the Shared Devices attribute with the `jbconfig` program.

### Idle Device Timeout attribute and DDS

A tape might remain mounted in a drive after a backup completes. Other requests for the drive from another device path must wait during this timeout period. Use the Idle Device Timeout attribute to adjust the timeout value.

The Idle Device Timeout attribute is not specifically a DDS attribute, but is useful in configuring shared drives. This attribute appears on the device **Properties** window on the **Advanced** tab when displayed in Diagnostic Mode. The default value is 0 (zero) minutes, which means that the device never times out and you must manually eject the tape.

If the device belongs to a library, you can also specify the Idle Device Timeout value for all devices in the library. However, the library value will take effect only on those devices whose **Idle Device Timeout** value is 0. The Idle Device Timeout value for a library is located on the **Timer** tab of the library **Properties** window.

## Max active devices

In a DDS environment, use the Max active devices attribute, on the **General** tab of the Storage Node resource to define the maximum number of active devices for a storage node.

This attribute sets the maximum number of devices that NetWorker may use from the storage node in a DDS configuration. In large environments with media libraries that have a large number of devices, storage nodes might not have the ability to optimize all the drives in the library. The Max active devices attribute allows you to limit the number of devices that the storage node uses at a specified time, which allows the storage node to have access to all the devices in the library, but does not limit the storage node to the number of devices it can fully optimize.

## Configuring NDMP devices

You can back up NDMP data to an NDMP or non-NDMP device in a standalone or library configuration. You can also back up NDMP data to ACSLS controlled silos.

### Configuring a standalone NDMP device

Use the NetWorker Management Console (NMC) to configure a standalone Network Data Management Protocol (NDMP) tape device for Direct NDMP backups.

#### Procedure

1. In the **Administration** window, click **Devices**.
2. In the navigation tree, right-click **Devices**, and then select **New**.
3. In the **Name** attribute, specify the NDMP device in the format:

```
rd=NAS_hostname:NAS_device_handle (NDMP)
```

where:

- *NAS\_hostname* is the hostname of the NAS that has the NDMP device attached.
- *NAS\_device\_handle* is the path of the device.

 **Note:** Configure the NDMP device as a remote device and add (NDMP) after the pathname. Otherwise, you receive a message similar to the following:

```
NDMP device name shall be in rd=snode:devname (NDMP) format
```

4. In the **Media Type** attribute, specify the device type.
5. Specify a valid NAS administrator account in the **Remote User** attribute.
6. Specify the password for the NAS administrator account in the **Password** attribute.
7. On the **Configuration** tab:
  - a. Select the **NDMP** checkbox. You can only set this attribute when you create the device. You cannot change the NDMP attribute after you create the device. To change the device configuration, delete and re-create the device.
  - b. Set the **Target Sessions** attribute to 1. NDMP devices do not support multiplexing.
  - c. The **Dedicated Storage Node** attribute must remain at the default value: **no**.
8. Under the **Advanced** tab, the **CDI** attribute must remain at the default value: **Not used**.
9. (Optional) Change the block size that is used by the NDMP device.

By default, NDMP devices use a block size of 60 KB. If required, select a different block size in the **Device block size** field. When you configure the NDMP client, set the *NDMP\_AUTO\_BLOCK\_SIZE* environment variable in the **Application Information** attribute.

10. Click **OK**.

## Configuring an NDMP autochanger

You can use an NDMP autochanger to manage Direct NDMP or Three-party backups with NDMP devices. To configure an NDMP autochanger, use NMC or the `jbconfig` command.

### Configuring an NDMP autochanger with NMC

When you configure an NDMP autochanger in NMC, the NetWorker software first detects the NDMP devices and then configures the library.

#### Procedure

1. In the **NetWorker Administration** window, click **Devices**.
2. Right-click the NetWorker Server, and then select **Configure All Libraries**.
3. On the **Provide General Configuration Information** window, accept the default library type, **SCSI/NDMP**, and then click **Next**.
4. On the **Select Target Storage Nodes** window, click **Create a new Storage Node**.
5. On the **Storage Node Name** field, specify the hostname of the NAS.  
If a DinoStor TapeServer manages the autochanger, specify the DinoStor hostname.
6. In the **Device Scan Type** attribute, select **NDMP**.
7. In the **NDMP User Name** and **NDMP Password** fields, specify the NAS administrator account. If DinoStor TapeServer manages the autochanger, specify the DinoStor username and password.
8. Click **Start Configuration**.
9. Click **Finish**.
10. Monitor the **Log** window for the status of the device scan.

When you specify an incorrect username and password combination:

- The Log status window reports:

```
No configured libraries detected on storage node
storage_node_name
```

- The `daemon.raw` file on the NetWorker server reports:

```
NDMP Service Debug: The process id for NDMP service is 0xb6c0b7b0
42597:dvdetect: connect auth: connection has not been authorized
42610:dvdetect: The NDMP connection is not successfully
authorized on host 'storage_node_name'
```

To resolve this issue, relaunch the **Configure All Libraries** wizard and correct the NDMP username and password combination.

**Note:** If the **Log** window reports that NetWorker cannot detect the serial numbers for the library, see [Configuring an NDMP autochanger by using the `jbconfig` command](#) on page 39 for detailed instructions.

## Configuring an NDMP autochanger using the `jbconfig` command

The NMC interface is the preferred method to configure an NDMP autochanger. Use the `jbconfig` command when you cannot configure the autochanger by using the NMC Configure Library wizard.

### About this task

The *NetWorker Command Reference Guide* or the UNIX man page provides more information about the `jbconfig` command.

### Procedure

1. Log in to the NetWorker server as root on UNIX, or Administrator on Windows.
2. At the command prompt, type `jbconfig`
3. At the **What kind of jukebox are you configuring** prompt, type 2 to configure an autodetected NDMP SCSI jukebox.
4. When prompted for an NDMP username, specify the NAS administrator account.  
If DinoStor Tape Server manages the jukebox, then specify the DinoStor account.
5. When prompted for an NDMP password, specify the NAS administrator password.  
If DinoStor manages the jukebox, then specify the DinoStor password.
6. When prompted for the NDMP Tape Server Name, specify the NAS filer hostname.  
If DinoStor manages the autochanger, then specify the DinoStor hostname.
7. At the **What name do you want to assign to this jukebox device** prompt, provide a name to identify the autochanger.
8. To enable auto-cleaning, accept the default value of **Yes**, otherwise type `no`.
9. At the **Is (any path of) any drive intended for NDMP use? (yes / no) [no]** prompt, type `yes`.
10. At the **Is any drive going to have more than one path defined? (yes / no) [no]** prompt, type `no` if you will not configure shared devices. Type `yes` to configure shared drives.
11. When prompted, for the first pathname for the NDMP devices in the jukebox, perform the following steps:
  - a. Specify the pathname in the following format:  
`NDMP_tape_server_name:device_path`  
 where:
    - `NDMP_tape_server_name` is the hostname of the NDMP or DinoStor Tape Server.
    - `device_path` is the first device path.
12. Complete the prompts for the second device.
13. In the **Enter the drive type of drive 1** prompt, specify the number that corresponds to the NDMP device type.
14. If each drive in the autochanger is the same model, then type `yes`. Otherwise, type `no`, and then specify the appropriate device types for each additional autochanger device.
15. When prompted to configure another autochanger, type `no`.

## Changing the block size of an NDMP device

By default, the block size that is used to write data to an NDMP backup is 60KB. With the exception of Celerra, when you specify the `NDMP_AUTO_BLOCK_SIZE=Y` variable for an NDMP client, an NDMP device can use the value that is defined in its Device block size attribute.

### About this task

To determine the block sizes that are supported by the NDMP filer before setting the block size for an NDMP device, consult the applicable vendor documentation.

To change the block size that is defined for the NDMP device, perform the following steps:

### Procedure

1. From the **View** menu, select **Diagnostic Mode**.
2. In the **Devices** window, right-click the NDMP device, and then select **Properties**.
3. On the **Advanced** tab, select a value in the **Device block size** field.

 **Note:** The selected block size must not exceed the block size that is configured on the NAS filer.

4. Click **Ok**.

## Message displayed when CDI enabled on NDMP or file type device

If you enable the CDI feature for an NDMP tape device or file type device (FTD), a message similar to the following appears:

```
nsrd: media notice: The CDI attribute for device "/dev/rmt/3cbn" has
been changed to "Not used".
```

To avoid this message, do not enable the CDI attribute for these device types.

## Configuring NDMP-DSA devices

When you use DSA, NetWorker sends the NDMP data to a NDMP-DSA device, which includes tape, virtual tape, AFTD, and Data Domain devices. The steps to configure a NDMP-DSA device for a specified device type is the same as configuring a non-NDMP device. The *NetWorker Administration Guide* provides detailed information.

## Configuring the Clone Storage Node

When cloning NDMP data, specify the destination storage node, called the clone “write source” (the device that receives the clone data), in the Clone storage nodes attribute. The *NetWorker Administration Guide* provides details.

## Pools requirements for NDMP

When you create a pool for non-NDMP devices, select only the devices that are required by the NDMP clients.

NetWorker cannot send bootstrap and index backups to an NDMP device. When you do not configure a non-NDMP devices or a non-NDMP device is not available to receive the index and bootstrap backups, the NDMP client backup appears to hang. Configure a separate pool to direct the index and bootstrap to a non-NDMP device.

Auto media verification in the Pool resource does not support NDMP.



When an NDMP client backup is a member of a clone-enabled group, configure a clone pool with non-NDMP devices that are local to the NetWorker server to receive the clone bootstrap and index.

## Configure NetWorker for NDMP backup and clone operations

This section explains how to configure NetWorker for NDMP backup and clone operations.

### Creating and configuring the NDMP client resource

Use the NMC Client Configuration wizard to create the NDMP client or create the client manually. It is recommended that you use the NMC Client Configuration wizard to create NDMP clients.

#### Using the Client Configuration wizard

##### About this task

Use the NMC Client Configuration wizard to create the NDMP client.

##### Procedure

1. From the **Administration** window in NMC, click **Protection**.
2. In the expanded left pane, select **Clients**, and then select **Protection > New Client Wizard**.
3. On the **Specify Client Information** window:
  - a. In the **Client Name** field, specify the hostname of the filer.
  - b. (Optional) Add comments in the **Comment** field.
  - c. (Optional) In the **Tag** field, specify the name of the tag for the dynamic group in which you want to add this client.
  - d. (Optional) In the **Groups** area, select an existing group, in which to add the client.
  - e. In the **Type** area, select **NDMP**, and then click **Next**.
4. On the **Specify the NDMP Client Credentials** window:
  - a. In the **NDMP User Name** field, specify a valid NAS administrator account.
  - b. In the **NDMP Password** attribute, specify the password for the NAS administrator account, and then click **Next**.
5. In the **Specify the NDMP Client Backup Options** window:
  - a. In the **NDMP backup type** attribute, select or specify the backup type:
    - BlueArc: `dump`
    - Mirapoint: `image`
  - b. In the **NDMP Array Name** field, specify the logical name that is assigned to the NDMP NAS array.

The **NDMP Array Name** field enables you to configure the same NAS device with multiple NDMP clients that have different host IDs.

 **Note:** NDMP clients that use the same NAS device must have the same NDMP array name.

- c. Review the **App Info** options and disable options, as required. It is recommended that the default options remain enabled.

**Table 77** Application information variable types

| App Info Type          | Description                                                                                                                                                                                                                                                                                            |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HIST                   | Enables the backup of index data. If you do not select this option, you can only perform full recoveries of the backup data.                                                                                                                                                                           |
| UPDATE                 | Enables the backup process to update the last backup dates in database on the NDMP client, after the backup completes. Only applies to NetApp and has no effect when backing up other NAS systems.                                                                                                     |
| DIRECT                 | Enables DAR or DDAR support. <a href="#">DAR and DDAR</a> on page 335 provides more information.                                                                                                                                                                                                       |
| Use Token-Based Backup | Enables the NDMP backup to use last backup time tokens to decide what files to backup. Not all NDMP clients support token based backup. When you select this option and the NDMP data server on the client does not support token based backups, NetWorker performs the backup by using backup levels. |

- d. In the **Advanced App Info** field, specify additional NAS specific environments variables, one per line. The following table provides a list of the available **Application Information** environment variables for each NAS.

 **NOTICE** Environment variables are case-sensitive. Use an equal (=) sign to separate the environment variable name from its value.

**Table 78** Vendor-specific Application Information variables

| NAS       | Variables                             | Definition                                                                                                                                                                                                                      |
|-----------|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Mirapoint | <i>MIRA_OPTIONS= (fromimagefull=)</i> | Required. The (fromimagefull=) value allows full image and message (file) based backups to use the date of the image when performing the selection.                                                                             |
|           | <i>NDMP_AUTO_BLOCK_SIZE=Y</i>         | Optional. Specify this variable to override the default block size of 60 KB when writing NDMP backups to an NDMP device. Uses the block size value defined in the Device block size attribute when you labeled the NDMP volume. |
| BlueArc   | <i>NDMP_BLUEARC_FH_NAMETYPE=UNIX</i>  | Required. This variable requests that the BlueArc filer provide UNIX-style names when backing up a CIFS share.                                                                                                                  |
|           | <i>NDMP_AUTO_BLOCK_SIZE=Y</i>         | Optional. Specify this variable to override the default block size of 60 KB when writing NDMP backups to an NDMP device. Uses the block size value defined in the Device block size attribute when you labeled the NDMP volume. |

**Table 78** Vendor-specific Application Information variables (continued)

| NAS | Variables                     | Definition                                                                                                                                                                                                                                                                                                                                                                                       |
|-----|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|     | <i>USE_TBB_IF_AVAILABLE=N</i> | Optional. NetWorker enables Token-Based Backup (TBB) automatically if TBB is enabled from the NAS side. Specify this variable to disable TBB support for incremental backups. This value reverts the backup to the native level based backup of the NAS. If TBB is disabled from the NAS side, NetWorker will not enable TBB regardless what the value of <i>USE_TBB_IF_AVAILABLE</i> is set to. |

6. Click **Next**.
7. On the **Select the NetWorker Client Properties** window:
  - a. In the **Priority** field, specify the order in which the NetWorker server contacts clients in a protection group for backup. The attribute can contain a value between 1 and 1,000. The lower the value, the higher the priority.
  - b. In the **Parallelism** attribute:
    - For Direct-NDMP, set the **Parallelism** attribute to 1.
    - For NDMP-DSA, the parallelism value depends on the NAS capabilities and set parallelism to a value that is appropriate for the NAS. Parallelism values of 4 to 8 are common. In general, the best parallelism setting depends on filer configuration and the amount of installed RAM.
  - c. In the **Remote Access** attribute:
 

Specify the root account on Linux/UNIX, and/or the administrator account on Windows, of any computer that you use to browse backups of the NAS. Specify each account on a separate line. For example:

```
administrator@windows_hostname
root@linux_hostname
```
  - d. Select the **Data Domain Interface**. This option specifies the protocol to use if you send the backup data to a Data Domain Device. Available selections are Any, Fibre Channel, or IP.
  - e. Do not select the **Block Based Backup** or **Client Direct** options, as they do not apply to NDMP backups.
8. On the **Specify the File System Objects** window, select or specify objects to backup.
  - When the NAS supports NDMP snapshot management extension, you can browse and mark individual file systems to back up. When the client supports browsing, by default, NetWorker selects all objects.
  - When the client does not support browsing, specify the save sets to back up.
    - To back up all the file systems on the client, type **ALL**.
    - NAS versions earlier than 3 do not support the **ALL** save set. List the file systems one per line.
    - When you do not use the **ALL** save set, specify the file system name, as configured on the NAS.

- File system names in the **Save set** field are case sensitive.
- You cannot specify a share name.

To back up large client file systems, optionally schedule each file system to back up separately. For example, create two separate clients with the same name, but with different save sets.

9. Click **Next**.
10. On the **Client Configuration Summary** window, review the attributes, and then click **Create**.
11. On the **Client Configuration Results** window, review the results, and then click **Finish** to exit the wizard.

[Troubleshooting NDMP configuration and backup failures for PowerStore, Unity, Celerra, VNX, and VNXe](#) describes how to resolve errors that you may experience when you configure the NDMP client.

## Performing post Client Configuration Wizard steps

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

### Modifying the Storage Node

#### About this task

On the **Globals (2 of 2)** tab, specify the storage node in the **Storage Nodes** attribute.

The attribute value depends on the type of backup:

- When you perform Direct-NDMP backups with NDMP devices, specify the hostname of the NAS that manages the tape device or autochanger.
- For three-party backups, specify the destination host first.
- For NDMP-DSA backups, specify the hostname of the storage node that manages the tape device or autochanger. If the NetWorker server is the storage node, specify *nsrserverhost*.
- For a DinoStor-managed NAS, specify the hostname of the DinoStor server first.

 **NOTICE** For NDMP-DSA backups, the NetWorker software uses the **Storage Node** attribute field of the NDMP client to determine which host receives the backup data. The `nsrndmp_save` command does not require the `-M` and `-P` options. If you specify the `-M` and `-P` options, they override the **Storage Node** attribute value.

## Adding NDMP Client Properties

### Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Clients**.
3. Right-click a **Client**, and then select **New Client Properties**. The **Client Properties** screen appears.
4. Select the **Apps & Modules** tab. Select the **NDMP** attributes of the **Client**. The **NDMP** options are as follows:
  - **NDMP**—Select this box to indicate whether this client is an NDMP client.
  - **NDMP multistreams enabled**—Do not select this box. Only Isilon supports multistreaming.
  - **NDMP log successful file recovery**—By default, NetWorker does not print each successfully recovered file name in the log messages, because this logging impacts

performance and takes up space. To enable the logging of successful recoveries for each file, select this checkbox.

- **Disable IPv6**—Check this box to disable IPv6 on NDMP backup and recovery.
- **NDMP array name**—This name is the logical name that is assigned to the array in NDMP NAS array configurations.
- **NDMP vendor information**—This attribute contains NDMP client vendor information.

5. Click **OK**.

## Configuring the NDMP client manually

It is recommended that you create Network Data Management Protocol (NDMP) clients by using the Client Configuration wizard. If you create the NDMP client manually, then the configuration details for each attribute in the Client Configuration wizard apply when you create the client manually.

Review this information before you configure an NDMP client manually:

- For an NDMP configuration that includes Storage Node resources, configure a Client resource for each storage node that you define for an NDMP backup and clone operation.
- For NDMP three-party storage nodes that use NDMP devices, repeat these steps for each NDMP storage node.
- For NDMP-DSA storage nodes, create the NetWorker Client resources in the same manner as non-NDMP clients. The *NetWorker Administration Guide* provides details on how to create a non-NDMP Client resource.
- NDMP does not support the use of directives including AES encryption. The NetWorker software ignores any value that you define in the **Directives** attribute for an NDMP client.
- When you select **Checkpoint enabled** on the **General** tab, do not modify the **Checkpoint granularity** attribute. NDMP backups do not support checkpoint granularity and the NetWorker software ignores any value that you define for this attribute.
- If the NAS supports the NDMP snapshot management extension, then you can browse and mark individual file systems for backup instead of specifying the save sets in the **Save set** attribute. You cannot use the **Save set browse** icon to browse the NDMP file system until you:
  - Select the **NDMP** checkbox, on the **Apps & Modules** tab.
  - Specify the NDMP username and password in the **Remote user and password** fields on the **Apps and Modules** tab.

### **Note:**

- Celerra, VNX, VNXe, and NetApp C-mode do not support Snapshot Management Extension. Only NetApp 7-Mode supports Snapshot Management Extension.
- Isilon, Celerra, VNX, VNXe, and NetApp C-mode filers do not allow you to browse. Only NetApp 7-Mode allows you to browse.

## Performing schedule backup and clone operations

Data Protection Policies provide you with the ability to schedule backup and clone operations, to protect NDMP data.

You can use the NDMP protocol to protect data on NAS devices.

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

## Overview of protection policies

A protection policy allows you to design a protection solution for your environment at the data level instead of at the host level. With a data protection policy, each client in the environment is a backup object and not simply a host.

Data protection policies enable you to back up and manage data in a variety of environments, as well as to perform system maintenance tasks on the NetWorker server. You can use either the **NetWorker Management Web UI** or the NMC **NetWorker Administration** window to create your data protection policy solution.

A data protection policy solution encompasses the configuration of the following key NetWorker resources:

### Policies

Policies provide you with a service-catalog approach to the configuration of a NetWorker datazone. Policies enable you to manage all data protection tasks and the data protection lifecycle from a central location.

Policies provide an organizational container for the workflows, actions, and groups that support and define the backup, clone, management, and system maintenance actions that you want to perform.

### Workflows

The policy workflow defines a list of actions to perform sequentially or concurrently, a schedule window during which the workflow can run, and the protection group to which the workflow applies. You can create a workflow when you create a new policy, or you can create a workflow for an existing policy.

A workflow can be as simple as a single action that applies to a finite list of Client resources, or a complex chain of actions that apply to a dynamically changing list of resources. In a workflow, some actions can be set to occur sequentially, and others can occur concurrently.

You can create multiple workflows in a single policy. However, each workflow can belong to only one policy. When you add multiple workflows to the same policy, you can logically group data protection activities with similar service level provisions together, to provide easier configuration, access, and task execution.

### Protection groups

Protection groups define a set of static or dynamic Client resources or save sets to which a workflow applies. There are also dedicated protection groups for backups in a VMware environment or for snapshot backups on a NAS device. Review the following information about protection groups:

- Create one protection group for each workflow. Each group can be assigned to only one workflow.
- You can add the same Client resources and save sets to more than one group at a time.
- You can create the group before you create the workflow, or you can create the group after you create the workflow and then assign the group to the workflow later.

### Actions

Actions are the key resources in a workflow for a data protection policy and define a specific task (for example, a backup or clone) that occurs on the client resources in the group assigned to the workflow. NetWorker uses a work list to define the task. A work list is composed of one or several work items. Work items include client resources, virtual machines, save sets, or tags. You can chain multiple actions together to occur sequentially or concurrently in a workflow. All chained actions use the same work list.

When you configure an action, you define the days on which to perform the action, as well as other settings specific to the action. For example, you can specify a destination pool, a retention period, and a target storage node for the backup action, which can differ from the subsequent action that clones the data.


When you create an action for a policy that is associated with the virtual machine backup, you can select one of the following data protection action types:

- **Backup** — Performs a backup of virtual machines in vCenter to a Data Domain system. You can only perform one VMware backup action per workflow. The VMware backup action must occur before clone actions.
- **Clone** — Performs a clone of the VMware backup on a Data Domain system to any clone device that NetWorker supports (including Data Domain system or tape targets). You can specify multiple clone actions. Clone actions must occur after the Backup action.

You can create multiple actions for a single workflow. However, each action applies to a single workflow and policy.

The following figure provides a high level overview of the components that make up a data protection policy in a datazone.

**Figure 25** Data Protection Policy



## 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 26** 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.

## Overview of configuring a new data protection policy

### About this task

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



**Example 4** Example of a data protection policy with 2 workflows (continued)

**Figure 27** 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*.

## Strategies for traditional backups

The primary considerations for a traditional backup strategy are the groups of Client resources, the workflows that define the series of actions that are associated with the backup, and the schedule for the backup.

### 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 traditional backups

A protection groups for traditional backups identifies the client resources to back up.

Traditional backups support the following types of protection groups:

- 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.

Create multiple groups to perform different types of backups for different Client resources, or to perform backups on different schedules. For example:

- Create one group for backups of clients in the Accounting department, and another group for backups of clients in the Marketing department.

- Create one group for file system backups and one group for backups of Microsoft Exchange data with the NetWorker Module for Microsoft.
- Create one group for a workflow with backups actions that start at 11 p.m., and another group for a workflow with backup actions that start at 2 a.m.

 **Note:** A Client resource can belong to more than one group.

### 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 Datanones** 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.

#### About this task

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](#).

## Actions sequences in traditional backup workflows

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


A workflow for a traditional backup can optionally include a probe or check connectivity action before the backup, and a clone action either concurrently with or after the backup.

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

### Workflow path from a traditional backup action

The only action that can follow a traditional backup is a clone action.

**Figure 28** Workflow path from a traditional backup action



## 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, monthly, or reference schedule for the action:
    - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
    - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
    - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
  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 79** 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.  

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.
  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.
- If a schedule is associated to an action, then override option is disabled.

20. Click **Next**.

The **Action Configuration Summary** page appears.

21. 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 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. In-built commands from NetWorker should be avoided as probe command.

#### 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, monthly, or reference schedule for the action:
  - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
  - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
  - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
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 80** 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 action 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 action if any of the probes 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.  
For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.  
If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.  
When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.  
Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.
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.

- Do not change the default selections for the Notification group box. NetWorker does not support notifications for probe actions and ignores and specified values.
- 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.
- 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.
- (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.
- (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.
- If a schedule is associated to an action, then override option is disabled.

- Click **Next**.  
The **Action Configuration Summary** page appears.
- Review the settings that you specified for the action, and then click **Configure**.

## Creating a traditional backup action

A traditional backup is a scheduled backup of the save sets defined for the Client resources in the assigned group for the workflow.

### Before you begin

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

### 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 **Backup**.
6. From the secondary action list, select the backup type, for example, **Traditional**.
7. (Optional) From the **Force Backup Level** list select a backup level.

For workflows that have more than one scheduled backup within a 24-hour period, use the **Force Backup Level** attribute to allow more than one backup to occur at two different backup levels in a 24-hour period. When you select a backup level in the **Force Backup Level** attribute, the first backup is performed at the scheduled backup level. Each subsequent occurrence of the backup action in the next 24 hours occurs at the level defined in the **Force Backup Level** attribute. For example, if the level defined by the schedule is Full and the **Force Backup Level** attribute is set to Incr, the first backup started by the action occurs at a level full and subsequent backups, within 24 hours of the start of the full backup are incremental. By default this option is cleared, which means that if the action runs multiple backup operations in a 24 period, all the backups occur at the scheduled backup level.

8. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
9. 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.
10. Specify a weekly, monthly, or reference schedule for the action:
- To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
  - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
  - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
11. To specify the backup level to perform when **Define** option under **Select Schedule** is selected, click the icon on each day. If it is **Select** option under **Select Schedule**, choose the customized NSR schedule using the drop down, after specifying the backup level within the NSR schedule resource.

The following table provides details about the backup level that each icon represents.

**Table 81** Schedule icons

| Icon                                                                                | Label                                                                                                                                                         | Description                                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|   | <b>Full</b>                                                                                                                                                   | Perform a full backup on this day. Full backups include all files, regardless of whether the files changed.                                                                                                               |
|  | <b>Incr</b>                                                                                                                                                   | Perform an incremental backup on this day. Incremental backups include files that have changed since the last backup of any type (full or incremental).                                                                   |
|  | <b>Cumulative Incr</b>                                                                                                                                        | Perform a cumulative incremental backup. Cumulative incremental backups include files that have changed since the last full backup.                                                                                       |
|  | <b>Logs Only</b>                                                                                                                                              | Perform a backup of only database transaction logs.                                                                                                                                                                       |
|  | <b>Incremental Synthetic Full</b><br> <b>Note:</b> Not supported for NDMP. | Perform an incremental synthetic backup on this day. An incremental synthetic full backup includes all data that changed since the last full backup and subsequent incremental backups to create a synthetic full backup. |

**Table 81** Schedule icons (continued)

| Icon                                                                              | Label       | Description                          |
|-----------------------------------------------------------------------------------|-------------|--------------------------------------|
|  | <b>Skip</b> | Do not perform a backup on this day. |

To perform the same type of backup on each day, select the backup type from the list and click **Make All**.

NetWorker does not support the use of synthetic full backup levels for NDMP data.

Celerra, Isilon, VNX, Unity, and NetApp filers with NDMP version 4 or later support token-based backups (TBB) to perform NDMP full and incremental backups. NetWorker supports the same number of incremental levels that the NAS vendor supports. Celerra, Isilon, and NetApp documentation provide the maximum number of incremental levels that the TBB incremental backup can support.

When you configure TBB after you update the NetWorker server from 7.6 SP1 or earlier, the first incremental backup does not occur until after one complete full backup.

Filers that do not support TBB, do not support incremental backups. If you select the level Incr, the NetWorker server performs a full backup.

Verify that the NAS storage vendor supports NDMP incremental backups before you use this feature.

12. Click **Next**.

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

13. From the **Destination Storage Node** box, select the storage node with the devices on which to store the backup data.
14. From the **Destination Pool** box, select the media pool in which to store the backup data.
15. From the **Retention** boxes, specify the amount of time to retain the backup data.

After the retention period expires, the save set is removed from the client file index and marked as recyclable in the media database during an expiration server maintenance task.

When you define the retention policy an NDMP client, consider the amount of disk space that is required for the client file index. NDMP clients with several thousands of small files have significantly larger client file indexes on the NetWorker server than a non-NDMP client. A long retention policy for an NDMP client increases disk space requirements on the file system that contains the client file indexes.

16. From the **Client Override Behavior** box, specify how NetWorker uses certain client configuration attributes that perform the same function as attributes in the Action resource:
  - **Client Can Override**—The values in the Client resource for **Schedule**, **Pool**, **Retention policy**, and the **Storage Node** attributes take precedence over the values that are defined in the equivalent Action resource attributes.
    -  **Note:** If the NetWorker policy action schedule is set to the **Skip** backup level, the **Client can Override** option is not honored. For NetWorker to consider the **Client can Override** option, change the action schedule to a different level.
  - **Client Can Not Override**—The values in the Action resource for the **Schedule**, **Destination Pool**, **Destination Storage Node**, and the **Retention** attributes take precedence over the values that are defined in the equivalent Client resource attributes.
  - **Legacy Backup Rules**—This value only appears in actions that are created by the migration process. The updating process sets the **Client Override Behavior** for the migrated backup actions to **Legacy Backup Rules**.

17. Click **Next**.

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

18. 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.

19. 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.

20. 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.

21. 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

22. 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.

23. 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.
24. 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.
25. (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.
26. (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.
- If a schedule is associated to an action, then override option is disabled.

27. 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**.
28. 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 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.

#### 29. Click **Next**.

The **Action Configuration Summary** page appears.

#### 30. 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 after the backup. A clone action is the only supported action after a backup action in a workflow.

## Cloning NDMP save sets

You can clone Direct-NDMP and NDMP-DSA save sets by using the same methods used to clone non-NDMP save sets.

Before you clone NDMP save sets, review these requirements:

- To clone Direct-NDMP or Three-party backup data:
  - The source NAS must run NDMP version 3 or later.
  - The destination NAS can run any version of NDMP, but you cannot clone a volume cloned with NDMP earlier than version 3 to another volume.
  - You cannot clone NDMP save sets to a non-NDMP device.
  - You can clone NDMP tapes from one NDMP host to another NDMP host of the same type. For example, you can clone tapes from a NetApp filer with an attached library to another NetApp filer or to the same filer.
- You require two NDMP devices to clone the NDMP save sets, one device to perform the read operation and one device to perform the write operation.
- You must clone NDMP-DSA backups to non-NDMP devices. You can however, clone NDMP-DSA save from one type of tape device to another. For example you can clone save sets on a DLT device to an AIT device.

- Use the nsrclone program to clone NDMP save sets from a command prompt. The *NetWorker Command Reference Guide* or the UNIX man pages provide more information on nsrclone usage.

### 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

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 **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, monthly, or reference schedule for the action:
  - To specify a schedule for each day of the week, select **Define** option under **Select Schedule** and period as **Weekly by day**.
  - To specify a schedule for each day of the month, select **Define** option under **Select Schedule** and period as **Monthly by day**.
  - To specify a customized schedule to the action, select **Select** option under **Select Schedule** and choose a customized schedule using the drop down menu that is already created under NSR schedule resource.
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 82** 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:
- From the **Destination Storage Node** list, select the storage node with the devices on which to store the cloned save sets.
  - 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.
  - From the **Destination Pool** list, select the target media pool for the cloned save sets.
  - 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:
- 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.
  - 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 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.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

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.

- 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**.
- 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 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.

- 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.
  - If a schedule is associated to an action, then override option is disabled.
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 traditional backup workflows

**Figure 29** Traditional backup workflow



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.

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**.

## Performing manual NDMP backups

After you configure the NetWorker server for NDMP backup data operations, you can perform manual NDMP backups.

On Windows, you can manually back up NDMP data by using the NetWorker User program. The method to backup NDMP data is the same as a non-NDMP local backup. You cannot perform a three-party backup with the NetWorker User program.

On Windows and UNIX, you can perform a manual backup from a command prompt by using the `nsrndmp_save` command.

Before performing a manual backup by using the `nsrndmp_save` command or the NetWorker User program, review these requirements:

- You can only perform manual Direct-NDMP backups from a NetWorker server.
- You can start a manual NDMP-DSA backup from a NetWorker server, storage node, or client. When you do not start the NDMP-DSA backup from the NetWorker server, the `servers` file on the NetWorker server and storage node, must contain the hostname of the host that initiates the backup.
- Before you perform a manual backup, you must configure the NDMP client on the NetWorker server. Manual backups use client configuration information for example, the variables that are defined in the **Application Information** attribute of an NDMP client.
- Direct-NDMP and three-party NDMP backups support manual DAR backups when the NDMP client contains the `DIRECT=Y` and `HIST=Y` environment variables in the **Application Information** attribute for the NDMP client.

 **NOTICE** To use DAR, the NAS filer must use NDMP version 4. The *NetWorker E-LAB Navigator* describes how to determine if a particular NDMP vendor supports DAR.

## Performing an NDMP backup from the command line

### About this task

Use the `nsrndmp_save` command to perform a manual command line NDMP backup.

The `nsrndmp_save` command does not back up the bootstrap. Without the bootstrap, you cannot perform a disaster recovery of the NetWorker server. To back up the bootstrap, run the `nsrpolicy -G nsrpolicy policy_name start` command from the NetWorker server. The `nsrpolicy` command uses the attribute values specified for the policy. For example, the pool and schedule values.

To perform an NDMP backup from the command prompt, use the following syntax:

```
nsrndmp_save -T backup_type -s NetWorker_servername -c clientname -l
backup_level -t date_time -g nsrpolicy_path
```



where:

- *backup\_type* is a supported backup type for the NAS filer:
  - BlueArc only supports the `dump` backup type.
  - MiraPoint only supports the `image` backup type.
- *backup\_level* is a full for a full backup, `incr` for an incremental backup. Each NAS supports full backups.
  - ① **Note:** NetWorker only supports or allows Full, Incremental, and Differential level 1 NDMP backups of all NAS.
- *date\_time* is the date and time of the last backup, which is enclosed in double quotes. Specify this value for `incr` level backups. When you do not specify the date and time, the backup is a native NDMP level-based backup.
  - ① **NOTICE** During a NetWorker scheduled policy backup, the NetWorker software supplies the date and the time information, and incremental and level backups work as expected.

Use one of these methods to determine the date and time of the last NDMP backup:

- Review the `daemon.raw` file on the NetWorker server or the `savegroup` completion report for a line similar to the following:

```
42920:nsrndmp_save: browsable savetime=1296694621
```

Use the value after `savetime=` with the `-t` option.

- Specify the date and time of the last backup reported by the `mminfo` command for the NDMP save set.

- ① **Note:** You can force the NDMP backup and recovery to ignore IPv6 and instead use IPv4 in one of two ways:

1. Add the `-f` option to the `nsrndmp_save` or the `nsrndmp_recover` commands as required.
2. On the Apps & Modules tab in the **Client Properties** window, select **Disable IPv6**.

## Example of an NDMP backup

To perform an incremental backup of a NetApp client that is named `mynetapp`, perform the following steps:

1. Determine the time of the last full backup:

```
mminfo -v -c mynetapp
```

**Table 83** NDMP backup

| client   | date     | time     | size   | ssid           | fl   | lvl  | name      |
|----------|----------|----------|--------|----------------|------|------|-----------|
| mynetapp | 08/16/15 | 15:23:58 | 1853MB | 38648127<br>01 | cbNs | full | /.../set1 |
| mynetapp | 08/17/15 | 15:39:58 | 815MB  | 38480364<br>30 | cbNs | incr | /.../set2 |

2. Specify the last backup time in `nsrndmp_save` command:

```
nsrndmp_save -T dump -s my_nwserver -c mynetapp -l incr -t "02/16/11
15:23:58" -g mygroup path
```

For NDMP-DSA backups, the NetWorker software uses the Storage Node attribute field of the NDMP client to determine which host receives the backup data. The `nsrndmp_save` command

does not require the `-M` and `-P` options. If you specify the `-M` and `-P` options, they override the **Storage Node** attribute value. The *NetWorker Command Reference Guide* and the `nsrndmp_save` man page on UNIX provide more information.

## Troubleshooting NDMP configuration and backup failures for other filers

This section provides a list of the possible causes and the resolutions for NDMP backup failures.

### Unable to connect to NDMP host *hostname*

This message appears when the NetWorker server cannot create or modify an NDMP client.

To resolve this issue ensure that the environment meets the following requirements:

- Username and password specified for the client is correct and has sufficient permissions to perform NDMP operations.
- NDMP service is running on the filer.

### Cannot perform NDMP backup after the NetWorker server licenses expire

If a NetWorker sever running in evaluation mode expires before you authorize the server, NDMP devices remain disabled after the addition of the required licenses and authorization of the NetWorker server.

To re-enable NDMP devices, perform the following steps:

1. To connect to the NetWorker server, use NMC, and then click the **Devices** button.
2. In the **Devices** windows, right-click the NDMP device, and then select **Properties**.
3. Click the **Configuration** tab, and then set the **Target Sessions** attribute to 1.
4. Click the **General** tab, and then in the **Enabled** section, select **Yes**.
5. Click **Ok**.

### Failed to store index entries

This error message occurs in the `daemon.raw` file when an index backups fails due to an insufficient amount of swap space.

To resolve this issue, increase the amount of swap space available to the NetWorker server.

 **NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

### IO\_WritePage write failed - No space left on device (28): No space left on device

This error message appears in the `daemon.raw` file when the index backup fails. There is insufficient temporary space to store the index entries before the NetWorker software commits the information into the client file index.

To resolve this issue, specify a new the temp directory with sufficient disk space in one of the following ways:

- Define the `NSR_NDMP_TMP_DIR` environment variable in the Application Information attribute of the client.
- Define the `NSR_NDMP_TMP_DIR` as an operating system environment variable on the NetWorker server.

[Memory and space requirements for NDMP FH updates](#) on page 23 describes how to determine the amount of disk space the NetWorker software requires to temporarily store client files index entries.

**NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

### nsrndmp\_save: get extension list: communication failure

This message appears during a NDMP local backup when NetWorker cannot determine the filer name.

To resolve this issue, perform the following steps:

1. From a command prompt on the NetWorker server, type:

```
ndmpsup -c NDMP_hostname -o output_filename
```

For example:

```
ndmpsup -c myfiler.mnd.com -o ndmpsup.txt
```

2. Edit the output file that the `ndmpsup` command generates and search for the string **Vendor Name**. Make note of the reported Vendor Name.

For example:

```
Vendor Name = BlueArc Corp
```

3. Change to the `/nsr/debug` directory on UNIX or the `NetWorker_installation_dir\nsr\debug` directory on Windows.

4. Create new empty file and name it with the following format:

```
ndmpgettextlist_disable_VENDOR_NAME
```

where you replace `VENDOR_NAME` with the vendor name of the filer reported in the `ndmpsup` output file.

For example, to create this file for a BlueArc filer on UNIX, type:

```
touch "ndmpgettextlist_disable_BlueArc Corp"
```

### Error reading the FH entries from save through stdin

This error message appears in the `daemon.raw` file of the NetWorker server when there is a communication error between the `nsrndmp_save` and `nsrndmp_2fh` processes.

Resolve any communication or connection issues, then retry the backup.

**NOTICE** You cannot use the NetWorker User program to perform file-by-file and save set recoveries from a backup when the corresponding index update failed.

### Cannot find file history info for file name...You may still be able to recover this file with a save set recovery

This error message appears in the `daemon.raw` file of the NetWorker server when file history (FH) information is missing or corrupted for the file that is specified in the error message. For example, NetWorker cannot update the client file index (CFI) with FH information when a backup process interruption occurs during the failover of a clustered NetWorker environment.

You cannot perform an Network Data Management Protocol (NDMP) file-by-file recover or a save set recover when the CFI does not contain the associated FH information.

To recover this file, perform a save set recover from the command prompt. [Performing an NDMP save set recovery from the command prompt](#) on page 105 provides for further information.

**NOTICE** The NetWorker server does not delete the FH files that are stored in the `tmp` directory when the CFI updates fail.

## nsrndmp\_save: data connect: failed to establish connection

This error message appears in the `daemon.raw` file of the NetWorker server for several reasons:

- Network connectivity or name resolution issues exist between the NetWorker server and the NDMP client.
- You specified an incorrect NDMP username or password specified for the NDMP client.
- The NDMP service is not started on the NAS filer.
- The NetWorker server cannot communicate with the NAS filer over port 10000.
- A free port in the NetWorker server's default port range (7937-9936) is not available during an NDMP-DSA backup.  
The *NetWorker Security Configuration Guide* provides more information about NDMP port requirements and configuration.
- A misconfigured loop router. For a Celerra filer, the server route command utility configures the loop router. For NetApp, the route utility configures loop back router. The value of this setup is network-specific and depends on the number of switches and hubs between the NAS filer, NetWorker server, and NetWorker storage node.
- On the host where DSA is running, if the hostname is present in the hosts file, the `nsrdsa_save` process uses this name during backup. The DSA host passes the loopback entry to the NDMP data server and the connection fails. To resolve this issue, remove the hostname from the localhost list.

Knowledge base articles on the Support website provides detailed troubleshooting information for this error message and other failed to establish connection failures that you might encounter during an NDMP backup.

## Monitoring NetWorker Server activities in the Administration window

The **Monitoring** window in the NetWorker **Administration** application enables you to monitor the activities of an individual NetWorker Server.

### About this task

The **Monitoring** window provides the following types of activity and status information:

- Data protection policies, workflows, and individual actions.
- Cloning, recovering, synthetic full backups, and browsing of client file indexes.
- Operations that are related to devices and jukeboxes.
- Alerts and log messages.

You can also perform some management operations from the **Monitoring** window, for example, starting, stopping, or restarting a data protection policy.

### Procedure

1. From the **NMC Console** window, click **Enterprise**.
2. In the **Enterprise** view, right-click the NetWorker Server, and then select **Launch Application**.  
The **Administration** window appears.
3. To view the **Monitoring** window, click **Monitoring**.

## About the Monitoring window

On the **Administration** window taskbar, select **Monitoring** to view the details of current NetWorker server activities and status, such as:

- Policies and actions.
- Cloning, recovering, synthetic backups, checkpoint restart backups, and browsing of client file indexes.
- Alerts and log messages, and operations that are related to devices and jukeboxes.

While the **Monitoring** window is used primarily to monitor NetWorker server activities, it can also be used to perform certain operations. These operations include starting, stopping, or restarting a workflow.

The **Monitoring** window includes a docking panel that displays specific types of information. Select the types of information you want to view from the docking panel.

A portion of the **Monitoring** window, which is known as the task monitoring area, is always visible across all windows. A splitter separates the task monitoring area from the rest of the window. You can click and move the splitter to resize the task monitoring area. The arrow icon in the upper right corner of the **Monitoring** window allows you to select which tasks you want to appear in this view.

Smaller windows appear within the **Monitoring** window for each window. Each smaller window, once undocked, is a floating window and can be moved around the page to customize the view. You can select multiple types from the panel to create multiple floating windows that can be viewed simultaneously. The following table describes the various types of information available in the docking panel, and the details each one provides.

**Table 84** Monitoring window panel

| Window                  | Information provided                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Policies/Actions</b> | The <b>Policies</b> tab provides you with status information about all configure policies and the associated workflows and actions. The <b>Actions</b> tab provides you with status information for all actions. <a href="#">Policies/Actions pane</a> on page 81 provides more information.                                                                                                                                   |
| <b>Sessions</b>         | Allows you to customize whether to display all session types, or only certain session types. The information that is provided depends on which session type you select. For example, if you select <b>Save Sessions</b> , the window lists clients, save sets, groups, backup level, backup start time, duration of the backup, devices, rate, and size. <a href="#">Sessions window</a> on page 85 provides more information. |
| <b>Alerts</b>           | Lists the priority, category, time, and message of any alerts. <a href="#">Alerts pane</a> on page 86 provides more information.                                                                                                                                                                                                                                                                                               |
| <b>Devices</b>          | Lists devices, device status, storage nodes, libraries, volumes, pools, and related messages. <a href="#">Devices pane</a> on page 87 provides more information.                                                                                                                                                                                                                                                               |
| <b>Operations</b>       | Lists the status of all library and silo operations, including <code>nsrjb</code> operations that are run from the command prompt. Also lists user input, libraries, origin, operation data, operation start time, duration of the operation, progress messages, and error messages.<br><br>When displaying Show Details from the <b>Operations</b> window, the length of time that the window is                              |

**Table 84** Monitoring window panel (continued)

| Window     | Information provided                                                                                                                                                                                                                                                                    |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|            | displayed depends on the value that is typed in the <b>Operation Lifespan</b> attribute on the <b>Timers</b> tab of the <b>Properties</b> dialog box for the corresponding library. To access library properties, click <b>Devices</b> in the taskbar. By default, this pane is hidden. |
| <b>Log</b> | Lists messages that are generated by the NetWorker server, including the priority of each message, the time the message was generated, the source of the message, and the category. <a href="#">Log window</a> on page 90 provides more information.                                    |

## Customizing the Monitoring window

This section describes how to customize the **Monitoring** window in the **Administration** interface.

### Customizing tables

You can customize the organization and display of tabular information in the **Monitoring** window.

#### Sorting tables

You can change the display of tabular information that appears in the window. You can sort Table grids by column heading, and then by alphabetic or numeric order within those columns.

1. Drag and drop the column heading to its new position.
2. Click the column heading to sort the items into alphabetic and numeric order. An arrow appears in the column heading to indicate the sort order.

#### Sorting selected rows in a table

Selected rows are sorted to the top of the table. This sorting is particularly useful when you select **Highlight All** from the Find panel to select all rows matching the Find criteria and then moving all selected rows to the top of the table to view the results.

1. From the **Edit** menu, select **Find**, or press **Ctrl + F** to view the **Find** panel.
2. To select the rows, click each row or use the Find criteria.
3. Select **Sort Selected**.

#### Sorting multiple columns in a table

You can select the column that you want to use as the tertiary sort key, the secondary sort key, and the primary sort key.

1. Click the column that you want to use as the last sort key.
2. Click the column that you want to use as the next-to-last sort key, and so on, until you select the primary column.

#### Displaying columns in a table

You can select which columns to display in a table.

1. From the **View** menu, select **Choose Table Columns**.

2. Click a column name to select or clear the column and then click **OK**. You can also select the columns to display by right-clicking a table header and selecting **Add Column** from the drop-down.

## Displaying panes

You can choose to show or hide panes in the **Monitoring** window.

### About this task

Perform the following steps to hide or show a pane in the **Monitoring** window.

### Procedure

1. From the **View** menu, select **Show**. A check mark appears beside the panes that appear in the **Monitoring** window.
2. To hide a pane, select a marked pane.  
A check mark does not appear beside the pane.
3. To show a pane, select an unmarked pane.  
A check mark appears beside the pane.

## Policies/Actions pane

The **Policies/Actions** pane provides you with the ability to review status information about policies and actions.

This pane has two tabs:

- **Policies**—Provides a navigation tree that displays all configured policies on the NetWorker Server. Expand each policy to display the workflows that are associated with each policy. Expand each workflow to display each action that is contained in the workflow.
- **Actions**—Provides a list of all Action resources.

### Policies pane

The **Monitoring** window in the **NetWorker Administration** window enables you to monitor activities for specific policies, workflows, and actions.

The **Policies/Actions** pane at the top of the **Monitoring** window lists the policies on the NetWorker Server by default. Click the + (plus) sign next to a policy in the list to view the workflows in the policy, and the + (plus) sign next to a workflow to view the actions for a workflow.

The **Policies** pane provides the following information for each item (where applicable):

- Overall status  
The following table provides details on the status icons that may appear in the **Policies** pane.

**Table 85** Policy status icons

| Icon                                                                                | Status    |
|-------------------------------------------------------------------------------------|-----------|
|  | Never run |
|  | Running   |
|  | Succeeded |

**Table 85** Policy status icons (continued)

| Icon                                                                              | Status                                                |
|-----------------------------------------------------------------------------------|-------------------------------------------------------|
|  | Failed                                                |
|  | Probing                                               |
|  | Interrupted                                           |
|  | Queued                                                |
|  | Cloning                                               |
|  | Consolidating (NetWorker Server 8.2.x and lower only) |

 **Note:** When the schedule for an action is skip, the status of the action appears as Never Run and the status of the Workflow is Succeeded.

- Most recent start time.
- Duration of the most recent run.
- Next scheduled runtime.
- Name of the assigned save set.
- Device on which the save set is stored.
- Backup level.
- Data transfer rate.
- Size of the save set.
- Messages that resulted from an action.

Right-click an action in the **Policies** pane and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

When you sort the items on the **Policies/Actions** pane by using the **Status** column, NetWorker sorts the items in alphabetical order that is based on the label of the icon.

Consider the following when a policy/action is in a probing state:

- A message is sent when the group starts and finishes the probe operation.
- The results of the probe operation (run backup/do not run backup) are also logged.
- Probes do not affect the final status of the group, and the group status does not indicate the results of the probe.
- If probing indicates that a backup should not run, then the group status reverts to its state before the group running.
- Check the results of the probe in the **Log** window to ensure that the probe indicates that the backup can be taken.

### Actions pane

To view a list of all actions, click the **Actions** tab at the bottom of the **Policies** pane. The **Policies** pane becomes the **Actions** pane.

The **Actions** pane provides the following information for each action:



- Overall status

 **Note:** The **Actions** pane displays the same status icons as the **Policies** pane.

- Name
- Assigned policy
- Assigned workflow
- Type
- Date and time of the most recent run
- Duration of the most recent run
- Percent complete, for actions that are in progress
- Next scheduled runtime

Right-click an action in the **Actions** pane and select **Show Details** to view details on currently running, successfully completed, and failed activities for the action.

## Workflow operations

This section describes how to use the **Monitoring** window to start, stop, and restart workflows.

### Starting, stopping, and restarting policies

The workflows in a policy can run automatically, based on a schedule. You can also manually start, stop, and restart specific workflows by using the the NMC **NetWorker Administration Monitoring** window.

#### About this task

You can restart any failed or canceled workflow. Note, however, that the restart must occur within the restart window that you specified for the workflow. Additionally, for a VMware backup, if you cancel a workflow from **NetWorker Administration** and then want to restart the backup, ensure that you restart the workflow from the **NetWorker Administration** window. If a workflow that was started from **NetWorker Administration** is restarted from the **vSphere Web Client**, the backup fails.

#### Procedure

1. In the **Monitoring** window, select the workflow or actions.
2. Right-click and then select **Start**, **Stop**, or **Restart**.

A confirmation message appears.

 **Note:** You cannot stop, restart, or start individual actions.

3. Click **Yes**.

### Viewing workflow backup details

Perform the following steps to view backup details for workflows.

#### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Policies** in the docking panel, and expand the Policy that you want to monitor.
3. Right-click the workflow, and then select **Show Details**. The **Workflow Summary** window appears.
4. In the **Workflow runs** pane of the **Workflow Summary** window, select the workflow.
5. Click **Show Messages**. In the **Show Messages** window, select one of the following options:

- Get Full Log—To display all messages.
- Print—To print the log.
- Save—To save the log to a local file.
- OK—To close the **Show Messages** window.

6. Click **OK** to close the **Workflow Summary** window.

## Viewing action backup details

Perform the following steps to view backup details for actions.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Actions** in the docking panel.
3. In the **Actions** pane, right-click the action, and then select **Show Details**. The details window for the action appears.
4. Review the information in the **Actions Messages** pane. To display detailed information from the action log file, click **Show Action Logs**, and then select one of the following options:
  - Get Full Log—To display all messages.
  - Print—To print the log.
  - Save—To save the log to a local file.
  - OK—To close the **Show Messages** window.
5. In one of the Actions detail panes, for example, the **Completed successfully** pane, select the action that you want to review.
6. Click **Show Messages**. In the **Show Messages** window, select one of the following options:
  - Get Full Log—To display all messages.
  - Print—To print the log.
  - Save—To save the log to a local file.
  - OK—To close the **Show Messages** window.
7. Click **OK** to close the **Details** window.

## Sessions window

Use the **Sessions** window to view the sessions that are running on a NetWorker server. You can change the view of this window to display these sessions:

The **Sessions** pane below the **Policies/Actions** pane provides details on individual save, recover, clone, and synthetic full sessions by client.

To view all sessions or to limit the list of sessions by the session type, click the tabs at the bottom of the **Sessions** pane. Session types include:

- Save
- Recover
- Clone
- Browse
- Synthetic Full/Rehydrated Sessions
- All

To change the displayed session types go to **View > Show**, and select the type of sessions to display. To display all sessions currently running on the NetWorker Server, regardless of type, select **All Sessions**.

You can stop a session (backup, synthetic full backup, clone, and recovery sessions) from the **Monitoring** window, even if the session was started by running the `nsrworkflow` command.

To stop a session, right-click the workflow in the pane, and select **Stop** from the list box.

## Changing displayed session types

The column headings that are displayed on this window differ depending on the type of sessions you chose to display.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Sessions** in the docking panel.
3. Select **View > Show** and then select the type of sessions to display. To display all sessions currently running on the NetWorker server, regardless of type, select **All Sessions**.

## Stopping a session

You can stop a session (backup, synthetic full backup, clone, and recovery sessions) from the **Monitoring** window, even if the session was started by running `savegrp`.

To stop a session, right-click the session in the window and select **Stop** from the drop-down.

The following table provides a list of actions that can be stopped from NMC.

**Table 86** Sessions that can be stopped from NMC

| Session type                             | Stop from NMC? |
|------------------------------------------|----------------|
| Save by Savegroup                        | Yes            |
| Synthetic Full by Savegroup              | Yes            |
| Clone by Savegroup                       | Yes            |
| Schedule Clone                           | Yes            |
| Manual Save                              | No             |
| Manual Clone via NMC                     | No             |
| Manual Clone via CLI                     | No             |
| Winworker and CLI Recovery               | No             |
| Recovery started from Recover wizard     | Yes            |
| VMware Backup Appliance Save and Recover | No             |

 **NOTICE** Stopping a session from NMC does not affect any other group operations running.

## Alerts pane

The **Alerts** pane displays alerts that are generated by a particular NetWorker server or Data Domain system that has devices that are configured on the NetWorker server. The **Alerts** pane includes priority, category, time, and message information.

An icon represents the priority of the alert. The following table lists and describes each icon.

**Table 87** Alerts window icons

| Icon                                                                              | Label        | Description                                                                                                                           |
|-----------------------------------------------------------------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
|  | Alert        | Error condition detected by the NetWorker server that should be fixed by a qualified operator.                                        |
|  | Critical     | Severe error condition that demands immediate attention.                                                                              |
|  | Emergency    | Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority. |
|  | Information  | Information about the current state of the server. This icon represents the lowest priority.                                          |
|  | Notification | Important information.                                                                                                                |
|  | Waiting      | The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.                                           |
|  | Warning      | A non-fatal error has occurred.                                                                                                       |

When items on the **Alerts** pane are sorted by the **Priority** column, they are sorted in alphabetical order based on the label of the icon.

## Removing alerts

Remove individual alert messages from the **Events** tables by removing them from the **Events** table. To delete a message in the **Events** table, right-click the message, and select **Dismiss**.

### About this task

 **Note:** The alert message remains in the **Log** window in the NetWorker **Administration** program.

## Devices pane

The **Devices** pane allows you to monitor the status of all devices, including NDMP devices. If the NetWorker server uses shared and logical devices, the window is adjusted dynamically to present a set of columns appropriate for the current configuration.

The **Devices** pane provides the following information:

- Status of the operation.
- Name of the device.
- Name of the storage node that contains the device.
- For tape devices, the name of the library that contains the device.
- Name of the volume in the device.
- Name of the pool that is associated with the volume.
- Last message generated for the device.
- Whether the operation requires user input.

For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape.

[Entering user input](#) on page 89 provides instructions on how to deal with a user input notification.

If the current server configuration includes a shared device, a **Shared Device Name** column appears on the **Devices** pane. The name of the shared device appears in the **Shared Device Name** column. If other devices for that configuration are not shared devices, then the **Shared Device Name** column is blank for those devices. Only a single device per hardware ID can be active at any particular moment. The information for inactive shared devices is filtered out, and as a result, only one device per hardware ID is presented on the window at any time.

An icon represents the device status. The following table lists and describes each icon.

**Table 88** Devices status icons

| Icon                                                                                | Label                       | Description                         |
|-------------------------------------------------------------------------------------|-----------------------------|-------------------------------------|
|    | Library device active       | The library device is active.       |
|    | Library device disabled     | The library device is disabled.     |
|    | Library device idle         | The library device is idle.         |
|    | Stand-alone device active   | The stand-alone device is active.   |
|    | Stand-alone device disabled | The stand-alone device is disabled. |
|  | Stand-alone device idle     | The stand-alone device is idle.     |

When you sort items in the **Devices** pane by the **Status** column, NetWorker sorts the devices in alphabetical order based on the label name of the icon.

## Operations window

The **Operations** window displays information about device operations. It provides the following information:

- Status of the operation.
- Name of the library.
- Whether the operation requires user input.  
For example, a labeling operation may want the user to acknowledge whether the system should overwrite the label on a tape. [Entering user input](#) on page 89 provides instructions on how to deal with a user input notification.
- The origin, or source, of the operation.  
For example, the interface, nsrjb or the NetWorker server.
- Time the operation started.
- Type of operation.
- Duration of the operation.
- Status messages from the operation.
- Any error messages.

 **NOTICE** Only the last error message of the operation appears in the **Error Messages** column. Move the mouse pointer over the cell containing the last error message to display the entire list of error messages.

The operation status is represented by an icon. The following table lists and describes each of the icons.

**Table 89** Operations window icons

| Icon                                                                              | Label      | Description                                          |
|-----------------------------------------------------------------------------------|------------|------------------------------------------------------|
|  | Failed     | The operation failed.                                |
|  | Queued     | The operation is waiting in the queue to run.        |
|  | Retry      | The operation failed, but may work if you try again. |
|  | Running    | The operation is running.                            |
|  | Successful | The operation completed successfully.                |
|  | User Input | The operation requires user input.                   |

When items on the **Operations** window are sorted by the Status column, they are sorted in alphabetical order based on the label of the icon.

## Viewing operation details

The **Operation Details** dialog box opens, providing information about the completion of the operation. The **Completion Time** displays the time that the operation finished. The time that it took to complete the operation is the difference between the completion and start times of the operation.

### About this task

To save operation details to a file, click **Save** in the **Operation Details** dialog box. When prompted, identify a name and location for the file.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation, then select **Show Details**.

## Stopping an operation

Certain operations can be stopped from the **Operations** window.

### Procedure

1. From the **Administration** window, click **Monitoring**.
2. Click **Operations** in the docking panel.
3. Right-click the operation to stop, then select **Stop**.
4. Click **Yes** to confirm the stop.

 **Note:** Operations that were started from a command line program, such as the `nsrjb` command, cannot be stopped from the **Operations** window. To stop these operations, press `Ctrl-C` from the window where the command was started.

## Entering user input

If the system requires user input, select the labeling operation in slow/verbose mode and the **Supply User Input** icon appears.

### Procedure

1. Right-click the operation, then select **Supply Input**.
2. Confirm the requirement to supply input.
  - If **Yes**, and input is supplied, the icon in the **User Input** column disappears.
  - If **No**, and input is not supplied, the operation will time out and fail.

 **Note:** If two users try to respond to the same user input prompt, the input of the first user takes precedence, and the second user receives an error message.

## Log window

To view the most recent notification logs, click the **Log** window from the docking panel in the **Monitoring** window. The **Log** window provides the priority, time, source, category, and message for each log.

 **Note:** If a particular log file is no longer available, check the log file on the NetWorker server. The log files are located in `NetWorker_install_path\logs` directory.

An icon represents the priority of the log entry. The following table lists and describes each icon.

**Table 90** Icons in the Log pane

| Icon                                                                                | Label        | Description                                                                                                                           |
|-------------------------------------------------------------------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
|  | Alert        | Error condition that is detected by the NetWorker server that should be fixed by a qualified operator.                                |
|  | Critical     | Severe error condition that demands immediate attention.                                                                              |
|  | Emergency    | Condition exists that could cause NetWorker software to fail unless corrected immediately. This icon represents the highest priority. |
|  | Information  | Information about the current state of the server. This icon represents the lowest priority.                                          |
|  | Notification | Important information.                                                                                                                |
|  | Waiting      | The NetWorker server is waiting for an operator to perform a task, such as mounting a tape.                                           |
|  | Warning      | Non-fatal error has occurred.                                                                                                         |

When you sort items on the **Log** pane by using the **Priority** column, NetWorker sorts the icons in alphabetical order based on the name of the label.

## Recover window

The **Recover** window displays information about recover configurations that are created with the NetWorker Management Console (NMC) Recovery wizard.

You can use this window to:

- Start the NMC Recovery wizard to create recover configurations or modify saved recover configurations.
- Identify the status of a recover configuration that is created with the NMC Recovery wizard.
- Start and stop a recover job.

The **Recover** window is divided into five sections:

- **Toolbar**—The toolbar is hidden by default. To display the recovery toolbar, select **View > Show toolbar**
- **Summary**
- **Configured Recovers**
- **Currently Running**

A splitter separates the **Configured Recovers** section from **Currently running** window. You can click and move the splitter to resize these two windows.

## Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. The following table summarizes the function of each toolbar button.

**Table 91** Recovery toolbar options

| Button                                                                              | Function                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|   | Starts the NMC <b>Recover</b> wizard to create recover configurations.                                                                                                                        |
|  | Displays the <b>Properties</b> window for the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                          |
|  | Deletes the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                                                            |
|  | Displays online help for the <b>Recover</b> window.                                                                                                                                           |
|  | Displays the <b>Find</b> window at the bottom of the <b>Recover</b> window. The <b>Find</b> window allows you to perform keyword searches for messages that appear in the <b>Logs</b> window. |
|  | Start the recover operation for a selected saved recover configuration. This option is only available for a recover configuration that has a Never run, or Failed status.                     |
|  | Stop in-progress recover operation that you selected in the <b>Currently Running</b> window.                                                                                                  |

 **Note:** The **Recover** toolbar does not appear by default. To display the **Recover** toolbar, select **View > Show toolbar**.

## Recover Summary

The Recover Summary section displays a high-level overview of recover jobs.

This section includes the following information:



- Total Recovers—The total number of successful recover jobs.
- Since—The number of successful recover jobs since this date.

## Configured Recovers

The **Configured Recovers** window displays a list of saved recover configurations in a tabular format. You can sort the information by column. The **Configured Recovers** table displays the following information for each saved recover configuration:

- Status—The job status of a saved recover configuration.
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB.
- Comment
- OS—The operating system of the source host.
- Recover requestor—The Windows or UNIX account used to create the recover configuration.
- Start Time
- End Time
- Start date

**Table 92** Save recover configuration job status

| Icon                                                                                | Description                                        |
|-------------------------------------------------------------------------------------|----------------------------------------------------|
|  | The last recover attempt failed.                   |
|  | The last recover attempt completed successfully.   |
|  | The recover job has never run.                     |
|  | The recover job is scheduled to run in the future. |
|  | The recover job has expired.                       |

## Currently running

The **Currently Running** window displays a list of in progress recover jobs in a tabular format. You can sort the information by column. The **Currently Running** table displays the following information for each job:

- Status
- Name
- Source client
- Destination client
- Recovery list

- Recover type—For example, file system or BBB
- Volume
- Comment
- Device
- Size
- Total size
- % complete
- Rate (KB/s)
- Start time
- Duration
- Currently running

## Find

The **Find** section appears along the bottom of the **Recover** window, after you select the **Find** button on the **Recover** toolbar. **Find** allows you to search for keywords in the **Configured Recovers** window. The following table summarizes the available find options.

**Table 93** Find options

| Find option          | Description                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Find</b>          | Highlight the first saved recover configuration that contains the specified keyword.                                                         |
| <b>Prev</b>          | Highlight the previous saved recover configuration that contains the specified keyword.                                                      |
| <b>Highlight All</b> | Highlights each saved recover configuration that contains the specified keyword.                                                             |
| <b>Sort Selected</b> | Sorts each highlighted recover configuration in the Configured Recover table so that they appear at the top of the Configured Recover table. |
| <b>Match case</b>    | Make the keyword search case sensitive.                                                                                                      |

## Monitoring checkpoint-enabled backups

You can view detailed information about a checkpoint-enabled backup.

### About this task

#### Procedure

1. From the **Administration** window, select **Monitoring > Groups**.
2. Right-click the group to which the checkpoint enabled client belongs, then select **Show Details**.
3. View the detailed information related to the group backups:
  - If the partial save set is in the work list for the group, the save set appears in the **Waiting to Run** section.
  - If the partial save set is running, the save set appears in the **Currently Running** section.

- If the entire partial save sets sequence of the savegroup is complete, the save set appears in the **Completed Successfully** section.
- If the entire partial save sets sequence of the savegroup is not complete, the save set appears in the **Failed** section.

 **NOTICE** If any messages are generated, the Show Messages button is enabled. Click Show Messages to view the messages.

4. Click **OK** to close the **Group Details** window.

## Query the media database for partial save sets

The savegrp completion report does not provide detailed information about partial save sets that might be necessary to perform a recovery.

### Querying partial save sets from the Console

You can view information about the partial save sets by using the NetWorker Console.

#### About this task

 **NOTICE** If no partial save sets are found that match the query, ensure that the backup of the partial save sets was started within the Save Time period. To change the values for the Save Time attribute, open the Save Set Query tab and select a date and time from the Save Time calendar.

#### Procedure

1. In the **Administration** window, click **Media**. Media-related topics appear in the navigation tree.
2. Select **Save Sets**. The following tabs appear in the **Save Sets** window:
  - Query Save Set
  - Save Set List
3. Select the **Query Save Set** tab, to query:
  - All partial save sets, select **Checkpoint Enabled**.
  - All partial save sets with the same Checkpoint ID, in the **Checkpoint ID** field, type the **Checkpoint ID** of the partial save set on which you want to perform the query.
4. Select the **Save Set List** tab to view the result of the save set query:
  - The **Checkpoint ID** column displays the partial save set **Checkpoint ID** and its Sequence ID. The **Checkpoint ID** is listed first followed by the **Sequence ID**, which is encased within brackets.
  - Sort the **Checkpoint ID** column to view the complete sequence of partial save sets.
  - The **Status** column displays the status of the partial save sets:
    - A Checkpoint browsable status indicates that the save sets can be browsed for recover.
    - A Checkpoint aborted status indicates that the backup of the partial save set was stopped or aborted. A save set recover is used to recover the partial save set. Consider the following:
      - When a checkpoint-enabled backup completes successfully, the status of the last partial save set is Checkpoint browsable.
      - When a checkpoint-enabled backup completes successfully, on the first backup attempt, the save set status is Checkpoint browsable. Only one Sequence id is

associated with the Checkpoint ID. The Sequence id is 1. If the Sequence id is 2, the first partial save set in the checkpoint-enabled backup is missing.

## Querying partial save sets by using the `mminfo` command

By default, the `mminfo` command output only displays the browsable save sets. The first and intermediate partial save sets are not displayed. Only complete checkpoint-enabled save sets or final partial save sets are displayed.

Use the `mminfo` command with specific queries to display more information about checkpoint-enabled save sets.

The following table lists the new media attributes support the Checkpoint Restart feature.

**Table 94** New Checkpoint Restart media attributes

| Media attribute    | Description                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------|
| checkpoint_id      | Displays the checkpoint restart id of the partial save set in the <code>chkpt_id</code> column. |
| checkpoint_seq     | Displays the partial save set sequence id in the <code>chkpt_seq</code> column.                 |
| checkpoint-restart | This flag attribute is used to only display checkpoint restart enabled save sets.               |

Also, several media sumflags are used with the Checkpoint Restart feature:

- `k`—Indicates this is a checkpoint enabled save set.
- `a`—The first and all intermediate partial save sets of a checkpoint sequence has aborted status.
- `b`—The last partial or complete save set of a checkpoint sequence is marked browsable.

## Displaying checkpoint enabled save sets

Display all checkpoint enabled save sets by using the following `mminfo` command:

```
mminfo -q 'checkpoint-restart' -r 'client,nsavetime,ssid(11),
sumflags(3),name,checkpoint_id,checkpoint_seq'
```

**Table 95** Checkpoint enabled save sets

| client | save time  | ssid           | ssflags | filename | chkpt_id   | chkpt_seq |
|--------|------------|----------------|---------|----------|------------|-----------|
| plapew | 1251910303 | 420470031<br>9 | cak     | /space   | 1251910303 | 1         |
| plapew | 1251910327 | 4187923127     | cbk     | /space   | 1251910303 | 2         |
| plapew | 1251910710 | 408726021<br>4 | cak     | /space   | 1251910710 | 1         |
| plapew | 1251910725 | 407048301<br>3 | cbk     | /space   | 1251910710 | 2         |

## Displaying all partial save sets for the checkpoint id

Display all partial savesets for the checkpoint id by using the following `mminfo` command:

```
mminfo -q "checkpoint_id=1251910303"
```

**Table 96** Partial save sets for the checkpoint id

| volume     | client | date     | size   | level | name   |
|------------|--------|----------|--------|-------|--------|
| plapew.001 | plapew | 08/02/15 | 17 MB  | full  | /space |
| plapew.001 | plapew | 08/02/15 | 799 MB | full  | /space |

## Reporting NDMP Data

The NetWorker software reports information about NDMP clients, data, and volumes in two ways:

- The NetWorker Management Console (NMC) reporting feature—Reports NDMP data in the same manner as non-NDMP data. The *NetWorker Administration Guide* provides more information.
- The `mminfo` command. Use the `mminfo` program to query the media database for NDMP volume or save set information.

## Querying the NDMP volumes by backup type with the `mminfo` command

You can query save sets by backup format (NDMP or DSA) to display volume information.

For example:

- To query NDMP volumes, type `mminfo -q ndmp`. Output similar to the following appears:

```
volume client date size level name
005D0000 simlcifs1 6/22/2011 1036 MB full /fs1
005D0001 simlcifs1 6/22/2011 173 MB full /fs1
005D0001 simlcifs1 6/22/2011 862 MB full /fs1
005D0002 simlcifs1 6/22/2011 348 MB full /fs1
```

- To query NDMP -DSA volumes, type `mminfo -q dsa`. Output similar to the following appears:

```
volume client date size level name
NDMP.001 10.8.67.219 12/13/2011 644 MB full /vol/vol0
NDMP.001 10.8.67.219 12/13/2011 402 MB full /vol/vol1
NDMP.001 10.8.67.219 12/13/2011 402 MB full /vol/vol1
NDMP.001 10.8.67.219 12/13/2011 36 MB full /vol/vol2
```

## Querying the NDMP save sets with the `mminfo` command

To determine which save sets are Network Data Management Protocol (NDMP) save sets and the status of an NDMP save set in the media database, query the media database. NDMP save set status information is important when performing NDMP recoveries:

- To perform a browsable NDMP recover, the `ssflags (f1)` field for an NDMP save set must contain a (b). The b value denotes a browsable save set.
- To perform a save set recover from the NetWorker User program, the `ssflags (f1)` field for an NDMP save set must contain either (r) or (b).
- An NDMP save set contains an N attribute in the `ssflags (f1)` field.
- An NDMP-DSA save set contains an s attribute in the `ssflags (f1)` field.

In the following example, the NDMP save set status is recoverable (r). To recover the data, you can only perform a save set recovery from a command line.

```
mminfo -av
```

```
volume type client date time size ssid fl lvl name
vol1 dlt clnt 6/22/2011 3:15:12 1036MB 3842140553 hrN full /fs1
```

In the following example, the NDMP-DSA save set status is browsable (b). Recover the data by using the NetWorker User program, or from the command line. A browsable NDMP-DSA save set supports browsable and save set recoveries.

```
mminfo -av
```

```
volume type client date time size ssid fl lvl name
vol1 dlt clnt 6/22/2011 3:15:12 36MB 4259813785 cbNs full /fs1
```

## Performing NDMP recoveries

NetWorker uses the `nsrndmp_recover` program to coordinate recover operations between the NetWorker software and the Network Data Management Protocol (NDMP) client. The `nsrndmp_recover` program does not move data to the NDMP client. When the `nsrndmp_recover` program identifies an NDMP-DSA save set, `nsrndmp_recover` automatically runs the `nsrdsa_recover` program on the same host that runs the `nsrndmp_recover` command.

To recover NDMP data, you can run the `nsrndmp_recover` program from a command prompt, or use one of following programs, which automatically starts `nsrndmp_recover`:

- `recover`—The command line program on Windows and UNIX.
- `winworkr`—The NetWorker User GUI on Windows.
- The NMC Recovery wizard.

During the recovery process, the `nsrndmp_recover` program passes `nlist` information to the NDMP client. There are three methods to recover NDMP backups:

- Index-based file-by-file recover—The `nlist` includes file offset and ACL information. When you recover many files, the recover process uses a significant amount of system resources on both the NetWorker server and the NDMP client to build and process the `nlist` information.
- Full save set recovery—The `nlist` only includes the path to the recovery directory, down to and including the mount point. When you recover many files, the recover process uses less system resources than an index-based NDMP recover to build and process the `nlist` information.
- NDMP directory restore—A partial save set recovery of a single file or single directory.

For example, when the NetWorker software writes NDMP data to a remote storage node, start the `recover` program on the NetWorker storage node to prevent the data from traversing the network.

**Note:** When you start the `recover` program on the NetWorker server, the data flows from the storage node to the NetWorker server and from the NetWorker server to the NDMP client, over the network.

## NDMP recovery requirements for other filers

The following list summarizes the requirements:

### scanner

You cannot use the `scanner` command with the `-i`, `-f` and `-r` options on an NDMP volume. You cannot use the `scanner` command on a volume that contains NDMP and non-NDMP save sets when you load the volume in an NDMP device. The *Scanner command usage*

technical note provides more information about using the `scanner` command with NDMP data.

### Cross platform recoveries

You can recover NDMP data to different NDMP client however, you cannot perform a cross platform recover. Recover NDMP data to an NDMP client that is the same brand, a compatible model, and the same operating system as the original NDMP client.

### Devices

Recover Direct-NDMP and Three-party backups performed to an NDMP device from an NDMP device. To improve recover performance from an NDMP tape device, configure the tape device to support variable length records. Recover NDMP-DSA backups from a non-NDMP device.

### Localized environments

When recovering data in a localized NDMP environment, the Index Recover status window shows the process in English and not the localized language.

### NDMP-DSA

For better recovery performance, start the recover process on the NetWorker host where the backup volume resides.

### Immediate recoveries

Run the `nsrndmp_recover` program on the storage node with the locally attached backup device to perform an immediate recovery of NDMP-DSA data.

### Blue Arc

The recover process creates a `$__NDMP__` directory at the root level of the recovery file system when you recover more than 1,024 files. The directory contains the file list that the NetWorker server uses for an index recovery. Do not change the directory and its contents during an active recovery operation. When a recovery is not in progress, you can delete the directory.

While performing NDMP backup and recover operations, a message similar to the following might appear:

```
NDMP session-Unknown environment variable name ignored.
```

You can ignore this message.

### Mirapoint

After a full backup recovery for a Mirapoint system, reboot the Mirapoint system. An incremental recovery does not require a restart.

## DAR and DDAR

By default, the Network Data Management Protocol (NDMP) recover process reads an entire tape from start to finish. The recover process extracts the data as it encounters the data on the tape. For large backup images, recovery is slow.

The Direct Access Recovery (DAR) and Directory DAR (DDAR) recovery process:

- Provides the ability to recover a file or directory from the exact location on a tape.
- DDAR only passes the directory path to the NAS filer. DAR passes the paths of each file individually.

- Reduces the size of the nlist information that the recover process stores in memory. During the recover process, the NAS filer (DDAR) assumes that the directory path includes all cataloged files and directories. However, DAR mentions each file that it wants recovered.
- Does not sequentially read the file or record numbers on the tape to locate the data, which reduces the amount of time that you require to recover specific files from a backup.

 **Note:** [Creating and configuring the NDMP client resource](#) on page 67 describes how to configure the DAR and DDAR Application Information attributes for NDMP clients.

#### When not to use DAR or DDAR

DAR and DDAR recoveries send multiple pathnames across the network to the NDMP Data Server and, in three-party configurations, to the NetWorker server. The recover process stores the pathnames in memory on the NDMP Data Server. Recoveries of a large amount of data from a large save set can negatively affect the network and the NDMP Data Server resources.

Do not use DAR and DDAR to recover the following objects:

- Several thousands of files in a single index-based recover operation.
- A specific directory structure containing several thousand or millions of files.

To perform a non-DAR-based recovery of a save set when you set the *DIRECT=y* at the time of backup, first define the *NSR\_NDMP\_RECOVER\_NO\_DAR=y* variable in the Application Information attribute of the NDMP client.

## Recover window

The **Recover** window displays information about recover configurations that are created with the NetWorker Management Console (NMC) Recovery wizard.

You can use this window to:

- Start the NMC Recovery wizard to create recover configurations or modify saved recover configurations.
- Identify the status of a recover configuration that is created with the NMC Recovery wizard.
- Start and stop a recover job.

The **Recover** window is divided into five sections:

- **Toolbar**—The toolbar is hidden by default. To display the recovery toolbar, select **View > Show toolbar**
- **Summary**
- **Configured Recovers**
- **Currently Running**

A splitter separates the **Configured Recovers** section from **Currently running** window. You can click and move the splitter to resize these two windows.

## Recover toolbar

The Recover toolbar provides you with the ability to quickly perform common recover operations. The following table summarizes the function of each toolbar button.

**Table 97** Recovery toolbar options

| Button                                                                              | Function                                                               |
|-------------------------------------------------------------------------------------|------------------------------------------------------------------------|
|  | Starts the NMC <b>Recover</b> wizard to create recover configurations. |



**Table 97** Recovery toolbar options (continued)

| Button                                                                            | Function                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Displays the <b>Properties</b> window for the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                          |
|  | Deletes the saved recover configuration that you selected in the <b>Configured Recover</b> window.                                                                                            |
|  | Displays online help for the <b>Recover</b> window.                                                                                                                                           |
|  | Displays the <b>Find</b> window at the bottom of the <b>Recover</b> window. The <b>Find</b> window allows you to perform keyword searches for messages that appear in the <b>Logs</b> window. |
|  | Start the recover operation for a selected saved recover configuration. This option is only available for a recover configuration that has a Never run, or Failed status.                     |
|  | Stop in-progress recover operation that you selected in the <b>Currently Running</b> window.                                                                                                  |

 **Note:** The **Recover** toolbar does not appear by default. To display the **Recover** toolbar, select **View > Show toolbar**.

## Recover Summary

The Recover Summary section displays a high-level overview of recover jobs.

This section includes the following information:

- Total Recovers—The total number of successful recover jobs.
- Since—The number of successful recover jobs since this date.

## Configured Recovers

The **Configured Recovers** window displays a list of saved recover configurations in a tabular format. You can sort the information by column. The **Configured Recovers** table displays the following information for each saved recover configuration:

- Status—The job status of a saved recover configuration.
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB.
- Comment
- OS—The operating system of the source host.
- Recover requestor—The Windows or UNIX account used to create the recover configuration.
- Start Time
- End Time

- Start date

**Table 98** Save recover configuration job status

| Icon                                                                              | Description                                        |
|-----------------------------------------------------------------------------------|----------------------------------------------------|
|  | The last recover attempt failed.                   |
|  | The last recover attempt completed successfully.   |
|  | The recover job has never run.                     |
|  | The recover job is scheduled to run in the future. |
|  | The recover job has expired.                       |

## Currently running

The **Currently Running** window displays a list of in progress recover jobs in a tabular format. You can sort the information by column. The **Currently Running** table displays the following information for each job:

- Status
- Name
- Source client
- Destination client
- Recovery list
- Recover type—For example, file system or BBB
- Volume
- Comment
- Device
- Size
- Total size
- % complete
- Rate (KB/s)
- Start time
- Duration
- Currently running

## Find

The **Find** section appears along the bottom of the **Recover** window, after you select the **Find** button on the **Recover** toolbar. **Find** allows you to search for keywords in the **Configured Recovers** window. The following table summarizes the available find options.

**Table 99** Find options

| Find option          | Description                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Find</b>          | Highlight the first saved recover configuration that contains the specified keyword.                                                         |
| <b>Prev</b>          | Highlight the previous saved recover configuration that contains the specified keyword.                                                      |
| <b>Highlight All</b> | Highlights each saved recover configuration that contains the specified keyword.                                                             |
| <b>Sort Selected</b> | Sorts each highlighted recover configuration in the Configured Recover table so that they appear at the top of the Configured Recover table. |
| <b>Match case</b>    | Make the keyword search case sensitive.                                                                                                      |

## Performing an NDMP index-based file-by-file data recovery

Perform an NDMP index based file-by-file recover in the same manner as a non-NDMP data recover. You can restore the data to the original NDMP client or perform a directed recovery to a different NDMP client.

Before you perform an index-based file-by-file recover, review the following information:

- Set the `HIST=y` in the Application Information attribute of the NDMP client at the time of the backup.
- The NDMP save set must be browsable. You cannot perform a browsable recover of a recoverable or recyclable save set. [Reporting NDMP Data](#) on page 93 describes how to determine the status of an NDMP save set.
- Do not use an index-based recovery to recover a large numbers of files or directories. For better recovery performance, use a save set recover. [Performing a Full or Directory Restore of NDMP data by using a save set recovery](#) on page 103 provides more information.
- To perform an index-based file-by-file recover:
  - Use the NetWorker User program on a Windows host. [Performing an NDMP index-based file-by-file recover using the NetWorker User program](#) on page 100 provides detailed information.
  - Use the `recover` program. [Performing an NDMP index-based file-by-file recover from a command prompt](#) on page 102 provides detailed information.

### Performing an NDMP index-based file-by-file recover using the NetWorker User program

On Windows, to recover data to the original NDMP client or to a different NDMP client, perform the following steps.

#### Procedure

1. Open the NetWorker User program and connect to the NetWorker server.

**i NOTICE** If you receive the error:

No file indexes were found for client `client_name` on server `server_name`

Try connecting to a different NetWorker server and you selected the correct NetWorker server, then ensure that you selected a browsable save set. Alternatively, perform a save set recover.

2. Select **Recover** to open the **Source Client** window.
3. Select source NDMP client and click **OK**. The local client is the default selection.
4. Select the destination client for the recovered data and click **OK**. If the destination client is not the source client, ensure the NAS filer is the same brand, a compatible model and the same operating system as the source NDMP client.
5. (Optional) Recover the data from an earlier backup time. The **Recover** window appears with the latest version of the backup files. To recover data from an earlier backup, change the date and time of backup using one of the following methods:
  - a. Change the browse time for all files in the recover window:
    - From the **View** menu, select **Change Browse Time**.
    - In the **Change Browse Time** window, select a new day within the calendar. Select **Previous Month** or **Next Month** to change from the current month.
    - In the **Time** field, change the time of day by typing an hour, a minute, and the letter a for A.M. or p for P.M. Use the 12-hour format.
    - Click **OK**.
  - b. View all versions of the selected file system object:
    - Highlight the file or directory for review.
    - From the **View** menu select **Versions**.
    - Once you locate the version to recover, change the browse time. To change the browse time, highlight the volume, directory, or file and click **Change Browse Time**. The **Version** window closes and the **Recover** window reflects the new browse time.
6. (Optional) Search for the files. To search for and recover the most recently backed-up version of a file or directory:
  - a. From the **File** menu, select **Find**.
  - b. Type the name of the file or directory. Use wildcards to expand the search; without wildcards, partial filenames do not provide any results.
7. Mark the data to recover. To select file system objects to recover:
  - a. In the left pane of the **Recover** window, click the appropriate directory folder.
  - b. Mark each directory or file to recover by selecting the checkbox next to each directory or file.
8. (Optional) Relocate the data to a different location. By default, the recover process recovers the selected files to the original location.

**i Note:** The NDMP protocol does not support name conflict resolutions. NetWorker will always overwrite existing files that have the same name as the recovered file. It is recommended that you recover the NDMP data to a different location, to avoid data loss.

To relocate the files to a different location:

- a. Select **Recover Options** from the **Options** menu.

NDMP recoveries do not support the following options:

- Rename recovered file
- Discard recovered file
- Prompt for every file conflict

NDMP recoveries will always overwrite existing files. It is recommended that you relocate the NDMP data to a different location, to avoid data loss.

- b. In the **Relocate Recovered Data To** field, type the full path name of the target directory, click **OK**.

The target directory is a literal string and must match the path as seen by the NAS filer in its native OS, exactly. Otherwise, the recover process uses the original location and overwrites existing files with the same name.

9. (Optional) To view the volumes required to recover the marked file system objects, from the **View** menu, select **Required Volumes**.
10. Click **Start** to begin the recovery. If any required volume is not available to the NetWorker server, a volume status warning appears.

When this warning appears:

- a. Click **No**.
- b. From the **View** menu, select **Required Volumes**.
- c. Ensure that the NetWorker software can mount each listed volume into an available device.
- d. Attempt the recover operation again.

The NetWorker server takes a few moments to recover the files, depending on file size, network traffic, server load, and tape positioning. During this time, messages appear so that you can monitor the progress of the recovery.

When the recovery completes successfully, a message similar to the following appears:

```
Received 1 file(S) from NSR server
server Recover completion time: Tue Jan 21 08:33:04 2009
```

## Performing an NDMP index-based file-by-file recover from a command prompt

This section applies to command line recoveries from a Windows and UNIX client.

### About this task

To avoid using the Windows version of `recover.exe` on Windows operating systems, perform one of the following actions:

- Specify the full path to the recover program. For example: `C:\Program Files\EMC NetWorker\nsr\bin\recover.exe`
- Ensure that the `$PATH` environment variable contains the `NetWorker_install_path\bin` directory before `%SystemRoot%\System32`

To recover Network Data Management Protocol (NDMP) data from a command prompt on a UNIX or Windows NetWorker host, perform the following steps.

## Procedure

1. From the command prompt, type:

```
recover -s NetWorker_servername -c client_name
```

where:

- *-s NetWorker\_servername* specifies a particular NetWorker server on the network to use when recovering data.

When you do not use the *-s* option, the `recover` program tries to connect to the first computer listed in the servers file. When the servers file does not contain any servers, or lists more than one server, the **Change Server** window appears, and you can select the server.

- *-c client\_name* specifies the source NDMP client.

2. When prompted, type the directory to browse, for example:

```
cd /mydirectory
```

3. Use the `add` command to add the required files or folders to the recover list. The *NetWorker Command Reference Guide* provides a complete list of options for the `recover` command.
4. When restoring NDMP data, it is recommended that you relocate the NDMP data to a different location.

**Note:** The NDMP protocol does not support name conflict resolutions. NetWorker will always overwrite existing files that have the same name as the recovered file. It is recommended that you recover the NDMP data to a different location, to avoid data loss.

- To relocate the data to a different directory, type:

```
relocate destination_directory_name
```

The target pathname for *destination\_directory\_name* is a literal string and must match the path as seen by the NAS filer in its native OS, exactly. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

- To recover the data to a different host, type:

```
relocate target_hostname::/mount_point
```

Data ONTAP may require you to add a backslash (\) after the mount point. For example, *target\_hostname::\mount\_point\*.

5. After you add all of the required files, type:

```
recover
```

## Performing a Full or Directory Restore of NDMP data by using a save set recovery

You perform a Network Data Management Protocol (NDMP) save set recover in the same manner as a non-NDMP save set recovery. You can recover data to the original NDMP client or perform a directed recovery of the data to a different NDMP client of the same platform.

Before you perform a full save set recover, review the following information:

- Use a full save set recovery to recover all files and folders in an NDMP data save set, or to recover an entire directory within an NDMP save set. You cannot use the NetWorker User program to perform an NDMP Directory Restore.
- To use the NetWorker User program on Windows, a client file index entry for the save set must exist. When the index entry for the save set does not exist, the recover fails with an `index not found` error. When the client file index entries do not exist for the save set, use the `nsrndmp_recover` program with the `-v off` option.
- You cannot perform a save set recover from the NetWorker User program when the save set status is eligible for recycling (E). The recover process requires a recoverable (r) or browsable (b) save set status. The *NetWorker Administration Guide* provides information on how to change the status of a save set. A save set recover reads the entire tape, from beginning to end, to find and recover the requested files. The recovery process completes when the recover operations reads all required tapes in their entirety.
- As each file recovers, the file name appears on the target share but the file size is 0 KB. The actual file size update occurs after the recovery completes.

## Performing an NDMP save set recover by using the NetWorker User in Windows

### About this task

**NOTICE** When the recover operations fails with the error:

```
Failed to propagate handle <number> to child process: Access is denied
```

The save set is not in the client file index of the NDMP client. Perform a save set recover from a command prompt. [Performing an NDMP save set recovery from the command prompt](#) on page 105 provides more information.

### Procedure

1. Start the NetWorker User program.
2. On the **Change Server** window, select the NetWorker server and click **OK**.
3. Select **Options > Recover Save Sets**.
4. On the **Source Client** window, select the appropriate NDMP client, and then click **OK**.
5. On the **Save Sets** window, select the name of the save set.
6. Select the version of the save set, if there are multiple versions. You can also select the cloned version of a save set, if applicable.
7. To recover specific files and directories instead of the entire save set:
  - a. Click **Files**.
  - b. Specify the files and directories, one per line.
  - c. Click **OK**.

**NOTICE** Do not use this method to mark tens of thousands of files. Instead, perform an NDMP Directory Restore. Marking many files and directories generates a large nlist and requires intensive resources on both the NetWorker server and the NAS filer.

8. Click **Recover Options**.

An NDMP data recovery does not support the following options:

- Rename recovered file

- Discard recovered file
- Prompt for every file conflict

 **NOTICE** It is recommended that you relocate the NDMP data to a different location. NDMP recoveries always overwrite existing files.

9. To recover the data to a pathname that is different from the original backup location, in the **Relocate Recovered Data To** field, type the full pathname of the destination directory, then click **Ok**.

For NDMP data recoveries, the target pathname is a literal string and must exactly match the path as seen by the native OS on the NAS filer. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

10. To recover the data to a different NDMP client, specify the name of the client to receive the NDMP data in the **Destination Client** field.
11. To view the volumes that are required to perform the recover, select **View > Required Volumes**
12. Click **OK** to begin the recovery. The recovery status appears in the **Recover Status** window.

## Performing an NDMP save set recovery from the command prompt

To perform a save set recovery to the original Network Data Management Protocol (NDMP) client or to a different NDMP client, use the `nsrndmp_recover` command.

### About this task

For example:

```
nsrndmp_recover -s NetWorker_server -c source_ndmp_client -S ssid/cloneid -v off
-m target_ndmp_client::/target_path /source_path
```

where:

- `source_ndmp_client` is the hostname of the source NDMP client.
- `target_ndmp_client` is the hostname of the destination NDMP client.
- `/source_path` is the original location of the data.
- `/target_path` is the location to recover the data.

 **NOTICE** It is recommended that you relocate the NDMP data to a different location. NDMP recoveries always overwrite existing files. The `/target_path` is a literal string and must exactly match the path as seen by the native OS on the NAS filer. Otherwise, the recover operation uses the original location and overwrites existing files with the same name.

- `-v off` allows you to restore data when client file index of the NDMP client does not contain information about the NDMP save set.

In the following examples, the NetWorker server is mars and the backup client is venus:

- To recover a mount point `/mnt` from a backup of NDMP host venus to a directory `/newmnt` on NDMP host jupiter, type:

```
nsrndmp_recover -s mars -c venus -S 123456789 -v off -m jupiter::/newmnt
```

- To recover a mount point `/mnt` from a backup of NDMP host venus to NDMP host pluto, type:



```
nsrndmp_recover -s mars -c venus -R pluto -S 123456789 -v off -m /mnt
```

Data ONTAP may require that you to add a slash (/) after the mount point. For example, *target\_hostname:/mount\_point/*.

## Troubleshooting NDMP recover

This section provides a list of the possible causes and the possible resolutions for NDMP recovery issues.

### RESTORE: could not create path *pathname*

This error message appears when restoring NetApp data. This error, when encountered, appears in the `daemon.raw` file of the NetWorker server and the recovery output.

To resolve this issue:

- Ensure that you specify a source and a target path during the recover that exists on the target filer.
- If you set the `UTF8=Y` application information variable during an NDMP client backup and the backup contains path names with non-ASCII characters, then perform a save set recover. Index-based recoveries will fail with this error message.

### These files were not restored (Restore failed with error, or file/directory specified but not found in backup)

This error message appears in the `daemon.raw` file of the NetWorker server and the in the recovery output.

To resolve this issue:

- Ensure that the file or directory specified during the recover, exists in the save set.
- Ensure that the pathname specified to relocate the data exists on the destination filer. For NDMP data recoveries, the target pathname is a literal string and must exactly match the path as seen by the native OS on the NAS filer.

Other filers