

Dell EMC NetWorker

Version 19.1

Snapshot Management Integration Guide

302-005-698

REV 01

Copyright © 2001-2019 Dell Inc. or its subsidiaries. All rights reserved.

Published May 2019

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

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

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

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

CONTENTS

Figures		9
Tables		11
Preface		13
Chapter 1	Overview of NetWorker Snapshot Features	19
	NetWorker Snapshot Management product description.....	20
	Snapshot operations.....	20
	Types of snapshot backups.....	21
	Types of snapshot recoveries.....	22
	NetWorker clone support.....	22
	Backup configuration methods.....	22
	Restore methods.....	23
	Monitoring and reporting snapshot operations.....	23
	Internationalization support.....	23
	Components of the snapshot environment.....	24
	Application host.....	24
	FC and iSCSI environments.....	24
	Storage arrays.....	24
	NetWorker server.....	25
	NetWorker storage node.....	25
	Snapshot mount host.....	25
	Backup storage media.....	25
	NetWorker application modules.....	25
	Third-party volume managers.....	26
	NetWorker snapshot licensing requirements.....	26
	Example NetWorker snapshot environments.....	26
	Example of a snapshot and clone to storage media.....	27
	Example of a restore from a snapshot backup.....	29
Chapter 2	Data Protection Policies	31
	Default data protection policies in NMC's NetWorker Administration window	32
	Strategies for storage array snapshot backups.....	32
	Overview of configuring a new data protection policy.....	33
	Creating a policy.....	34
	Create a workflow for a new policy in NetWorker Administration.....	36
	Protection groups for snapshot backups.....	38
	Creating a basic client group.....	38
	Creating a dynamic client group.....	39
	Creating a save set group.....	39
	Creating a query group.....	40
	Actions supported in snapshot backups.....	42
	Supported actions in snapshot workflows.....	43
	Creating a check connectivity action.....	43

	Creating a probe action.....	46
	Creating a snapshot backup action.....	50
	Creating a clone action.....	53
	Visual representation of snapshot workflows.....	59
Chapter 3	Software Configuration	61
	Backup group resource migration.....	62
	Roadmap for snapshot configurations.....	64
	Snapshot configuration prerequisites.....	65
	Storage array specific prerequisites.....	65
	Application host prerequisites.....	66
	Mount host prerequisites.....	66
	Storage node prerequisites.....	67
	Configuring the user privileges.....	67
	Configuring snapshot backups with the client wizard.....	68
	Configuring snapshot backups manually.....	74
	Configuring the Client resource manually for the application host....	74
	Configuring the Client resource manually for a mount host.....	76
	Configuring the Application Information variables.....	76
	Configuring preprocessing and postprocessing scripts.....	76
Chapter 4	Configuring ProtectPoint on VMAX	79
	Overview.....	80
	ProtectPoint on VMAX3 prerequisites.....	80
	Configuring ProtectPoint.....	80
	Enabling vDisk on a Data Domain system.....	81
	Provisioning protection devices on Data Domain systems.....	81
	Completing the VMAX system configuration.....	82
	Considerations for ProtectPoint device and NetWorker ProtectPoint enabled pools.....	85
	Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard.....	85
	VMAX3 SRDF/S support.....	87
	Rollbacks in the SRDF/S environment.....	88
	Configuring Data Domain NsrSnapSG device groups for intelligent pairing....	88
	Intelligent Pairing vDisk selection decision tree.....	88
	Intelligent Pairing allocates vDisk for mount, validate, and restore....	89
Chapter 5	Configuring ProtectPoint on RecoverPoint with XtremIO	91
	Overview.....	92
	Basic backup workflow.....	92
	Basic restore workflow.....	92
	ProtectPoint for RecoverPoint on XtremIO prerequisites.....	93
	Enabling vdisk on the Data Domain.....	94
	Provisioning protection devices on Data Domain systems.....	95
	Configuring RecoverPoint and XtremIO storage.....	96
	Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard.....	97
	Configuration for restore to secondary VMAX with the wizard.....	99
	Considerations for ProtectPoint device and NetWorker ProtectPoint enabled pools.....	100

	Configuring Data Domain NsrSnapSG device groups for intelligent pairing....	100
	Intelligent Pairing vDisk selection decision tree.....	101
	Intelligent Pairing allocates vDisk for mount, validate, and restore....	101
Chapter 6	Configuring snapshots on XtremIO arrays	103
	Snapshot support for XtremIO.....	104
	Snapshot operation with XtremIO REST API.....	104
	Prerequisite for XtremIO configurations.....	104
	Supported XtremIO features.....	104
	Snapshot management policy with XtremIO.....	104
	Snapshot backups with XtremIO.....	105
	XtremIO configuration methods.....	105
Chapter 7	Configuring snapshots on VMAX Storage Arrays	107
	Snapshot support of VMAX storage arrays.....	108
	Snapshot operations with TimeFinder software.....	108
	Prerequisites and support for VMAX configurations.....	108
	Types of supported mirror devices.....	108
	Pairing source LUNs to mirror LUNs.....	109
	Intelligent pairing.....	109
	Configuring NsrSnapSG storage groups for intelligent pairing.....	110
	Manual pairing LUNs with the symm.res file.....	110
	Configuring the symm.res file.....	111
	VMAX SRDF/S support.....	112
	Rollbacks in the SRDF/S environment.....	112
	Solutions Enabler Client and Server mode configuration.....	113
	Solutions Enabler in Client and Server mode configuration.....	113
	Known limitation for VMAX.....	114
Chapter 8	Configuring snapshots on VNX Block Storage Arrays	115
	Snapshot support of VNX Block storage arrays.....	116
	Snapshot operations with SnapView software.....	116
	Prerequisites and support for VNX configurations.....	116
	Configuring the Navisphere security file.....	116
	Creating the Navisphere file manually on UNIX systems.....	117
	Creating the Navisphere file manually on Windows systems.....	117
	Configuring Unisphere CLI on VNXe3200.....	118
	UEMCLI Windows registry setup.....	118
Chapter 9	Configuring snapshots on RecoverPoint	121
	Snapshot support of RecoverPoint.....	122
	Snapshot operations with RecoverPoint software.....	122
	Prerequisite for RecoverPoint configurations.....	123
	Restrictions for RecoverPoint configurations.....	123
	Supported RecoverPoint features.....	123
	Snapshot management policy.....	123
	RecoverPoint configuration methods.....	124
	RecoverPoint snapshot retention.....	124
Chapter 10	Configuring snapshots in a Cluster Environment	127

	NetWorker support of cluster environments.....	128
	Failover with snapshots in a cluster environment.....	128
	Configuring a cluster environment for snapshots.....	128
	AIX systems in a cluster environment.....	129
	ProtectPoint restore and rollback for VCS on Solaris.....	130
	Performing a ProtectPoint VCS restore.....	130
	Performing a ProtectPoint VCS rollback.....	132
Chapter 11	Data Management and Recovery	135
	Snapshot lifecycle management.....	136
	Management and recovery of file system snapshot data.....	136
	Save set IDs and expiration policies.....	136
	Browsing snapshot and clone save sets.....	136
	Snapshot recovery support and limitations.....	137
	Raw partitions and raw devices.....	137
	NetApp restore fails	137
	Restoring from a snapshot with the Recovery Wizard.....	138
	Restoring a snapshot by rollback.....	141
	Rollback considerations.....	141
	Configuring the <code>psrollback.res</code> file.....	144
	Rollbacks with Veritas Volume Manager.....	145
	Rollbacks with IBM AIX Volume Manager.....	145
Chapter 12	Troubleshooting	147
	NetWorker snapshot backup issues.....	148
	Snapshot backup on Unity fails.....	148
	NAS Isilon snapshot mount fails on Linux	148
	Backup on Windows fails with a Delayed Write Failed error.....	148
	Backup fails and hangs when NMC user has insufficient privileges....	148
	Snapshots fail to mount for AIX managed file systems.....	149
	Snapshots fail for Linux Volume Manager on VNX with PowerPath...	149
	Linux Logical Volume Manager snapshots fail with an error.....	149
	NetWorker to Media-Clone stops responding and the backup fails....	149
	NetWorker snapshot restore issues.....	150
	File-by-file or saveset restore fails.....	150
	Restore of raw devices fails on Linux with permission issue.....	150
	Command <code>nrsnap_recover -l</code> runs but fails to restore a file.....	150
	Restore fails with disk signature error.....	150
	Directed restore files and folder permission issue.....	151
	Snapshot mount might fail because VMAX does not release the lock	
	on Restore FTS LUN.....	151
	NSM with XtremIO leaves snapshots mounted	151
Appendix A	Application Information Variables	153
	Using Application Information variables.....	154
	Common Application Information variables.....	154
	Application Information variables for VMAX arrays.....	155
	Application Information variables for VNX Block arrays.....	158
	Application Information variables for RecoverPoint appliances.....	159
	Application Information variables for XtremIO arrays.....	159

Appendix B	Command-Line Operations for Snapshot Management	161
	Using CLI commands for snapshot operations.....	162
	Using nsrsnapadmin for snapshot operations.....	162
	Example nsrsnapadmin operations.....	163
	Querying snapshot save sets.....	163
	File-by-file browsing and restore.....	163
	Rollback restore.....	163
	Deleting a snapshot save set.....	164
	Modifying the retention period of a snapshot save set.....	164
	Querying with the <code>mminfo</code> command.....	165
Appendix C	Migrating Legacy PowerSnap Configurations	167
	Migrating legacy PowerSnap configurations to NSM.....	168
	Removing PowerSnap on UNIX systems.....	168
	Removing PowerSnap on Microsoft Windows systems.....	168
	Deprecated Client resource attributes.....	169
	Migrating VMAX (Symmetrix) arrays.....	169
	Migrating VNX (CLARiiON) arrays.....	169
	Migrating RecoverPoint appliances.....	169
	Starting the nsrpsd process.....	170
	Licensing.....	170
Glossary		171

CONTENTS

FIGURES

1	Snapshot and clone operation with the storage node as the mount host.....	27
2	Snapshot and clone operation with the application host as the mount host.....	27
3	Restore from a snapshot with the storage node as the mount host.....	29
4	Platinum policy configuration.....	32
5	Data protection policy example.....	34
6	All possible workflow actions for a snapshot backup.....	43
7	Sample snapshot workflow.....	59
8	Snapshot and clone in a cluster environment.....	128

FIGURES

TABLES

1	Revision history.....	13
2	Style conventions.....	15
3	Save set criteria.....	41
4	Schedule icons.....	44
5	Schedule icons.....	47
6	Backup type icons.....	51
7	Schedule icons.....	54
8	Migration of Group attributes.....	62
9	vDisk object hierarchy mapping.....	81
10	vdisk object hierarchy mapping.....	95
11	Common Application Information variables.....	154
12	Application Information variables for VMAX arrays.....	156
13	Application Information variables for VNX Block arrays.....	158
14	Application Information variables for RecoverPoint appliances.....	159
15	Application Information variables for XtremIO arrays.....	159
16	Commands and options supported in nrsnapadmin interactive mode.....	162

TABLES

Preface

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

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

Note

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

Purpose

This document provides planning, practices, and configuration information for the use of the NetWorker Snapshot Management features within a NetWorker backup and storage management environment.

Audience

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

- Identifying the different hardware and software components that make up the NetWorker datazone.
- Configuring storage management operations by following procedures.
- Locating problems and implement solutions by following guidelines .

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

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

Related documentation

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

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

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

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

Special notice conventions that are used in this document

The following conventions are used for special notices:

NOTICE

Identifies content that warns of potential business or data loss.

Note

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

Typographical conventions

The following type style conventions are used in this document:

Table 2 Style conventions

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

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

Where to find product documentation

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

Where to get support

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

To access a product-specific page:

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

Knowledgebase

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

To search the Knowledgebase:

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

Live chat

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

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

Service requests

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

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

Note

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

To review an open service request:

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

Online communities

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

How to provide feedback

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

CHAPTER 1

Overview of NetWorker Snapshot Features

This chapter includes the following topics:

- [NetWorker Snapshot Management product description](#)..... 20
- [Components of the snapshot environment](#).....24
- [NetWorker snapshot licensing requirements](#).....26
- [Example NetWorker snapshot environments](#).....26

NetWorker Snapshot Management product description

The NetWorker Snapshot Management (NSM) feature works with replication and mirror technologies on supported storage arrays or storage appliances to create and manage snapshot and ProtectPoint copies of production data, with minimal disruption to the production host processes. The NetWorker server catalogs the snapshots, provides snapshot recovery, and clones the snapshots to Data Domain (ProtectPoint), or to conventional storage media, such as disk or tape. The Snapshot Management feature is available as part of the NetWorker extended client software package.

The NetWorker extended client installation provides all the functionality that the NetWorker PowerSnap Module previously handled. The *NetWorker Installation Guide* provides more details. [Migrating Legacy PowerSnap Configurations](#) provides examples of how to migrate legacy PowerSnap configurations to NetWorker Snapshot Management.

Before you plan, configure, and administer the snapshot environment, become familiar with the concepts in this chapter. You should have an advanced working knowledge of the storage array technology that you use with NetWorker Snapshot Management.

The *NetWorker E-LAB Navigator* provides details on the versions that NetWorker supports, including volume managers, NetWorker modules, and cluster environments.

Snapshot operations

NetWorker Snapshot Management supports the application host, which is a NetWorker client that writes production data to volumes on a supported storage array or storage appliance. These production volumes consist of one or more logical units (LUNs) of storage, which the array or appliance replicates to a mirror LUN or snapshot pool. The mirror LUN can be local or a LUN on a remote array or remote appliance.

NetWorker supports the following storage array and storage appliance configurations:

- ProtectPoint—VMAX3 or XtremIO to Data Domain vdisk snapshot operations.
- VMAX arrays—TimeFinder Clone, VDEV, BCV, VP Snap, SnapVX, and Symmetrix Remote Data Facility (SRDF) operations.
- VNX and VNXe Block arrays—SnapView Copy-on-write (COW/Snapshot), Mirror (clone), and VNX Snap operations.
- Unity arrays— VNXe-SNAP (VNXe/VNXe2 Snapshots) operation.
- RecoverPoint appliances that are configured on supported VMAX, VNX Block, XtremIO, and VPLEX storage arrays—Continuous Local Replication (CLR) and Continuous Remote Replication (CRR).

NetWorker uses the replication and splitting or the cloning capabilities of the array to create point-in-time (PIT) copies of specified production data onto a storage array volume. These PIT copies are called snapshots. In the case of ProtectPoint, NetWorker copies the snapshot to the DD vdisk.

To manage the snapshots, NetWorker mounts the snapshot volume on a mount host, which can be the application host, a NetWorker storage node, or a remote NetWorker client host. NetWorker uses the mount host for clone operations that save the snapshot to conventional storage media such as disk or tape, and for restore operations from the snapshot or conventional storage media.

NetWorker policies manage the lifecycles of the snapshot backups, and the backup copies that are cloned to conventional storage volumes from snapshots.

[Example NSM snapshot environments](#) provides illustrations of typical snapshot environments and describes the snapshot, clone, and recovery processes.

Types of snapshot backups

The type of NetWorker snapshot backup that you configure depends on where you intend to create and store the snapshot:

- Snapshot backup—NetWorker creates a snapshot of the specified files on the application host and stores the snapshot only on the storage array. The NetWorker server catalogs the snapshot as a backup in the media database. The NetWorker server can perform a restore from the snapshot.
- ProtectPoint backup—NetWorker creates a snapshot of the specified files on the application host and stores the snapshot only on a Data Domain device. The NetWorker server catalogs the snapshot as a backup in the media database. The NetWorker server can perform a restore from the snapshot.

Note

Snapshot refers to both storage array only snapshots, and snapshots copied to a Data Domain device with ProtectPoint.

Client file system layout considerations

The following are considerations for nested file systems:

- Taking a snapshot backup of both the parent file system and any file system that is mounted under the parent in the same backup is not supported.
- It is possible to take snapshots of the parent and any file system that is mounted under the parent in separate backup configurations.
- Support for the rollback workflow in nested file system configurations is limited.
- The rollback of both the parent and the underlying mounted file systems simultaneously is not supported.
- The rollback of any of the file systems that are mounted under the parent directory is supported.
- Rollback is not permitted of the parent file system while underlying file system is in mounted state.

Note

The NetWorker Module for SAP (NMSAP) has configuration parameters that allow you to exclude the snapshot of specific files. The *NetWorker Module for SAP Administration Guide* provides details.

Unsupported backup configurations

NetWorker does not support the following backup configurations:

- Containers (zones) on Solaris operating systems.
- Solaris ZFS file systems.
- Windows GPT, Dynamic volumes, and Windows volume management.

The *NetWorker E-LAB Navigator* provides details.

Types of snapshot recoveries

The type of recovery that you can perform for snapshot-based data depends on the location of the data and the following factors:

- Restore from a snapshot—NetWorker mounts the snapshots to a mount host, browses, and selects the directories to restore.
- Restore from a clone copy on conventional media—NetWorker performs a conventional restore from the backup storage media.
- Rollback restore—NetWorker restores the snapshots by using the storage array capabilities. The process unmounts the original source volumes on the application host and the rollback replaces the entire contents with the contents of a selected snapshot.

Note

NetWorker does not support rollback in the VMAX3 ProtectPoint workflow when restore devices are exported directly from the Data Domain.

A ProtectPoint RecoverPoint rollback of the clone snapshot that has been copied using NetWorker's cloned controlled replication is not supported.

NetWorker does not support rollback on XtremIO storage arrays.

NetWorker does not support rollback on VNXe arrays (Unity).

NetWorker clone support

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

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

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

Cloning is also supported to Data Domain CloudTier and CloudBoost.

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

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

The *NetWorker Administration Guide* provides details.

Note

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

Backup configuration methods

You can configure snapshot backups by using the NetWorker Management Console (NMC) interface. All the supported storage arrays support the following configuration methods:

- NetWorker Client Configuration Wizard—It is recommended that you use the wizard to create and modify the configurations for snapshots. The wizard accommodates the most common snapshot workflows by providing the correct sequence of steps and the verification of configuration dependencies.

- NMC Client Properties windows—Provides a user interface that you can use to manually create or modify configurations. For example, you can use the Client Properties Window to specify the uncommon directives or options that the wizard interface does not support, such as the variables listed in [Application Information Variables](#).

Note

To create necessary lockbox entries, RecoverPoint and XtremIO require that you type username and password information. These workflows do not support manual client configuration.

The *NetWorker Module for Databases and Applications Administration Guide* and the *NetWorker Module for SAP Administration Guide* provide details on the supported application backup and recovery interfaces.

Restore methods

You can use one of the following interfaces to restore snapshot-based data for file system backups:

- The NMC Recovery Wizard—It is recommended that you use the wizard to restore data from the snapshots and conventional storage media.
- The `nsrsnapadmin` command utility—An interactive Command Line Interface (CLI) tool that you can use for various snapshot-related operations, including restore from a snapshot. [Using nsrsnapadmin for snapshot operations](#) on page 162 provides details.
- The `nsrsnap_recover` command—A CLI method that you can use to restore data from a snapshot or conventional storage media.

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

The *NetWorker Module for Databases and Applications Administration Guide* and the *NetWorker Module for SAP Administration Guide* provide details on the supported application backup and recovery interfaces.

Monitoring and reporting snapshot operations

NetWorker enables you to monitor snapshot operations for each NetWorker client. You can monitor the progress of the snapshot creation, mounting, deletion, and cloning operations.

The NetWorker `nwsnap.raw` log file, on the application and mount host, provides detailed information about snapshot operations.

Internationalization support

The standard NetWorker client support for non-ASCII international character sets applies to snapshot management.

Components of the snapshot environment

You can deploy various required and optional hosts, devices, connectivity, and applications in a NetWorker datazone for snapshot management.

[Example NetWorker snapshot environments](#) on page 26 provides illustrations of typical snapshot environments and describes the snapshot, NetWorker clone, and recovery processes.

Application host

An application host in the snapshot environment is a computer with production data that resides on storage array volumes and requires snapshot services. The production data can consist of file systems and databases.

NSM supports snapshots of a VMware guest OS for raw device mapped volumes on a VMAX and iSCSI volumes on VNX. If you use VNX storage, when a RecoverPoint appliance controls the volumes, NSM supports RDM when you use VMAX storage or iSCSI. XtremIO snapshots and RecoverPoint are supported.

Each application host must be configured as a NetWorker host and must have the following software installed:

- NetWorker client 19.1.
- NetWorker extended client 19.1.
- NMDA/NMSAP 19.1 if you are protecting IBM DB2 data, Oracle data, or SAP with Oracle data.
- Must be configured as a NetWorker client.

FC and iSCSI environments

All hosts that are involved in the movement of production data within the NetWorker snapshot environment must use Fibre Channel (FC) connectivity, which is deployed as a storage area network (SAN). NetWorker Snapshot Management (NSM) supports iSCSI for VMAX, VNX, and XtremIO. NetWorker snapshots do not support Fibre Channel over Ethernet (FCoE) environments.

Note

NSM supports the VMware guest OS for VNX when the VNX volumes are using iSCSI. NSM does not support RDM volumes with VNX. The VMAX VMware, guest OS only supports raw device mapped (RDM) volumes.

Storage arrays

For snapshot operations, one or more supported storage arrays must provide logical units (LUNs) to store the application host's production data and the snapshots of this data.

NetWorker supports the following storage array and data management technologies:

- VMAX (Symmetrix) storage array
- VNX Block (CLARiiON) storage array
- VNXe Block
- RecoverPoint storage appliance

- XtremIO native snapshot

The *NetWorker E-LAB Navigator* provides details about the supported storage arrays.

NetWorker server

The NetWorker server manages the snapshot clients and the configuration settings that are required to create the snapshots and perform the cloning operations.

NetWorker storage node

The NetWorker storage node manages the devices for backups to conventional storage media, such as AFTD, DD Boost devices, Cloud Boost, Cloud Tier, and tape. Snapshot management requires a storage node for all the clone operations and to restore data from clones.

If you plan to create and restore snapshots and do not plan to clone snapshots, then the use of a storage node is optional.

Snapshot mount host

NetWorker requires a client host to mount the storage array's snapshot volumes for snapshot restore operations, snapshot validation, and for cloning to conventional storage media.

The mount host can be the local application host, a NetWorker storage node, or a remote NetWorker client host. The choice of mount host depends on the storage network configuration. A well-planned configuration includes consideration of the data processing speed and the bandwidth load on the different possible hosts.

The mount host must use the same operating system with the same third-party volume manager (if any) as the application host. Synchronize the system clocks of the mount host and the application host.

Note

Rollback operations do not use a mount host. Rollback is not supported in a nested file system environment.

Backup storage media

NetWorker can clone snapshots to conventional backup storage media, such as AFTD, Data Domain Boost devices, tape, Data Domain CloudTier, and CloudBoost.

NetWorker application modules

NSM supports application hosts integrated protection, with NetWorker Module for Databases and Applications (NMDA) and NetWorker Module for SAP (NMSAP) on VMAX, VNX, XtremIO storage arrays, and RecoverPoint.

The *NetWorker E-LAB Navigator* provides details on supported versions.

The *NetWorker Module for Databases and Applications Administration Guide* and the *NetWorker Module for SAP Administration Guide* provide details on application configurations.

Third-party volume managers

NetWorker supports the use of third-party volume managers for managing the storage array data, such as Veritas Volume Manager (VxVM) and Linux Logical Volume Manager (LVM). However, for VxVM managed volumes, NetWorker does not support the following configurations:

- If the production file system and the snapshot file system are simultaneously visible to the same host, then the backups can fail. Some operating systems or LVMs require that the production file system and the snapshot file system must be visible only on separate hosts, such as the application host and a different mount host.
- If multiple LUNs with the same disk signature or the same volume ID are visible to the same host, then the backups can fail. For example, if multiple mirrors or both the source and mirror LUNs are visible to the same host, then the backups can fail.
- VxVM on Microsoft Windows systems.

The *NetWorker E-LAB Navigator* provides support details.

NetWorker snapshot licensing requirements

The following types of licensing can enable snapshot management:

- NetWorker capacity licenses
- Traditional licenses

For both types of licensing, the NetWorker software reports on capacities that are consumed for the standard (nonsnapshot) backups and the snapshot backups.

If NetWorker detects valid PowerSnap licenses, then NetWorker honors the licenses.

The NetWorker source capacity enabler enables the use of snapshot management within the datazone up to the purchased total source capacity. The number of clients that you can protect within the datazone has no restriction.

For traditional licensing, you need a capacity-based NetWorker license and other required licenses, such as the client connection license, storage node license, and application module license for the NetWorker clients under protection.

The *NetWorker Licensing Guide* and your NetWorker sales representative can provide details about the types of licensing for NetWorker Snapshot Management.

Example NetWorker snapshot environments

Plan the NetWorker snapshot environment to manage data efficiently as illustrated by the following examples. [Snapshot operations](#) on page 20 describes basic snapshot concepts.

Example of a snapshot and clone to storage media

The following figures illustrate two variations of the flow of data during a snapshot and clone operation in a typical NetWorker snapshot environment.

Figure 1 Snapshot and clone operation with the storage node as the mount host

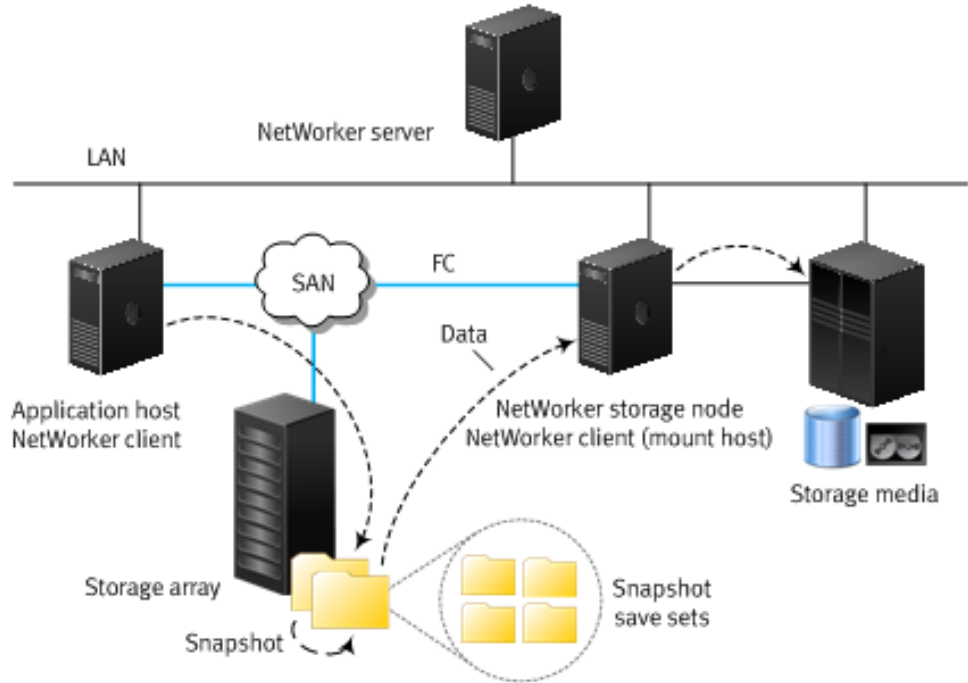



Figure 2 Snapshot and clone operation with the application host as the mount host



The process flow is as follows:

1. The application host processes its production data by writing to one or more source volumes on an attached storage array.

Note

The application host can run NMDA or NMSAP. As a common practice for these modules, the application host can have its own NetWorker storage node, which makes the application host also the mount host.

2. At a scheduled time, NetWorker creates a snapshot of the production data on a different volume on the storage array or on a different array:
 - a. NetWorker policies and client resource settings identify which data on the application host requires a snapshot.
 - b. NetWorker synchronizes the source LUNs with the target LUNs. The source LUNs contain the production data volumes. To ensure consistency, NetWorker together with NMDA or NMSAP quiesces the database, or if just file systems, quiesces and flushes data before taking the snapshot.
 - c. The storage array splits or fractures the target snapshot LUN from the production LUNs. This process creates a fully usable snapshot on the snapshot volume.
3. NetWorker optionally mounts the completed snapshot on the mount host to validate that the snapshot can be restored.

The choice of mount host depends on the storage network configuration with consideration of the data processing speed and the bandwidth load on the different possible hosts. For example, the mount host can be one of the following:

 - A NetWorker storage node or the application host as shown in [Example of a snapshot and clone to storage media](#) on page 27.
 - A remote mount host with the NetWorker client and extended client software installed.

Note

If the NetWorker client resource settings specify the Client Direct and Data Domain Boost options, on-client data deduplication processing occurs on the mount host during the clone operations to conventional media.


4. NetWorker manages the snapshot according to the options in the client resource settings and NetWorker policy resource:
 - If NetWorker clones the snapshot to conventional storage media, the snapshot data becomes available for additional NetWorker clone operations and conventional NetWorker restore operations.

The *NetWorker Administration Guide* provides details of storage media configurations.
 - If NetWorker does not clone the snapshot but retains the snapshot on the storage array or the Data Domain (ProtectPoint snapshot), the snapshot is available for restore, rollback, or clone. NetWorker retains the snapshot on the storage array only until it expires or until NetWorker must delete it to create snapshots, as specified by the NetWorker policy.

Example of a restore from a snapshot backup

The following figure illustrates the data flow for a selective restore of files from a snapshot save set. The NetWorker storage node restores data from the snapshot target volume to the production source volume.

Figure 3 Restore from a snapshot with the storage node as the mount host



The process flow is as follows:

1. NetWorker selects the snapshots that contain the data that you want to restore. NetWorker then mounts the snapshot on the mount host.
2. If you are restoring file system data, you locate the files, file systems, or volumes that you want to restore. If NetWorker is restoring application data, it requests specific application files that are required for the recovery.
3. You specify where to restore the data on the application host or alternatively on a different host.
4. When you start the restore, NetWorker contacts the mount host and the application host or an alternative restore host.
5. NetWorker copies the data from the snapshot volume to the specified volume:
 - For a file level restore, the data restore path is over the LAN.
 - For a rollback recovery, the storage array capabilities perform the recovery from the snapshot LUNs to the production LUNs. The snapshot containing the data is not mounted on the mount host for a rollback recovery.
 - NetWorker restores the snapshot using the saveset restore. In this instance NetWorker uses the entire saveset to restore data.

CHAPTER 2

Data Protection Policies

This chapter includes the following topics:

- [Default data protection policies in NMC's NetWorker Administration window](#)... 32
- [Strategies for storage array snapshot backups](#)..... 32
- [Overview of configuring a new data protection policy](#) 33
- [Creating a policy](#) 34
- [Create a workflow for a new policy in NetWorker Administration](#)..... 36
- [Protection groups for snapshot backups](#)..... 38
- [Actions supported in snapshot backups](#)..... 42
- [Supported actions in snapshot workflows](#)..... 43
- [Visual representation of snapshot workflows](#)..... 59

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 4 Platinum policy configuration



Gold policy

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

Silver policy

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

Bronze policy

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

Strategies for storage array snapshot backups

Multiple strategies for data protection policies are available to help you optimize how NetWorker Snapshot Management (NSM) performs snapshot backups.

When you protect storage array devices by using snapshot technology, the snapshot workflow supports the following actions:

- Probe
- Check connectivity
- Snapshot backup
- Clone

[Actions supported in snapshot backups](#) on page 42 provides more details.

Overview of configuring a new data protection policy

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

Procedure


1. Create a policy resource.
When you create a policy, you specify the name and notification settings for the policy.
2. Within the policy, create a workflow resource for each data type.
For example, create one workflow to protect file system data and one workflow to protect application data. When you create a workflow, you specify the name of the workflow, the time to start the workflow, notification settings for the workflow, and the protection group to which the workflow applies.
3. Create a protection group resource.
The type of group that you create depends on the types of clients and data that you want to protect. The actions that appear for a group depend on the group type.
4. Create one or more action resources for the workflow resource.
5. Configure client resources, to define the backup data that you want to protect, and then assign the client resources to a protection group.

Example 1 Example of a data protection policy with 2 workflows

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

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

Figure 5 Data protection policy example



Note

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

Creating a policy

Procedure

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

Note

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

4. In the **Comment** field, type a description for the policy.
5. From the **Send Notifications** list, select whether to send notifications for the policy:

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

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

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

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

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

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

```
mail -s subject recipient
```

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

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

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

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

where:

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

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

After you finish

Create the workflows and actions for the policy.

Create a workflow for a new policy in NetWorker Administration

Procedure

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

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

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

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

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

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

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

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

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


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


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

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

- On Windows, type the following command:

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

where:

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

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

After you finish

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

Protection groups for snapshot backups

A protection group for a snapshot backup identifies the client resources to back up.

Snapshot backups supports the following types of 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.

Creating a basic client group

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

Before you begin

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

Procedure

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

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

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

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

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

Note

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

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

Note

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

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

8. Click **OK**.

After you finish

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

Creating a dynamic client group

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

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

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

Procedure

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

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

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

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

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

Note

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

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

Creating a save set group

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

Before you begin

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

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, right-click **Groups**, and then select **New**.

The **Create Group** dialog box appears, starting with the **General** tab.

- In the **Name** field, type a name for the group.

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

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

Note

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

- From the **Group Type** list, select **Save Set ID List**.
- In the **Comment** field, type a description of the group.
- (Optional) To associate the group with a workflow, from the **Workflow (Policy)** list, select the workflow.
You can also assign the group to a workflow when you create or edit a workflow.
- In the **Clone specific save sets (save set ID/clone ID)** field, type the save set ID/clone ID (ssid/clonid) identifiers.
To specify multiple entries, type each value on a separate line.
- To specify the Restricted Data Zone (RDZ) for the group, select the **Restricted Data Zones** tab, and then select the RDZ from the list.
- Click **OK**.

Creating a query group

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

Procedure

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

Note

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

- From the **Group Type** list, select **Save Set Query**.
- In the **Comment** field, type a description of the group.
- (Optional) To associate the group with a workflow, from the **Workflow (Policy)** list, select the workflow.
You can also assign the group to a workflow when you create or edit a workflow.

7. Specify one or more of the save set criteria in the following table.

Note

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

Table 3 Save set criteria

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

Table 3 Save set criteria (continued)

Criteria	Description
Name	In the Filter save sets by name field, specify the name of the save set. <hr/> Note You cannot use wildcards to specify the save set name.

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

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

Actions supported in snapshot backups

The snapshot workflow supports the following actions:

Probe

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

Only a backup action can follow a probe action.

Check connectivity

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

Snapshot backup

A snapshot backup action performs a snapshot of the data on a supported storage device.

Clone

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

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

Note

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

In NetWorker 19.1, you can clone a snapshot backup to a cloud-enabled media pool.

Supported actions in snapshot workflows


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

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

All possible workflow actions for a snapshot backup

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

Figure 6 All possible workflow actions for a snapshot backup



You can configure an action to run concurrently with an existing action in a workflow. If you configure a clone action to run concurrently with a snapshot backup action, NetWorker clones the snapshot backup save sets for each client. For example, if a Protection Group has two clients (client1 and client2), and the group is assigned to a workflow that contains the snapshot backup action and the clone action, then after the client1 backup is completed, NetWorker clones the save sets for client1. When the client2 backup completes, NetWorker clones the save sets for client2.

Creating a check connectivity action

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

Before you begin

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

Procedure

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

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

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

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

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

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

from running when the policy or workflow that contains the action is started, clear this option.



Note

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

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

The following table provides details about the icons.

Table 4 Schedule icons

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

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

13. (Optional) Configure advanced options and schedule overrides.

Note

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

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

Note

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

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

Note

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

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

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

- **Disabled**—Do not enforce an action start time. The action will start at the time defined by the workflow.
 - **Absolute**—Start the action at the time specified by the values in the spin boxes.
 - **Relative**—Start the action after the period of time defined in the spin boxes has elapsed after the start of the workflow.
19. (Optional) Configure overrides for the task that is scheduled on a specific day.
- To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:
- Select the day in the calendar, which changes the action task for the specific day.

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

Note

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

20. Click **Next**.

The **Action Configuration Summary** page appears.

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.

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.

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

The following table provides details on the icons.

Table 5 Schedule icons

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

- Click **Next**.

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

11. Specify when to start the subsequent backup action:
 - To start the backup only if all the probes associated with client resources in the assigned group succeed, select the **Start backup only after all probes succeed** checkbox.
 - To start the backup if any of the probes are associated with a client resource in the assigned group succeed, clear the **Start backup only after all probes succeed** checkbox.

12. Click **Next**.

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

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

Note

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

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

Note

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

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

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

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

Note

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

16. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.
17. From the **Failure Impact** list, specify what to do when a job fails:
 - To continue the workflow when there are job failures, select **Continue**.

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

Note

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

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

Note

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

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

Note

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

23. Click **Next**.

The **Action Configuration Summary** page appears.

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

Creating a snapshot backup action

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

Before you begin

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

Procedure

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

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



2. From the **Action Type** list, select **Backup**.
3. From the secondary action list, select **Snapshot**.
4. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
5. 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.
6. Specify a weekly or monthly schedule for the action:
 - To specify a schedule for each day of the week, select **Weekly by day**.
 - To specify a schedule for each day of the month, select **Monthly by day**.
7. Specify the type of backup to perform on each day:
 - To specify a level on a specific day, click the backup type icon on the day.
 - To specify the same type of backup on each day, select the backup type from the list, and then click **Make All**.

Note

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

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

Table 6 Backup type icons

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

8. Click **Next**.

The **Snapshot Options** page appears.

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

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

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

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

13. Click **Next**.

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

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

Note

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

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

Note

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

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

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

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

Note

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

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

Note

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

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

Note

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

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

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

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

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

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

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

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

```
mail -s subject recipient
```

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

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

where:

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

21. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
22. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
23. Click **Next**.
The **Action Configuration Summary** page appears.
24. Review the settings that you specified for the action, and then click **Configure**.

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

The following table provides details on the icons.

Table 7 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.

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.
-
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 snapshot workflows

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

Figure 7 Sample snapshot workflow



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

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

CHAPTER 3

Software Configuration

This chapter includes the following topics:

- [Backup group resource migration](#).....62
- [Roadmap for snapshot configurations](#).....64
- [Snapshot configuration prerequisites](#).....65
- [Configuring the user privileges](#).....67
- [Configuring snapshot backups with the client wizard](#).....68
- [Configuring snapshot backups manually](#).....74
- [Configuring the Application Information variables](#).....76
- [Configuring preprocessing and postprocessing scripts](#).....76

Backup group resource migration

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

Resource migration for group resources when Snapshot is enabled

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

Table 8 Migration of Group attributes

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

Table 8 Migration of Group attributes (continued)

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

Note

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

Considerations about migration

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

Roadmap for snapshot configurations

The following high-level road map outlines the sequence of NetWorker snapshot configuration tasks that you must perform.

1. Verify the configuration prerequisites.
[Snapshot configuration prerequisites](#) on page 65 provides details.
2. Configure the user privileges on the application host and the storage node.
[Configuring the user privileges](#) on page 67 provides details.
3. Configure the NetWorker client for snapshots by using the Client Configuration Wizard or the manual method. The following topics provide details:
 - [Configuring snapshot backups with the client wizard](#)
 - [Configuring snapshot backups manually](#)
4. Configure any necessary Application Information variables.
[Configuring the Application Information variables](#) provides details.
5. Configure any necessary preprocessing or postprocessing scripts.
[Configuring preprocessing and postprocessing scripts](#) provides details.
6. Based on the array or appliance that you use for snapshot backups, follow the appropriate configuration instructions:
 - [Configuring snapshots on VMAX Storage Arrays](#)
 - [Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard](#) on page 85
 - [Configuring snapshots on VNX Block Storage Arrays](#) on page 115
 - [Configuring snapshots on RecoverPoint](#) on page 121
 - [Configuring snapshots in a Cluster Environment](#)
7. Test the configuration.

Snapshot configuration prerequisites

Verify the basic compatibility of all systems used for NetWorker snapshot operations. [Components of the snapshot environment](#) provides details.

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

Storage array specific prerequisites

Ensure that you install the application host and the mount host with the prerequisite software for the storage array that you use for NetWorker snapshot operations:

- VMAX storage array must have the following software installed:
 - Solutions Enabler on the application host and (if cloning snapshots) the NetWorker server.
- VNX Block storage array must have the following software installed:
 - Unisphere® host agent, also known as Navisphere®, on the application host and the mount host.—The agent synchronizes the host device with the VNX devices. NetWorker also uses the agent to determine whether the LUNs are visible on the application host or the mount host.
 - SnapCli on the mount host and optionally on the application host—This CLI is responsible for making the VNX snapshot LUN visible to the mount host.
 - AdmSnap on the mount host and optionally on the application host—This CLI is responsible for making the SnapView snapshot LUN visible to the mount host.
 - AdmHost on the mount host for Microsoft Windows systems only—This CLI is responsible for activating and mapping the SnapView clone to a specific drive letter on a Microsoft Windows mount host.
 - Naviseccli or UEMCLI on the application host and the mount host—This CLI is responsible for LUN discover and snapshot sync/split operations. The *NetWorker E-LAB Navigator* provides more details.
 - UEMCLI on the mount host and application host—This CLI is used for later VNX releases and replaces other CLI. Check the array specific requirements.
 - For snapshot (COW) backups, create the snapshot in advance on the VNX and mount the snapshot to the proxy host. This is a mandatory setting, if you do not perform this requisite, the backup fails.
 - For clone (MIRROR) backups, create the clone on the VNX. The clone group should be in a synchronized state. The clone must be mounted to the proxy host. This is a mandatory setting, if you do not perform this requisite, the backup fails.
 - For VNX -snap backups, create the snapshot mount point on the source LUN, mount the snapshot to the proxy host. This is a mandatory setting, if you do not perform this requisite, the backup fails.
- RecoverPoint appliance must have the following software installed:
 - For RecoverPoint 4.0 and later, Solutions Enabler is not required.
 - RecoverPoint Continuous Data Protection (CDP), Continuous Local Replication, and RecoverPoint Continuous Remote Replication (CRR) configured on the RecoverPoint appliance.

The support matrix for the storage array or appliance that you use, available from the Support website, provides details on system and software requirements.

Application host prerequisites

Verify the following prerequisites before you configure the application host for snapshot operations:

- The application host has installed the supported NetWorker client and NetWorker extended client software.
- The application host has the support NetWorker Module for Databases and Applications (NMDA) and NetWorker Module for SAP (NMSAP) software, if it is protecting DB2, Oracle, or SAP with Oracle.
- The NetWorker server recognizes the application host as a client.
- The application host has completed at least one successful NetWorker backup.
- Synchronize the application host system clock with the mount host system clock.
- Set up the application host connection to the storage array:
 - VMAX, VNX, and XtremIO storage arrays, by themselves or with a RecoverPoint appliance, require a SAN connection. For a VMware guest OS, the source and target LUNs are visible as raw device mapped (RDM). VMAX supports both iSCSI and traditional device mapping. VNX, VNX2, VNX2e, and XtremIO only support RDM via iSCSI.
 - A RecoverPoint, XtremIO, and VNX appliance also require a LAN connection for communication with the application host.
- The volume or device pathnames of the production LUNs on the storage array are visible to the application host.
- If you use a separate mount host, then the volume or device pathnames of the mirror (target) LUNs on the array are visible to the mount host.

Mount host prerequisites

A NetWorker client must mount the storage array volumes for the snapshot restore or clone operations to conventional media. You can configure any of the following hosts as a mount host:

- Local application host
- NetWorker storage node
- Remote NetWorker client host

The choice of mount host depends on the storage network configuration. A well-planned configuration considers the data processing speed and the bandwidth load on the different possible hosts.

Before you configure the mount host for snapshot operations, ensure that you perform the following prerequisites:

- Confirm that the mount host runs the same versions of the operating system and the volume manager (if any) as the application host.
- Install the NetWorker client and NetWorker extended client software.
- Ensure that a NetWorker server recognizes the mount host as a client.
- Synchronize the system clock of the mount host with the system clock of the application host.

- Set up the mount host connection to the storage array:
 - VMAX, VNX, and XtremIO storage arrays, by themselves or with a RecoverPoint appliance, require a SAN connection.
 - A RecoverPoint, VNX, and XtremIO appliance also require a LAN connection for communication with the mount host.
- The volume or device pathnames of the snapshot target LUNs on the storage array are visible to the mount host.

Note

Windows supports 26 drive letters (A to Z). NetWorker Snapshot Management (NSM) uses a drive letter when it mounts target volumes for restoring or moving data to media. If the mount host does not have enough available drive letters to mount all the target volumes, then the operations use Windows mount points to mount the target volumes.

Storage node prerequisites

If you plan to clone the snapshot save sets to conventional storage media, then use a NetWorker storage node as the mount host.

Note

If you prefer to perform clones by using the local application host as the mount host, consider upgrading the NetWorker client on the application host to a NetWorker storage node.

Ensure that you complete the following prerequisites:

- NetWorker storage node 19.1 or later software is installed.
- Backup storage devices are configured on the storage node for the clone operations.

Configuring the user privileges

Specify the NetWorker User Groups privileges on the application host and the mount host for snapshot operations.

Procedure

1. Run NMC, and in the **Enterprise** view, select the NetWorker server that manages the snapshots, and then select **Enterprise > Launch Application**.
2. On the **Server** tab, in the resources tree, click **User Groups**.
3. In the **User Groups** table, right-click the group that you want to modify and select **Properties**.
4. In the **Users** attribute, to specify the user as root, administrator, or system on the application host and the mount host, type the following information:

- Microsoft Windows systems:

```
user=administrator,host=application_hostname
user=administrator,host=mount_hostname
user=system,host=application_hostname
user=system,host=mount_hostname
```

- UNIX systems:


```
user=root,host=application_hostname
user=root,host=mount_hostname
```
5. In the **Privileges** attribute, select **Operate NetWorker**.

Note

The **Operate NetWorker** privilege can require the selection of additional privileges as indicated in a pop-up message.

6. Click **OK**.

Configuring snapshot backups with the client wizard

The NMC Client Configuration Wizard helps you to configure a client resource for snapshot, clone, and backup operations.

Before you begin

Ensure that the system meets the necessary prerequisites. [Snapshot configuration prerequisites](#) on page 65 provides details.

Note

The following steps are only required for scheduled backups not using the ProtectPoint workflow. NetWorker also supports client-initiated ProtectPoint backups, and the following steps are not required if you perform only a client-initiated or manual backup.

Procedure

1. Run NMC, and in the **Enterprise** view, select the NetWorker server that manages the snapshots, and select **Enterprise > Launch Application**.
2. In the NetWorker server's **Protection** view, in the navigation tree, right-click **Clients**, and then select **New Client Wizard**.
3. The wizard displays the **Specify Client Information** page. Complete the following fields:
 - a. In the **Client Name** field, type the hostname of the application host on which NetWorker captures in the snapshots.
 - b. (Optional) In the **Comment** box, add notes for the Client.
 - c. (Optional) In the **Tag** box, type one or more tags to identify this Client resource for the creation of dynamic client groups for data protection policies. The tags are user defined.
Type each tag on a separate line.
 - d. To add the current client to an existing group of clients that use the same workflow for snapshot, clone, or backup actions, from the **Group** list select the group.

Alternatively, you can add the client to one or more groups later, after you complete the wizard.
 - e. In the **Type** area, select **Traditional**.

f. Click **Next**.

4. The **Specify Backup Configuration Type** page displays. Complete the following fields:

a. In the **Available Applications** table, you can select **Filesystem**, **SmartSnap** or another supported NetWorker application type that is installed on the client.

The **SmartSnap** option allows you to specify array LUNs World Wide Names (WWNs).

Note

If you select **SmartSnap** the **Enable NetWorker Snapshot Management on the selected application** is automatically selected. Also, the SmartSnap LUN does not need to be mounted or visible to the host.

b. Select **Enable NetWorker Snapshot Management on the selected application**.

c. Click **Next**.

5. The **Select the Snapshot Management Options** page displays. Complete the following fields:

a. Select the type of storage array or storage appliance that the client uses for primary storage and where the snapshots are created. Only storage arrays that are compatible with NetWorker and the client operating system appear. The options are:

- **VMAX/Symmetrix**
 - **ProtectPoint for VMAX3**
 - **ProtectPoint for RecoverPoint**
 - **VNX/CLARiiON**
 - **RecoverPoint**
 - **XtremIO**
-

Note

SmartSnap is only supported for VMAX/Symmetrix and ProtectPoint for VMAX3.

b. To use the application host as the snapshot mount host, select **Use the current client as the mount host**.

If the selected array supports the use of the current client as the mount host for the backup type, NetWorker mounts the array's mirror volume on the current host for snapshot restore operations and for clone actions. Otherwise, the wizard uses the storage node that you select on this page as the mount host. The mount host must use the same operating system as the current client host.

Note

Alternatively, you can manually create a NetWorker Client resource on a different host and specify this host as the value of `NSR_DATA_MOVER=hostname`. NetWorker uses this host as the mount host, eliminating the need to have a storage node or the application host as the mount host. [Common Application Information variables](#) provides details.

- c. From the drop-down list, select the hostname of the NetWorker storage node.
 - d. Click **Next**.
-

Note

The storage array that you selected on this page determines the next wizard page.

6. If you selected the **VMAX/Symmetrix** storage array option, complete the **Select the VMAX Mirror Policy** page that appears:
 - a. Select the **VMAX Mirror Policy**:
 - If both the source device and the target mirror devices reside on the same VMAX array, select **Local Operation**.
 - If the source device and the target mirror devices reside on different VMAX arrays that are connected by a Symmetrix Remote Data Facility (SRDF), select **Remote (SRDF) Operation**.
 - b. Select a snapshot **Mirror Type**:
 - To create snapshots using TimeFinder SnapVX functionality, select **SNAPVX**. This option is available only on VMAX version 3 and later storage arrays.
 - To create snapshots by using the TimeFinder clone functionality, select **CLONE**.
 - To create snapshots by using the TimeFinder/Snap (COW) functionality, select **VDEV**.
-

Note

VMAX version 3 and later does not support VDEV functionality.

- To create snapshots by using the TimeFinder VP Snap functionality, select **VPSNAP**.
- To create snapshots by using the TimeFinder split-mirror functionality, select **BCV**.
- If you selected **Remote (SRDF) Operation**, select **R2** to create a non-snapshot backup to media directly from the remote R2 device.

[Snapshot operations with TimeFinder software](#) on page 108 provides details on mirror operations.

- c. Click **Next**.

The **Select the NetWorker Client Properties** page appears. Go to step 12.

7. If you selected the **ProtectPoint for VMAX3** storage array option, complete the **Select ProtectPoint Destination** page that appears:
 - a. In the **ProtectPoint Remote Operation** area:
 - If both the source device and the target mirror devices reside on the same ProtectPoint array, select **Local Operation**.
 - If the source device and the target mirror devices reside on different ProtectPoint arrays that are connected by a Symmetrix Remote Data Facility (SRDF), select **Remote (SRDF) Operation**.
 - b. Click **Next**. The **Select the NetWorker Client Properties** page appears. Go to step 12.
8. If you selected the **ProtectPoint for RecoverPoint** storage array option, complete the **Specify the RecoverPoint Storage Array Options** page that appears:
 - a. From the drop-down list, select or type the **RecoverPoint Appliance Hostname / IP** that the client uses for snapshot communications.
 - b. If required, provide **Username** and **Password** credentials for the array that the client uses for snapshot operations.
 - c. Click **Next**. The **Select the NetWorker Client Properties** page appears. Go to step 12.
9. If you selected the **VNX/CLARiiON** storage array option, then complete the **Specify the VNX Mirror Policy and Storage Array Options** page that appears:
 - a. Select the snapshot **Mirror Type**:
 - To create the snapshots by using the SnapView copy-on-write functionality, select **Copy on Write (COW/Snapshot)**.
 - To create the snapshots by using the SnapView clone functionality, select **Mirror (SnapView Clone)**.
 - To create the snapshots by using the Redirect on Write (ROW) functionality, select **VNX-SNAP (VNX Snapshot)** or **VNXe-SNAP (VNXe/VNXe2 Snapshot)**.
[Snapshot operations with SnapView software](#) on page 116 provides details on mirror operations.
 - b. Specify the **VNX Storage Processor Options**:
 - If required, provide **Username** and **Password** credentials for the array that the client uses for snapshot operations.
 - Specify the **VNX storage array hostname or IP address** that the client uses for snapshot communications.
 - c. Click **Next**. The **Select the NetWorker Client Properties** page appears. Go to step 12.
10. If you selected the **RecoverPoint** storage array option, complete the **Specify the RecoverPoint replication type and Storage Array Options** page that appears:
 - a. Select the **Replication Type**:
 - **CDP** (Continuous Data Protection)
 - **CRR** (Continuous Remote Replication)

[Snapshot operations with RecoverPoint software](#) provides details on mirror operations.

b. Specify the **RecoverPoint Appliance Credentials**:

- Specify the **RecoverPoint Appliance Hostname/IP** that the client uses for snapshot operations.
- If the **Username** and **Password** credentials do not exist for the RecoverPoint appliance that the client uses for the snapshot operations, provide them.

Note

NetWorker requires a username and password after an upgrade from the RecoverPoint PowerSnap module to NetWorker Snapshot Management. The NetWorker server stores these credentials in a lockbox.

c. Click **Next**. The **Select the NetWorker Client Properties** page appears. Go to step 12.

11. If you selected the **XtremIO** storage array option, complete the **Specify the XtremIO storage process's credentials** page that appears:

- a. Specify the **XtremIO Hostname/IP** that the client uses for snapshot operations.
- b. If the **Username** and **Password** credentials do not exist for the XtremIO appliance that the client uses for the snapshot operations, provide them.
- c. Click **Next**.

12. If you plan to clone the snapshots to conventional storage media, complete the **Select the NetWorker Client Properties** page that appears:

- a. (Optional) Change the **Client Direct** setting based on the workflow preferences and the data processing bandwidth on the hosts involved.

Client direct enables the NetWorker client on the mount host to bypass the storage node and directly clone the snapshot data to supported AFTD or DD Boost devices or Cloud devices. If this process is not available on the client, then the storage node performs the clone operation.

- b. (Optional) In the **Advanced Options** area, enter **Debug** or **Extra Options** for the client that is separated by commas. For example,


```
NSR_PS_DO_PIT_VALIDATION=FASLE,
NSRATTER_DD_VDISK_POOLNAME=Pool ABC, NSR_VERBOSE=Yes.
```

- c. Click **Next**.

13. If you selected **Filesystem** in Step 4a, the **Select Snapable Filesystem Objects** page displays.

Note

If you are using NMDA/NMSAP, the application-specific wizard pages appear. Complete the pages accordingly.

For applications supported, go through application specific pages in the wizard and select the application objects to be restored and recovered:

- a. In the browse tree, select the file systems or files that you want to include in the snapshot.

The tree lists all the file systems that are mounted on the application host that are compatible with the previous selections in this wizard. If no compatible file systems are available, then an error appears.

On UNIX systems, the browse tree displays only the file systems that are added to the list of mountable file systems on the local application host.

Note

For VMAX arrays on systems with third-party volume managers, avoid selecting file systems or application objects from different volume groups for a single backup. If more than one snapshot is required to complete the backup and mirrors are unavailable, the backup from multiple volume groups can fail. However, you can greatly reduce this risk by using intelligent pairing and allocating sufficient devices to the storage group (NSRSnapSG).

- b. If you need to update the view of mounted file systems, click **Refresh**. This process can take some time.
 - c. Click **Next**.
14. If you selected **SmartSnap** in Step 4a, the **Add Snapable Lun WWN's** page appears. Perform the following steps:
- a. Type the WWN you want to include, and then click **Add**. You can add multiple WWN's.

The system validates the entry. If the validation is successful, the WWN appears in the list. If the WWN entered is invalid, then an error message is displayed indicating an invalid WWN.
 - b. If you want to remove a WWN select it, and then click **Delete**.
 - c. Click **Next**.
15. The **Client Configuration Summary** page, perform the following steps:
- a. Review the attributes and values that are listed in the summary.

To modify a setting, click **Back** or click the link in the step panel, and then make changes.
 - b. (Optional) Click **Snapshot Validation**. This choice causes NetWorker to verify the likely, but not guaranteed, success of a backup that uses this configuration, provided the backup runs unimpeded by other backups on this client and mount host.

Note

If you selected **SmartSnap** in Step 4a, the **Snapshot Validation** do not display.

The validation can take some time. If the validation encounters any errors an NMC pop-up message appears, which displays each problem but does not prevent the wizard from creating the Client resource.

Note

To validate the snapshot configuration of a Client resource, on the **Protection** screen in the Clients area, right-click the Client resource and select **Check Snapshot Configuration**.

- c. To accept and create the configuration, click **Create**.
16. In the **Check Results** page, perform the following steps:
 - a. Ensure that the client configuration is successfully completed.
 - b. Click **Finish**.

Configuring snapshot backups manually

It is recommended that you use the NetWorker Management Console (NMC) Client Configuration wizard to create and modify NetWorker client backup configurations.

However, in some situations you can use manual methods to modify a configuration. For example, you can modify a client resource if you must specify uncommon directives or options that the wizard does not support, such as the variables described in [Application Information Variables](#) on page 153.

Configuring the Client resource manually for the application host

You can manually create or modify a VMAX or VNX Block storage array configuration for an application host. You can manually modify a RecoverPoint appliance or XtremIO configuration, but you cannot create the configuration with the new NetWorker Management Console (NMC) Client Properties window, because the wizard is used to type username and password information into the NetWorker server lockbox.

You can manually specify uncommon directives or options that the wizard interface does not support, such as the variables listed in [Application Information Variables](#).

Procedure

1. Ensure that the prerequisites are met.

[Application host prerequisites](#) provides details.
2. Run NMC, and in the NMC **Enterprise** view, select the NetWorker server name, and then click **Launch Application**.
3. Click **Protection** and in the browse tree, select **Clients**, and then specify the application client:
 - To create a Client resource, click the **Clients** icon. From the **File** menu, select **New**.
 - To modify a Client resource from the list in the right panel, select the client name. From the **File** menu, select **Properties**.

Note

The *NetWorker Module for Databases and Applications Administration Guide* and the *NetWorker Module for SAP Administration Guide* provide details on configuring the additional attributes that are required for applications.

4. Click the **General** tab, and then perform the following steps:
 - a. In the **Name** field, verify or type the hostname of the application client.
 - b. For file systems, in the **Save sets** field, type or browse and select all the file systems, directories, or individual files that you want to include in the snapshot.

When you type the file system objects, specify each item on a separate line with a fully qualified pathname. Pathnames are case sensitive.

Note

Due to operating system limitations, a backup might fail when you specify file system pathnames that are longer than 996 characters or 275 directories deep.

The *NetWorker Administration Guide* provides details on the **General** tab settings.

5. Click the **Apps & Modules** tab. In the **Application information** field, specify the snapshot attributes with the values that you want the configuration to use.

[Application Information Variables](#) provides details.

Note

NetWorker does not validate these attributes. Type the correct attribute name in uppercase characters with the proper value specified, which depending on the attribute can also be uppercase.

6. On the **View** menu, to put NMC in Diagnostic mode, select **Diagnostic mode**.
7. If you plan to clone snapshots to conventional storage media, click the **Globals (2 of 2)** tab, and then in the **Remote Access** field specify the mount host in the following format:

- On Microsoft Windows systems:

```
system@mount_host
```

- On UNIX systems:

```
root@mount_host
```

The mount host is the host that mounts the storage array volume that contains the snapshots. Typically, the mount host is the application host or the storage node.

8. When you have completed the client configuration, click **OK**.
9. To verify the likely, but not guaranteed, success of the backup configuration, provided the backup runs unimpeded by other backups on this client or the mount host, in NMC right-click the Client resource, and then select **Check Snapshot Configuration**. The validation can take some time. If the validation encounters any errors an NMC pop-up message appears, which displays each problem.

The *NetWorker Administration Guide* provides details on NetWorker client configurations that are not specific to snapshot management.

Note

The **Check Snapshot Configuration** option does not appear when you use NetWorker Snapshot Management (NSM) with NetWorker Module for Databases and Applications (NMDA) and NetWorker Module for SAP (NMSAP).

Configuring the Client resource manually for a mount host

If you manually configured the Client resource for the application host, then manually configure a NetWorker Client resource for the snapshot mount host.

Procedure

1. If one does not exist, create a NetWorker Client resource for the mount host.
2. Ensure that the mount host has completed at least one successful NetWorker backup.
3. Ensure that the prerequisites are met.

[Mount host prerequisites](#) on page 66 provides details.

4. If you use the mount host only for snapshot restores and clones, clear the selection of its scheduled backups and all its groups as follows:
 - a. Run NMC, and in the **Enterprise** view, select the NetWorker server name, and then launch the NetWorker application.
 - b. From the View menu, select **Diagnostic Mode**.
 - c. In the **Protection** view, in the browse tree, select **Clients**, In the right panel, right-click the mount host, and select **Modify Client Properties** to view its **Properties**.
 - d. Click the **General** tab, and then clear **Scheduled Backup** checkbox.
 - e. Click **OK**.

Configuring the Application Information variables

Special Application Information variables provide specific control of snapshot processes. The Client Configuration Wizard cannot configure these variables. Manually configure these variables in the NetWorker Client resource for the application host using the NetWorker Client Properties windows. [Application Information Variables](#) on page 153 provides details.

Configuring preprocessing and postprocessing scripts

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

Note

Technical Support does not support the contents of the user-defined scripts. Scripts for a particular operation must return the correct exit code to NetWorker Snapshot Management (NSM).

You can use preprocessing scripts and postprocessing scripts for operations such as application quiescing, shutdown, or startup. The following steps provide guidelines for configuring the scripts.

Procedure

1. The scripts can produce output such as log files, but the scripts must return an exit status of 0, which means that the script did not fail and the backup can run. Any other exit code for a preprocessing script causes the backup to fail.
2. Provide the script files with the following security:
 - On Microsoft Windows systems, provide the script files with security that grants full control only to the local SYSTEM, local Administrators, or Backup Operators groups. Otherwise, the scripts will not run.
To set this security in Windows Explorer, right-click the script file, select **Properties**, click the **Security** tab, and click **Advanced**.
 - On UNIX systems, the root user must own the script files. The scripts can set only owner access permissions, and the scripts must at least have run access. Otherwise, the scripts will not run. The parent directory of the scripts must have at least owner run permissions, and must not have write permissions for the group and world.
3. Place the scripts in a directory where a user must have administrator/root privileges to add, modify, or run the resident scripts. Otherwise, any backups that use the scripts fail.

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

4. Include the pathnames of user-defined scripts in the **Application Information** attribute of the property window of the application Client resource by typing the following variables:

```
NSR_PRE_SNAPSHOT_SCRIPT=pre-mirror_split-script_path
NSR_POST_SNAPSHOT_SCRIPT=post-mirror_split-script_path
```

5. After a backup is completed, verify the log files that are generated in the `/nsr/logs` (UNIX) directory on the application client host. The log file name is in the form of `script_name_LOGFILE.txt`. The script output appears in the log file.

CHAPTER 4

Configuring ProtectPoint on VMAX

Note

This chapter provides basic information for VMAX and RecoverPoint with XtremIO and Data Domain configurations for ProtectPoint operations. The Data Domain, VMAX and RecoverPoint documentation provides details about the corresponding system configurations.

This chapter includes the following topics:

- [Overview](#)80
- [ProtectPoint on VMAX3 prerequisites](#) 80
- [Enabling vDisk on a Data Domain system](#) 81
- [Provisioning protection devices on Data Domain systems](#)81
- [Completing the VMAX system configuration](#)82
- [Considerations for ProtectPoint device and NetWorker ProtectPoint enabled pools](#) 85
- [Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard](#) 85
- [VMAX3 SRDF/S support](#)87
- [Configuring Data Domain NsrSnapSG device groups for intelligent pairing](#) 88

Overview

The ProtectPoint™ solution integrates primary storage on a VMAX3 array, and protection storage for backups on a Data Domain® system. ProtectPoint provides efficient block movement of the modified tracks containing data on the application source LUNs to encapsulated Data Domain LUNs for deduplicated snapshot backups.

You can create ProtectPoint backups by using one of the following methods:

- To use local storage array as the first line of protection, create a policy workflow with a snapshot action followed by a clone action that sends the data to a NetWorker pool containing a DD ProtectPoint device. Set `NSR_SNAP_TYPE=symm-dmx` and `SYMM_SNAP_TECH=snapvx` in the client resource settings.
- To use ProtectPoint backup as the first line of protection, create policies with a snapshot action that send data to a destination pool that contains ProtectPoint device. Optionally, create a clone action that sends data to other ProtectPoint devices or other media devices. Set `NSR_SNAP_TYPE=protectpoint` in the NetWorker client resource settings.

You can provision Data Domain backup and restore devices, create a NetWorker ProtectPoint device, and label the device for a NetWorker pool. When you create a snapshot backup or clone action in a workflow, you can select the destination pool that you used when you labeled the device.

ProtectPoint on VMAX3 prerequisites

Ensure that the following prerequisites are met before you use ProtectPoint devices:

- Use a VMAX3 array.
- Ensure that Solutions Enabler 8.2-2153 or later is installed on the application host, mount host, storage node, mount host, and in some configurations on the NetWorker server.

Configuring ProtectPoint

Procedure

1. On Data Domain, enable vDisk.
2. Provision protection vDisk devices on Data Domain.
3. Configure the protection devices on the VMAX3.
4. (Optional) Create restore vDisk devices on Data Domain, if you are restoring directly from Data Domain.
5. Create a NetWorker device, and then label it under a NetWorker pool using the wizard.
6. To create snapshot policies, follow the steps in the [Data Protection Policies](#) section, and then set the destination pool to a pool that contains a ProtectPoint device.
7. Follow the steps in the [Configuring snapshot backups with the client wizard](#) on page 68 or [Configuring snapshot backups manually](#) on page 74 sections, and then set the options accordingly.

Enabling vDisk on a Data Domain system

Enable vDisk on a Data Domain system through the `vdisk enable` command. Use the Data Domain command line interface to complete the required administration tasks. The *Data Domain Operating System Command Reference Guide* provides details about the commands.

Procedure

1. Log in to the Data Domain system, as an administrator.
2. To verify that the vDisk license is enabled, type the following command:

```
# license show
```

```
Feature licenses:
## License Key Feature
-----
1 ABCD-EFGH-IJKL-MNOP DDBOOST
-----
```

If the Data Domain vDisk license is disabled, to add the vDisk license by using the provided license key, type the following command:

```
# license add license_key
```

```
License "ABCE-BCDA-CDAB-DABC" added.
```

3. To enable the vDisk service subsystem, type the following command:

```
# vdisk enable
```

```
DD VDISK enabled
```

4. To verify that the vDisk service is enabled, type the following command:

```
# vdisk status
```

```
VDISK admin state: enabled, process is running, licensed
```

Provisioning protection devices on Data Domain systems

The Data Domain administrator must create the required vDisk device pool and device group and configure the backup and restore devices in the vDisk device pool and device group. You (DD Administrator) can use the vDisk mappings in the following table to plan the configuration.

Table 9 vDisk object hierarchy mapping

vDisk storage object	Mapping level
Device pool	NetWorker Datazone
Device group	Application
Device	Source LUN

Note

A NetWorker ProtectPoint device represents a specific vdisk pool.

The *Data Domain Operating System Administration Guide* provides the latest information about the vDisk configuration and any limitations.

Procedure

1. To create the vDisk device pool and device group, type the following vDisk commands:

```
vdisk pool create pool_name user username
```

```
vdisk device-group create device_group_name pool pool_name
```

The *Data Domain Operating System Command Reference Guide* provides details about the vDisk command and options.

2. For each VMAX3 source LUN that contains data, ensure that you have one backup device and one restore device. You might need additional restore devices for each source LUN if you also plan to restore data to different hosts.
3. To create the backup devices and restore devices with the same geometry as the VMAX source LUNs, and provision the devices accordingly, type the following vDisk command:

```
vdisk device create [count count] heads head_count cylinders  
cylinder_count sectors-per-track sector_count pool pool_name  
device-group device_group_name
```

Note

The heads, cylinders, and sectors-per-track information that defines the device geometry must match the VMAX source LUNs.

4. Add the vDisk devices to the access group on the Data Domain system:
 - a. Create a vDisk access group:

```
scsitarget group create group_name service vdisk
```
 - b. Add all the vDisk devices to the vdisk access group:

```
vdisk group add group_name pool vdisk_pool_name device-group  
device_group_name
```
5. Verify that the VMAX3 DX ports and the Data Domain endpoint ports are zoned together.
6. View the list of VMAX3 initiators on the Data Domain system:

```
scsitarget initiator show list
```
7. Add the VMAX3 initiators to the access group on the Data Domain system.
8. Add Symm initiators to the group:

```
vdisk group add group_name initiator initiator_name
```

Completing the VMAX system configuration

The VMAX storage administrator must complete the required steps to configure and provision the VMAX storage resources for the ProtectPoint operations.

Procedure

1. View the back-end ports (DX ports) on the VMAX array, and display the WWNs of the Data Domain devices zoned to the VMAX DX ports being viewed.
2. Display the LUNs that are visible for a specific Data Domain WWN.
3. List the disk groups that are available on the VMAX array.
4. Encapsulate all the Data Domain backup and restore devices to create the required eLUNs through the Federated Tiered Storage (FTS) software, and provision the restore LUNs to the restore host.

The *Solutions Enabler Symmetrix CLI Command Reference Guide* provides details about the commands and options for FTS operations.

5. Mask all the restore devices to the recovery host.

Note

The application host can be the recovery host.

6. You can either create a VMAX storage group named NSRSNAPSG for the restore devices or you can create your own VMAX storage group. You must then specify NSM_SNAP_SG attribute in the recovery operation, with value of VMAX storage group name. This group name is case-insensitive:

```
# symmsg -sid SymmID create NsrSnapSG
```

```
# symmsg -sid SymmID -sg NsrSnapSG add dev <SymDevName>
```

7. Create the initial relationship between the source LUNs and the target backup eLUNs through a SnapVX link operation by using the `symsnapvx` commands.

The latest *Solutions Enabler Symmetrix CLI Command Reference Guide* provides details about the `symsnapvx` commands and options. The following example shows the `symsnapvx` commands that are used to create an initial relationship between the source devices 28 and 29 and the target backup devices 9D and 9E:

- a. Create a snapshot name:

```
# symsnapvx -sid 493 establish -devs 00028:00029 -name DD_SNAPVX
```

- b. Link the source LUNs and backup device:

```
# symsnapvx -sid 493 link -devs 00028:00029 -lndevs 009D:009E -snapshot_name DD_SNAPVX -copy
Execute Link operation for Device Range (y/[n])? y
```

- c. Wait for the copy operation to finish, and check the status:

```
ledma114:/software/build147 # symsnapvx list -linked -dev 00028 -sid 493 -detail
```

Configuring ProtectPoint on VMAX

```
Symmetrix ID : 000196700493 (Microcode Version: 5977)
```

```
-----  
Sym Link Flgs Remaining Done
```

```
Dev Snapshot Name Gen Dev FCMD Snapshot Timestamp (Tracks) (%)
```

```
-----  
00028 DD_SNAPVX 0 0009D .D.X Tue Jul 22 14:36:15 2014 0 100
```

Considerations for ProtectPoint device and NetWorker ProtectPoint enabled pools

Before you configure ProtectPoint devices and pools, review the following points:

- For ProtectPoint operations, the destination pool must contain a single ProtectPoint device, and one or more devices that are not ProtectPoint devices.
- NetWorker preserves the Data Domain credentials (hostname, username, and password), which were provided during the completion of the ProtectPoint device creation wizard in the device resource. NSM retrieves and uses this information for all Data Domain communication.
- NetWorker supports only one mounted Protection Point device in a media pool.
- When the source snapshot is a ProtectPoint snapshot, then one of the listed clone operations occurs. The type of clone operation depends on the type of devices that are labeled for the destination clone pool:
 - Data Domain Replication or clone—This clone operation occurs when the source snapshot is SnapVX (V3) or RecoverPoint with XtremIO and the destination clone pool is ProtectPoint enabled. NetWorker clones or replicates the snapshot to the (remote) Data Domain that the ProtectPoint device in the destination clone pool points to. However, if the source snapshot is from a file system snapshot backup, then NetWorker clones to an alternate media device.
 - Data Domain to Data Domain (CCR/MFR) Replication—This clone operation occurs when the source snapshot is a ProtectPoint snapshot and the destination clone pool is ProtectPoint enabled. NetWorker clones or replicates the snapshot to the (remote) Data Domain.
 - Clone to media—This clone operation occurs when the destination clone pool is not ProtectPoint enabled. NetWorker clones the source snapshots to an alternate destination media. This operation replaces the *rollover* capabilities that were available in earlier versions of NSM.

Note

When you clone to a ProtectPoint device and perform Data Domain operations, configure the destination vDisk pool to use the same **Device Group Name** as the source ProtectPoint device. If the name does not exist, the NetWorker clone operation automatically creates the **Device Group Name** for the destination ProtectPoint device.

Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard

You can create or modify a ProtectPoint, RecoverPoint, and VMAX device by using the Device Configuration Wizard.

Procedure

1. In the NMC **Enterprise** view, select the NetWorker server name, and then double-click the **NetWorker** application to launch it.
2. In the **NetWorker Administration** window, click the **Devices** tab.

3. In the left panel, right-click **Devices**, and then select **New Device Wizard**.

Use the wizard to specify the options and values you need for the backup configuration.

Note

To modify completed wizard pages, click the links in the steps panel. The number of steps can vary according to the type of configuration.

4. On the **Select the Device Type** page, select **ProtectPoint** device type, and then click **Next**.
 5. On the **Data Domain Preconfiguration Checklist** page, review the requirements.
Configure the Data Domain system for ProtectPoint, and then define a DD Boost username.
 6. On the **Specify the Data Domain Configuration Options** page, configure the following attributes:
 - a. In the **Data Domain System Name** field, specify one of the options:
 - In the **Use an existing Data Domain System** field, select an existing system.
 - In the **Create a New Data Domain System** field, specify the fully qualified domain name (FQDN) or IP address of the Data Domain system.
 - b. In the **vDISK Credentials** attribute, provide the **vDISK Username** and the **vDISK Password** credentials that you used when creating vDisk device pool on the Data Domain system.
 - c. Click **Next**.
 7. On the **Select Folders to use as Devices** page, select any device folder. Each folder represents a vDisk pool on the Data Domain.
The table displays the NetWorker Device Name and Disk Pool information. Click **Next**.
 8. On the **Configure Pool Information** page, specify the following settings:
 - a. Select **Configure Media Pools for Devices**.
 - b. Specify the **Pool Type** that targets clients to the devices. For backups, select **Backup**. For cloning or staging operations, or to create a pool, select **Backup Clone**.
-

Note

If you created a media pool or if you want to use an existing media pool, ensure that the media pool does not have a media type required restriction.

If you create a pool, do not select an existing pool.

- c. Select **Label and Mount device after creation**.
 - d. Click **Next**.
9. On the **Select Storage Nodes** page, specify the following settings:

- a. Select or create the NetWorker storage node that handles the devices.
If you create a storage node, select **Dedicated Storage Node**.
 - b. Click **Next**.
10. On the **SNMP Monitoring Options** page, type the name of the Data Domain SNMP community string and specify the events that you want to monitor.
If you do not know the name of the community, then clear **Gather Usage Information**.
SNMP monitoring enables NMC to display the Data Domain system status and to list the backup and the recovery events. The monitoring feature also provides a link for launching the Data Domain interface.
 11. On the **Review the Device Configuration Settings** page, review the settings, and then click **Configure**.
NetWorker configures, mounts, and labels the ProtectPoint device for the specified pool.
 12. The **Device Configuration Results** page is informational only. To exit the wizard, click **Finish**.
 13. In the NMC **Devices** view, verify that NetWorker labeled and mounted the device, and it is ready for use. This view also lists the volume name for the device.

VMAX3 SRDF/S support

SRDF/S is a VMAX feature that maintains a synchronous, realtime copy of data at the LUN level between two VMAX storage arrays, one of which is local and the other remote.

You can configure ProtectPoint on a remote VMAX array with SRDF/S functionality, by associating a source LUN (referred to as R1) on the local array with a source LUN (R2) on the remote array. The SRDF/S software maintains continuous synchronization of the two sources by copying all the changes on one LUN device to the other.

The following SRDF/S requirements and support for snapshot operations apply to SRDF/S configurations:

- At runtime NetWorker automatically determines the state of the SRDF/S link. NetWorker does not require you to manually configure environment variables or application variables.
- If no SRDF/S link is available at the beginning of an operation, the backup or restore operation fails.
- NetWorker does not support any changes to the SRDF/S link mode made during backup or restore operations.
- If the SRDF link is in a failed over or failed back state, the snapshot operations fail.
- Mirror replication cannot transition between asynchronous and synchronous modes during any NSM operation. The mode must remain constant.
- NetWorker does not support the creation of snapshots of file systems or of volume groups that cross SRDF/RA groups.
- NetWorker supports only single-hop remote connections.

[Rollbacks in the SRDF/S environment](#) on page 88 provides details on rollback operations in this environment.

Note

NetWorker does not support SRDF/A.

Rollbacks in the SRDF/S environment

During the rollback, NetWorker automatically performs the following operations:

Note

The term *link* in this rollback process refers to the replication state, not the physical connection between the R1 and R2 devices on the separate VMAX arrays.

- Transitions the link to split between the R1 and R2 devices.
- Rolls back the data from the DD restore device to the R2 device.
- Synchronizes the R2 device to the R1 device by using reverse synchronization.
- Transitions the link to the synchronized state.
- Leaves the RDF link in a synchronized state, after the rollback is completed.

Configuring Data Domain NsrSnapSG device groups for intelligent pairing

RecoverPoint/ProtectPoint leverage Data Domain vDisk for PiT validation, mount, and restore operations. Designate or create a Data Domain device group and populate with an appropriate number of restore vDisk devices. The number of devices depends on the number of source XtremIO devices that are used by the application, and if the device group is to be shared between multiple hosts the number of potential concurrent restore or mount operations.

Intelligent Pairing vDisk selection decision tree

RecoverPoint/ProtectPoint selects vDisk devices based on a cascading decision tree and is controlled by you.

The following criteria are used:

- Determine whether there is a unique instance of a Data Domain (DD) disk group name "NsrSnapSG" (or a user-provided name) on the DD, and get its DD pool name as follows:
 - If the user provides the Data Domain pool name.
 - If the ddvdisk pool contains a single disk group, use it, regardless of its name.
 - If the pool contains a NsrSnapSG disk group, use it.
 - Otherwise, fail.
- If you provide the Data Domain disk group name:
 - Get the list of pools containing this disk group name.
 - If there is 0 or > 1 pool, fail, otherwise, use it.
- When neither pool nor disk group name is provided:

- Get the list of pools containing disk group name "NsrSnapSG."
- If there is 0 or > 1 pool, fail, otherwise use it.

Intelligent Pairing allocates vDisk for mount, validate, and restore

When over subscribing the vDisk devices to multiple hosts a vDisk is locked once Intelligent Pairing has selected it. This locking prevents the device to allocate to a concurrent mount/validation or restore operations. The lock is released to the pool for subsequent operations.

Locking of vDisk LUN for restore/mount operations is done by adding a specific key-value metadata key to the vDisk. Intelligent pairing inspects the vDisk for the presence of the key-value pair and if absent considers the vDisk "available." It removes the key at the end of the operation.

Both the Data Domain pool and disk group that is used by IP can optionally be specified using client resource attributes. These must be typed manually into the client resource at the end of client configuration:

- `NSR_DD_VDISK_RESTORE_POOLNAME=<dd vdisk restore pool>`
- `NSR_DD_VDISK_RESTORE_DEVGRPNAME=<dd vdisk restore devgrp>`

CHAPTER 5

Configuring ProtectPoint on RecoverPoint with XtremIO

This chapter includes the following topics:

- [Overview](#)92
- [Basic backup workflow](#)92
- [Basic restore workflow](#)92
- [ProtectPoint for RecoverPoint on XtremIO prerequisites](#) 93
- [Enabling vdisk on the Data Domain](#) 94
- [Provisioning protection devices on Data Domain systems](#)95
- [Configuring RecoverPoint and XtremIO storage](#) 96
- [Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard](#) 97
- [Configuration for restore to secondary VMAX with the wizard](#) 99
- [Considerations for ProtectPoint device and NetWorker ProtectPoint enabled pools](#)100
- [Configuring Data Domain NsrSnapSG device groups for intelligent pairing](#)100

Overview

The ProtectPoint for XtremIO system is an integration of XtremIO, Data Domain, and RecoverPoint technologies. This system consists of all RecoverPoint appliances (RPAs) that are used to replicate and protect data at the production site, an XtremIO journal at the production site, and a catalog at the Data Domain copy site. The solution enables an application administrator to leverage the ProtectPoint workflow to protect applications and application data. The storage administrator configures the underlying storage resources on the primary storage system and the Data Domain system. With this storage configuration information, RecoverPoint, and the ProtectPoint feature from NetWorker, the application administrator can trigger the workflow to protect the application or file system using familiar NetWorker NMDA/NMSAP, file system integration methods. To use ProtectPoint backup as the first line of protection, create policies with a snapshot action that send data to a destination pool that contains a ProtectPoint device.

This section explains how to provision Data Domain backup and restore devices, create a NetWorker ProtectPoint device, and label the device for a NetWorker pool.

Basic backup workflow

In the basic backup workflow, data is transferred from the primary storage system to the Data Domain system. NetWorker NSM integration with ProtectPoint manages the data flow, but does not modify the data.

After creating the snapshot, RecoverPoint moves the snapshot to the Data Domain system. The primary storage system keeps track of the data that has changed since the last update to the Data Domain system, and only copies the changed data. Once all the data that is captured in the snapshot has been moved, the Data Domain system creates a static image of the data that reflects the application-consistent copy that is initially created on the primary storage system. The snapshot is then recorded by the NetWorker server.

The backup workflow consists of the following steps:

1. The file system or application backup is started using a NetWorker snapshot enabled policy or from the application host using NMDA or NMSAP integration.

Note

File system backups cannot be client initiated.

2. On the primary storage system, ProtectPoint creates a snapshot of the primary storage device.
3. RecoverPoint analyzes the data, and then copies the changed data to a Data Domain storage device.
4. The Data Domain system creates, and then stores a static-image of the snapshot.
5. NetWorker NSM creates a catalog record of the snapshot.

Basic restore workflow

ProtectPoint with RecoverPoint allows application administrators to start an object level or granular file-by-file restore directly from the restore devices on the Data Domain system to the AR host without involving the primary storage or the

RecoverPoint cluster. ProtectPoint with RecoverPoint supports object-level restores, and rollback restores.

For an object-level restore, after selecting the backup image on the Data Domain system, the application administrator restores the data to a new set of Data Domain block services for ProtectPoint devices (restore devices) to present to the AR host, then copies individual files back to the production devices.

For a rollback restore, after selecting the backup, NSM passes control to the RecoverPoint appliance to start the XtremIO array level restore.

Note

You cannot perform a ProtectPoint RecoverPoint rollback using a cloned CCR copy (secondary DD copy). A RecoverPoint rollback restore is always performed by RecoverPoint, and it is only aware about the copy on the local Data Domain. The CCR/clone copy is created by NSM, by directly copying from one DDR to the remote DDR. You can use the CCR/clone copy only for an NSM PIT restore or clone to media [rollover].

ProtectPoint for RecoverPoint on XtremIO prerequisites

Ensure that you meet the following prerequisite before you use ProtectPoint devices:

- Use an XtremIO array as the backend to the RecoverPoint appliance.

Ensure that you meet the following prerequisites for all ProtectPoint operations:

Data Domain

- Have a Data Domain system that is supported by ProtectPoint.
- Block services for ProtectPoint must be enabled on the Data Domain system.
- Data Domain Boost must be enabled on the Data Domain system.

RecoverPoint

- RecoverPoint backup (BK) license.
- The RecoverPoint cluster must have Gen5 or later RecoverPoint appliances (RPAs), running RecoverPoint version 4.4.x or later.
- Port 443 must be open between the RecoverPoint appliances, the XtremIO Management System, and the XtremIO System-wide Management (SYM) module on X1-Storage Controller 1 (X1-SC1) and X1-Storage Controller 2 (X1-SC2) IP addresses.
- Port 11111 must be open between the RPAs and XtremIO SYM module on X1-SC1 and X1-SC2.
- IP connectivity must be configured between the RPA and the Data Domain system.
- FC zoning must be configured between the RPA and the XtremIO cluster, and is optional between the RPA and the Data Domain system.

Note

FC zoning between the RPA and the Data Domain system is only required if DD Boost over FC is used for communication between the RPA and the Data Domain system.

- Zone at least two initiators from the RPA to the Data Domain system.
- Create one zone per fabric between the RPA and the XtremIO cluster, and include all the RPA ports that are intended for XtremIO connectivity and all the XtremIO ports that are intended for RPA connectivity in the zone.
- Zone the RecoverPoint initiators to multiple targets on the Data Domain system or XtremIO cluster in accordance with RecoverPoint best practices.

XtremIO

- FC zoning must be configured between the XtremIO cluster and the AR host:
 - Use a single-initiator per single-target (1:1) zoning scheme. If the FC switchzone count limitation has been reached, it is also possible to use single initiator per multiple-target (1:many) zoning scheme.
 - The optimal number of paths depends on the operating system and server information. To avoid multipathing performance degradation, do not use more than 16 paths per device.
 - Enable MPIO if two or more paths are zoned to a Windows AR host.

Configuring ProtectPoint with NetWorker

1. On the Data Domain, enable vDisk.
2. On Data Domain, provision protection of the vDisk devices.
3. On Data Domain, create restore vDisk devices.
4. Create a NetWorker device, and then label it under a NetWorker pool using the wizard.
5. Create snapshot policies, follow the steps in the [Data Protection Policies](#) section, and then set the destination pool to a pool containing a ProtectPoint device.
6. Follow the steps in the [Configuring snapshot backups with the client wizard](#) on page 68 or [Configuring snapshot backups manually](#) on page 74 sections, and then set the options accordingly.

Enabling vdisk on the Data Domain

Enable vDisk on a Data Domain system through the vDisk enable command. Use the Data Domain command line interface to complete the required administration tasks.

The *Data Domain Operating System Command Reference Guide* provides details about the commands.

Procedure

1. Log in to the Data Domain system as an administrative user.
2. To verify that the vDisk license is enabled, type the following command:

```
# license show
```

```
Feature licenses:
## License Key Feature
-----
1 ABCD-EFGH-IJKL-MNOP DDBOOST
-----
```

If the DD vDisk license is disabled, type the following command to add the vDisk license by using the license key provided:

```
# license add license_key
```

```
License "ABCE-BCDA-CDAB-DABC" added.
```

3. To enable the vDisk service subsystem, type the following command:

```
# vdisk enable
```

```
DD vDisk enabled
```

4. To verify that the vDisk service is enabled, type the following command :

```
# vdisk status
```

```
vDisk admin state: enabled, process is running, licensed
```

Provisioning protection devices on Data Domain systems

The Data Domain administrator must create the required vdisk device pool and device group and configure the backup and restore devices in the vdisk device pool and device group. You (DD Administrator) can use the vdisk mappings in the following table to plan the configuration.

Table 10 vdisk object hierarchy mapping

vdisk storage object	Mapping level
Device pool	NetWorker Datazone
Device group	Application
Device	Source LUN

Note

A NetWorker ProtectPoint device represents a specific vdisk pool.

The *Data Domain Operating System Administration Guide* provides the latest information about the vdisk configuration and any limitations.

Procedure

1. To create the vdisk device pool and device group, type the following vdisk commands:

```
vdisk pool create pool_name user username
```

```
vdisk device-group create device_group_name pool pool_name
```

The *Data Domain Operating System Command Reference Guide* provides details about the vdisk command and options.

2. For each XtremIO source LUN that contains data, ensure that you have one backup device and one restore device. You might need additional restore devices for each source LUN if you also plan to restore data to different hosts.
3. To create the backup and restore devices with the same geometry as the XtremIO source LUNs, and provision the devices, type the following vdisk command:

```
vdisk device create capacity <n> {GiB} pool pool_name device-
group device_group_name
```

Note

By default NSM's Intelligent Pairing looks for a Data Domain device group of name NSRSnapSG. It is recommended to create the device group with this name. Optionally any device group name can be created but the user must then specify this name and the value to NSR_DD_VDISK_RESTORE_POOLNAME and NSR_DD_VIDKS_RESTORE_DEVGRPNAME attributes in the NetWorker client resource.

4. On the Data Domain system, add the vdisk devices to the access group:

- a. Create a vdisk access group:

```
scsitarget group create group_name service vdisk
```

- b. Add all the vdisk devices to the vdisk access group:

```
vdisk group add group_name pool vdisk_pool_name device-group
device_group_name
```

5. Add XtremIO initiators to the group:

```
vdisk group add group_name initiator initiator_name
```

Configuring RecoverPoint and XtremIO storage

RecoverPoint and XtremIO storage configuration for ProtectPoint consists of provisioning storage on the XtremIO cluster, integrating the XtremIO cluster with the RecoverPoint cluster, and creating the backup and restore devices on the Data Domain system. The following tasks are for configuring storage for new installations. Some of these tasks may not be required always.

Procedure

1. Complete the following tasks on the XtremIO cluster:
 - a. Create an Initiator Group, and register all the RecoverPoint cluster FC ports to the Initiator Group.
 - b. Provision Volumes on the XtremIO cluster, and then make the Volumes available to the AR host.
2. Complete the following tasks on the RecoverPoint cluster:
 - a. Register the XtremIO Management Server.
 - b. Register the Data Domain system.
 - c. Create a RecoverPoint consistency group that contains all the XtremIO Volumes created for the ProtectPoint environment.
3. On the Data Domain system, complete the following tasks:
 - a. Login to an SSH session.
 - b. If it is not already enabled, enable block services for ProtectPoint.
 - c. If it is not already enabled, enable DD Boost.
 - d. Create a block-services for ProtectPoint pool where the ProtectPoint static-images will reside.

- e. To function as the destination for restore data from the RecoverPoint cluster, on the ProtectPoint MTree, create a storage unit.
- f. Create two block services for ProtectPoint device pools, one for backups and one for restores.

Note

After creating the block services for ProtectPoint device pool, RecoverPoint automatically creates the block services for ProtectPoint device-groups and devices that are required for ProtectPoint with RecoverPoint backups.

- g. Create a block-services for ProtectPoint device-group for restores that resides in the restore pool that is created in Step 3 e.
- h. Populate the restore device-group with restore devices of the same size as the production LUNs you are backing up.

Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard

You can create or modify a ProtectPoint, RecoverPoint, and VMAX device by using the Device Configuration Wizard.

Procedure

1. In the NMC **Enterprise** view, select the NetWorker server name, and then double-click the **NetWorker** application to launch it.
2. In the **NetWorker Administration** window, click the **Devices** tab.
3. In the left panel, right-click **Devices**, and then select **New Device Wizard**.

Use the wizard to specify the options and values you need for the backup configuration.

Note

To modify completed wizard pages, click the links in the steps panel. The number of steps can vary according to the type of configuration.

4. On the **Select the Device Type** page, select **ProtectPoint** device type, and then click **Next**.
5. On the **Data Domain Preconfiguration Checklist** page, review the requirements.
Configure the Data Domain system for ProtectPoint, and then define a DD Boost username.
6. On the **Specify the Data Domain Configuration Options** page, configure the following attributes:
 - a. In the **Data Domain System Name** field, specify one of the options:
 - In the **Use an existing Data Domain System** field, select an existing system.
 - In the **Create a New Data Domain System** field, specify the fully qualified domain name (FQDN) or IP address of the Data Domain system.

- b. In the **vDISK Credentials** attribute, provide the **vDISK Username** and the **vDISK Password** credentials that you used when creating vDisk device pool on the Data Domain system.
 - c. Click **Next**.
7. On the **Select Folders to use as Devices** page, select any device folder. Each folder represents a vDisk pool on the Data Domain.

The table displays the NetWorker Device Name and Disk Pool information. Click **Next**.

8. On the **Configure Pool Information** page, specify the following settings:
 - a. Select **Configure Media Pools for Devices**.
 - b. Specify the **Pool Type** that targets clients to the devices. For backups, select **Backup**. For cloning or staging operations, or to create a pool, select **Backup Clone**.

Note

If you created a media pool or if you want to use an existing media pool, ensure that the media pool does not have a media type required restriction.

If you create a pool, do not select an existing pool.

- c. Select **Label and Mount device after creation**.
 - d. Click **Next**.
9. On the **Select Storage Nodes** page, specify the following settings:
 - a. Select or create the NetWorker storage node that handles the devices.

If you create a storage node, select **Dedicated Storage Node**.
 - b. Click **Next**.
10. On the **SNMP Monitoring Options** page, type the name of the Data Domain SNMP community string and specify the events that you want to monitor.

If you do not know the name of the community, then clear **Gather Usage Information**.

SNMP monitoring enables NMC to display the Data Domain system status and to list the backup and the recovery events. The monitoring feature also provides a link for launching the Data Domain interface.
11. On the **Review the Device Configuration Settings** page, review the settings, and then click **Configure**.

NetWorker configures, mounts, and labels the ProtectPoint device for the specified pool.
12. The **Device Configuration Results** page is informational only. To exit the wizard, click **Finish**.
13. In the NMC **Devices** view, verify that NetWorker labeled and mounted the device, and it is ready for use. This view also lists the volume name for the device.

Configuration for restore to secondary VMAX with the wizard

To restore the snapshot to secondary VMAX, perform the following:

Procedure

1. Add the parameter
`NSM_SNAP_SG=<secondary_VMAX_symmid>:restore_storage_group`
in the advanced option in the restore wizard of NMC.
-

Note

For the file level restore, restore host must have the primary VMAX LUNs to mount the snapshot.

Considerations for ProtectPoint device and NetWorker ProtectPoint enabled pools

Before you configure ProtectPoint devices and pools, review the following points:

- For ProtectPoint operations, the destination pool must contain a single ProtectPoint device, and one or more devices that are not ProtectPoint devices.
- NetWorker preserves the Data Domain credentials (hostname, username, and password), which were provided during the completion of the ProtectPoint device creation wizard in the device resource. NSM retrieves and uses this information for all Data Domain communication.
- NetWorker supports only one mounted Protection Point device in a media pool.
- When the source snapshot is a ProtectPoint snapshot, then one of the listed clone operations occurs. The type of clone operation depends on the type of devices that are labeled for the destination clone pool:
 - Data Domain Replication or clone—This clone operation occurs when the source snapshot is SnapVX (V3) or RecoverPoint with XtremIO and the destination clone pool is ProtectPoint enabled. NetWorker clones or replicates the snapshot to the (remote) Data Domain that the ProtectPoint device in the destination clone pool points to. However, if the source snapshot is from a file system snapshot backup, then NetWorker clones to an alternate media device.
 - Data Domain to Data Domain (CCR/MFR) Replication—This clone operation occurs when the source snapshot is a ProtectPoint snapshot and the destination clone pool is ProtectPoint enabled. NetWorker clones or replicates the snapshot to the (remote) Data Domain.
 - Clone to media—This clone operation occurs when the destination clone pool is not ProtectPoint enabled. NetWorker clones the source snapshots to an alternate destination media. This operation replaces the *rollover* capabilities that were available in earlier versions of NSM.

Note

When you clone to a ProtectPoint device and perform Data Domain operations, configure the destination vDisk pool to use the same **Device Group Name** as the source ProtectPoint device. If the name does not exist, the NetWorker clone operation automatically creates the **Device Group Name** for the destination ProtectPoint device.

Configuring Data Domain NsrSnapSG device groups for intelligent pairing

RecoverPoint/ProtectPoint leverage Data Domain vDisk for PiT validation, mount, and restore operations. Designate or create a Data Domain device group and populate with an appropriate number of restore vDisk devices. The number of devices depends on the number of source XtremIO devices that are used by the application, and if the device group is to be shared between multiple hosts the number of potential concurrent restore or mount operations.

Intelligent Pairing vDisk selection decision tree

RecoverPoint/ProtectPoint selects vDisk devices based on a cascading decision tree and is controlled by you.

The following criteria are used:

- Determine whether there is a unique instance of a Data Domain (DD) disk group name "NsrSnapSG" (or a user-provided name) on the DD, and get its DD pool name as follows:
 - If the user provides the Data Domain pool name.
 - If the dddisk pool contains a single disk group, use it, regardless of its name.
 - If the pool contains a NsrSnapSG disk group, use it.
 - Otherwise, fail.
- If you provide the Data Domain disk group name:
 - Get the list of pools containing this disk group name.
 - If there is 0 or > 1 pool, fail, otherwise, use it.
- When neither pool nor disk group name is provided:
 - Get the list of pools containing disk group name "NsrSnapSG."
 - If there is 0 or > 1 pool, fail, otherwise use it.

Intelligent Pairing allocates vDisk for mount, validate, and restore

When over subscribing the vDisk devices to multiple hosts a vDisk is locked once Intelligent Pairing has selected it. This locking prevents the device to allocate to a concurrent mount/validation or restore operations. The lock is released to the pool for subsequent operations.

Locking of vDisk LUN for restore/mount operations is done by adding a specific key-value metadata key to the vDisk. Intelligent pairing inspects the vDisk for the presence of the key-value pair and if absent considers the vDisk "available." It removes the key at the end of the operation.

Both the Data Domain pool and disk group that is used by IP can optionally be specified using client resource attributes. These must be typed manually into the client resource at the end of client configuration:

- NSR_DD_VDISK_RESTORE_POOLNAME=<dd vdisk restore pool>
- NSR_DD_VDISK_RESTORE_DEVGRPNAME=<dd vdisk restore devgrp>

CHAPTER 6

Configuring snapshots on XtremIO arrays

This chapter includes the following topics:

- [Snapshot support for XtremIO](#).....104
- [Snapshot operation with XtremIO REST API](#)..... 104
- [Prerequisite for XtremIO configurations](#)..... 104
- [Supported XtremIO features](#)..... 104
- [Snapshot management policy with XtremIO](#)..... 104
- [Snapshot backups with XtremIO](#)..... 105
- [XtremIO configuration methods](#)..... 105

Snapshot support for XtremIO

This section describes NetWorker snapshot support, practices, and configurations that are specific to XtremIO arrays.

Note

For NetWorker Snapshot Management (NSM) with native XtremIO, a virtual machine cannot be used as a proxy or storage node.

Snapshot operation with XtremIO REST API

The XtremIO REST API provides snapshot capability of application volumes. During production operations, XtremIO consistency groups and non-consistency based snapshot support. For consistency groups based snapshot, all members of the consistency group will be snap regardless of what client save set you specified.

Prerequisite for XtremIO configurations

The NetWorker Snapshot Management (NSM) prerequisites for XtremIO arrays are as follows:

- XtremIO Storage Array version 4.0.1 or later is required for RESTful API snapshot support.
- Create the Proxy/Mount host initiator name on the XtremIO array, and type it into the client resource as NSR_XTREMIO_PROXY_INITIATOR_NAME attribute.

Supported XtremIO features

The NSM integration with XtremIO supports the following backup and restore capabilities:

- Snapshot backup (PIT)
- Pit mount on mount host
- Clone (backup of PIT)
- Snapshot restore
- Snapshot management
- Rollback is not supported

Snapshot management policy with XtremIO

NSM XtremIO creates snapshots and makes the snapshot available to NetWorker by saving them into the media database as snapshot save sets (snapsets). The backup administrator can use the NMC Client Configuration Wizard to manage volumes that are protected by XtremIO.

Before you configure a NetWorker client with the configuration wizard, have the following information ready:

- The XtremIO Storage array hostname or IP address.

- The username and password for the XtremIO storage array.
- The mount host initiator name created on the XtremIO array.

Snapshot backups with XtremIO

During scheduled backups, NSM for XtremIO creates PIT copies of the production LUN. During the snapshot process, NSM extracts the XtremIO array credentials from the NetWorker Client resource to discover production LUN information to create a snapshot for.

XtremIO configuration methods

The NMC Client Configuration Wizard supports the creation of NSM XtremIO configurations. However, after you have created an XtremIO configuration, modify the configuration through the **NMC Properties** windows to add the proxy host initiator name, or to make any other changes. You can also add the proxy host initiator name in the NMC Client Configuration Wizard from **Advanced options > Extra Options**.

CHAPTER 7

Configuring snapshots on VMAX Storage Arrays

This chapter includes the following topics:

- [Snapshot support of VMAX storage arrays](#)..... 108
- [Pairing source LUNs to mirror LUNs](#)..... 109
- [VMAX SRDF/S support](#)..... 112
- [Solutions Enabler Client and Server mode configuration](#)..... 113
- [Known limitation for VMAX](#)..... 114

Snapshot support of VMAX storage arrays

NetWorker provides snapshot support and configurations that are specific to VMAX (Symmetrix) storage arrays.

[Migrating Legacy PowerSnap Configurations](#) provides information on migrating PowerSnap VMAX implementations to NetWorker snapshot management.

Snapshot operations with TimeFinder software

To perform snapshot operations on VMAX storage arrays, the TimeFinder Solutions Enabler package must be installed on both the application host and the mount host (if separate from the application host). Solutions Enabler manages all NetWorker TimeFinder client and server operations.

The TimeFinder software maintains multiple, host-independent copies of production data by generating synchronous real-time mirrors of the production data. TimeFinder can use locally mirrored LUNs on the same VMAX array or local LUNs with remotely mirrored LUNs on a separate VMAX array with an SRDF connection.

The NetWorker intelligent pairing feature automatically matches TimeFinder source LUNs with appropriate target mirror devices. This feature replaces a manually configured `symm.res` file on the application host. However, if a `symm.res` file is present, the file takes priority.

Prerequisites and support for VMAX configurations

The following prerequisites and support apply to VMAX configurations:

- Ensure that Solutions Enabler is installed on the application host and the mount host. [Storage array prerequisites](#) provides details.
- Mask the mirror devices to the mount host. You can run the `symdev show mirror_device` command on the mount host to verify that the device has a physical pathname.
- Create a SYM Access Group which contains the application host and the mount host.
- To enable snapshots you must add the following privileges to the Access Groups:
 - BASE
 - BASECTRL
 - BCV
- VMAX configurations support all thin-provisioned LUNs.
- Whenever a storage layout change affects the application host or the mount host, run the `symcfg discover` command on these hosts to rediscover the storage layout. If snapshots fail, you can run this command on the mount host to find mirror devices that are not visible to the mount host.
- Solutions Enabler Client/Server mode is supported. [Solutions Enabler Client and Server mode configuration](#) on page 113 provides more details.

Types of supported mirror devices

NetWorker supports the following types of mirror devices on VMAX storage arrays:

- TimeFinder BCV—These devices are full physical copies and appear as mirrors of the standard device.
- TimeFinder Snap/VDEV—These devices use the copy-on-write (COW) snapshot creation method.
- TimeFinder Clone—These devices create high performance, full source copies. The following limitations apply for rollback operations that use TimeFinder Clone copies:
 - Rollback operations fail for a snapshot that is created with the application variable SYMM_CLONE_FULL_COPY=FALSE.
 - Rollback operations fail for a source LUN that has another established BCV mirror unless you set SYMM_RB_OVERRIDE_OTHER_TGTS=TRUE.
 - Rollback operations fail for a source LUN that has an active relationship with more than one snapshot or mirror.

[Application Information variables for VMAX arrays](#) provides details.

- TimeFinder VP Snap—These devices create space-efficient snapshots for virtual thin pool devices. The following limitations apply:
 - All VP Snap target devices that are paired to the same source LUN must be bound to the same thin pool.
 - A source device cannot simultaneously run both a VP Snap session and a Clone No Full Copy session.
- TimeFinder SnapVX—These devices are a fundamentally new TimeFinder, with snapshots now existing as pointers rather than as physical devices:
 - You need not specify a target device and the source and target pairs.
 - When you perform a rollback with SnapVX snapshots, the snapshot is not deleted from the NetWorker media database and is available again for subsequent rollback operations.
 - SnapVX only supports target LUN selection using Intelligent Pairing. If a `symm.res` is used the selection is ignored.

Pairing source LUNs to mirror LUNs

NetWorker snapshot operations require the use of paired source and target mirror LUNs. NetWorker intelligent pairing can automatically determine these LUN pairs or you can manually specify the pairs in the `symm.res` file on the application host.

Intelligent pairing

Intelligent pairing is a NetWorker feature that automatically chooses an available mirror LUN, based on the mirror that is the least expensive to synchronize with the source LUN.

Intelligent pairing selects only mirrors that are visible and usable by the snapshot mount host, which can be separate from the application host. This feature eliminates the potential error in manual configuration, which can have new LUNs masked only to the application host when they should also be masked to the mount host.

Intelligent pairing is now user configurable. Use the `NSM_SNAP_SG` client resource attribute. Add this attribute manually to the client resource after you create the client. Using the `NSM_SNAP_SG` client resource attribute increases performance and the number of LUNs are reduced.

Intelligent pairing selects mirror LUNs from a pool of LUNs that you specify in a VMAX storage group (NsrSnapSG) on each VMAX array that NetWorker uses:

- Each storage group can contain a maximum of 4096 LUNs.
- This storage group can contain any type of LUN, except source LUNs.
- Ensure that you add sufficient numbers, types, and sizes of devices to a storage group so that intelligent pairing can find compatible pairs. For example, for Clone and VP Snap operations, a source LUN requires the use of STD or BCV devices as mirrors.
- If NetWorker cannot find a valid mirror, then the snapshot fails with the following message:

```
Not enough resources.
```

NOTICE

Do not use the device LUNs in the NsrSnapSG storage groups for any purpose other than as NetWorker snapshot mirror devices. The snapshot operations destroy the contents of any device selected from an intelligent pairing storage group. Do not add source LUN devices to a storage group.

The snapshot operations can pair a mirror LUN with only one source LUN at a time. On rare occasions, more than one application host can simultaneously try to use the same free mirror LUN for a backup operation. One backup will succeed and the competing backup will fail. Retry the failed backup and NetWorker will use a different mirror.

Configuring NsrSnapSG storage groups for intelligent pairing

You can create a maximum of 11 intelligent pairing storage groups on a VMAX array. The names of the storage groups can be NsrSnapSG and NsrSnapSG0 through NsrSnapSG9. Using the NSM_SNAP_SG client resource attribute increases performance and the number of LUNs are reduced.

Each VMAX storage group can contain up to 4096 devices. If you need to specify more than 4096 mirrors for intelligent pairing, create more than one storage group.

Note

NsrSnapSG storage groups are created on the VMAX arrays, not on the application host where you run the command. Any application host that can see the VMAX array can see its storage groups.

To create a storage group for intelligent pairing, run the following command: `symmsg -sid vmax_id create NsrSnapSG`

To add a device to this group that NetWorker can use as a mirror, run the following command: `symmsg -sid vmax_id -sg NsrSnapSG add dev device_id`

Note

These examples use mixed case for clarity. The characters in the storage group names are not case-sensitive.

Manual pairing LUNs with the `symm.res` file

You can use the `symm.res` file instead of, or also with, intelligent pairing. This file enables you to manually select mirrors to pair to specific source LUNs. SnapVX does

not support `symm.res`. SnapVX requires targets for mounting the snapshot which you can select only by using IP.

Note

If you do not correctly configure the `symm.res` file to mask the mirrors to the mount host, NetWorker creates snapshots that are not available for restore or for clone operations.

Consider the following features of the `symm.res` file and intelligent pairing:

- If the `symm.res` file specifies a source LUN, that LUN cannot use intelligent pairing, even if none of its specified mirrors are valid or even if the `symm.res` file is incorrect.
- Source LUNs are available to intelligent pairing only if the `symm.res` file does not specify them or the `symm.res` file does not exist.
- You can disable intelligent pairing by specifying the Application Information variable `NSR_PS_SYMM_IP=FALSE`.

Configuring the `symm.res` file

The default location of the `symm.res` file is `/nsr/res/symm.res`. You can specify an alternate location by using the Application Information variable `SYMM_SNAP_POOL=pathname`.

The file consists of one or more lines, each in the following format:

```
vmax_id:source_dev_id
[vmax_id:]mirror_dev_id1[vmax_id:]mirror_dev_id2
```

When you specify one or more mirror LUNs for a single source LUN, NetWorker pairs the best mirror LUN with the source LUN. If NetWorker cannot find a pair, the backup fails with the following message:

```
Not enough resources.
```

NetWorker ignores blank lines in the `symm.res` file and lines starting with `#`.

The `vmax_id` for the source LUN is mandatory. If the selected mirror LUN does not have a VMAX-id, the mirror LUN uses the VMAX-id of the source LUN.

For SRDF configurations, the `vmax_id` of the mirror LUNs is mandatory and must be different from the `vmax_id` of the source LUN.

For example, a simple `symm.res` file for a source LUN ABC with 3 mirror LUNs 123, 456, and 789 can be as follows:

```
# LUNs for /critical_filesystem
000194901248:ABC 123 456 789
```

For an SRDF configuration, the same `symm.res` file could be as follows:

```
000194901248:ABC 0001949017BA:123 0001949017BA:456 0001949017BA:789
```

VMAX SRDF/S support

SRDF/S is a VMAX feature that maintains a synchronous, real-time copy of data at the LUN level between two VMAX storage arrays, one of which is local and the other remote.

To configure snapshots on a remote VMAX array with SRDF/S functionality, you must associate a source LUN (referred to as R1) on the local array with a source LUN (R2) on the remote array. The SRDF/S software maintains continuous synchronization of the two sources by copying all changes on one LUN device to the other.

For typical snapshot operations, the remote R2 LUN has its own mirror that NetWorker uses for snapshot creation and snapshot clones. A mirror is optional on the local R1 LUN. When NetWorker creates an SRDF/S snapshot, it validates the synchronization of the R1 and R2 devices and then syncs/splits the mirror of the R2 device. This split of the mirror creates the snapshot, which represents a third copy of the data. If the NetWorker policy specifies a clone, then NetWorker performs the clone from this mirror.

NetWorker also supports a direct R2 backup, with no snapshot. In this operation, NetWorker ensures the synchronization of the R1 and R2 devices, suspends the link between the R1 and R2 devices, and performs a clone directly from the R2 source LUN. After the clone completes, NetWorker reestablishes the link. The snapshot policy must have a Retain Snapshots value of 0 or the clone will fail.

SRDF/S requirements and support for snapshot operations are as follows:

- When using a direct R2 backup, R2 mirrors must be visible and available to the mount host.
- NetWorker automatically determines the state of the SRDF/S link at runtime. There is no requirement for manually configured environment variables or application variables.
- If there is no SRDF/S link at the beginning of an operation, then the backup or restore operation will fail.
- NetWorker does not support any changes to the SRDF/S link mode made during backup or restore operations.
- If the RDF link is in a failed over or failed back state, the snapshot operations will fail.
- Mirror replication cannot transition between asynchronous and synchronous modes during any NSM operation. The mode must remain constant.
- NetWorker does not support the creation of snapshots of file systems or of volume groups that cross SRDF/RA groups.
- NetWorker supports only single-hop remote connections.
- NetWorker supports only TF/Mirror (RBCV) in asynchronous copy (SRDF/A) environments and does not support concurrent RDF and STAR configurations and modes.

[Rollbacks in the SRDF/S environment](#) on page 88 provides specific details on rollback operations in this environment.

Rollbacks in the SRDF/S environment

During the rollback, NetWorker automatically performs the following operations:

Note

The term *link* in this rollback process refers to the replication state, not the physical connection between the R1 and R2 devices on the separate VMAX arrays.

- Transitions the link to split between the R1 and R2 devices.
- Rolls back the data from the DD restore device to the R2 device.
- Synchronizes the R2 device to the R1 device by using reverse synchronization.
- Transitions the link to the synchronized state.
- Leaves the RDF link in a synchronized state, after the rollback is completed.

Solutions Enabler Client and Server mode configuration

You can run SYMCLI as a client to a remote SYMAPI server to manage a remotely-controlled Symmetrix array.

The *Solutions Enabler Installation Guide* provides more information on SYMAPI Client/Server configuration.

Solutions Enabler in Client and Server mode configuration

To enable VMAX backups with the Solutions Enabler in remote server mode, ensure that you meet the following requirements:

On the client or mount or both host:

- The Solutions Enabler clients are installed.
- Gatekeepers are not configured and masked to the client, mount or both host.
- The `netcnfg` file in the SYMAPI configuration directory includes the `NSM_SERVER` setting, which specifies the `SYMAPI_SERVER` name
`NSM_SERVER - TCP/IP <SYMAPI SERVER HOST NAME> <SYMAPI SERVER IP> <PORT NUMBER> ANY`
- The connection environment variable `SYMCLI_CONNECT_TYPE` which defines the local or remote mode of the local host (client), is set to the REMOTE value
`SYMCLI_CONNECT_TYPE=REMOTE.`
- To use SYMCLI through a remote SYMAPI service, the environment variable `SYMCLI_CONNECT` is set to an available service name of the server connection as defined in the `netcnfg` file. For the service name `NSM_SERVER`, set the environment variable as `SYMCLI_CONNECT=NSM_SERVER.`

On the SYMAPI server node:

- The operating system must support the Solutions Enabler as configured with the SYMAPI server.
 - The Solutions Enabler server is installed and configured in server mode.
 - In-band SAN connectivity is enabled to the VMAX array.
-

Note

Solutions Enabler in client or server mode works with Solutions Enabler 8.4.0.7 and Solutions Enabler 9.0. If Networker is running before Solutions Enabler configuration in remote mode, then we need to restart Networker snapshot management service on client, mount or both host.

Known limitation for VMAX

When you use SUSE LINUX server as a mount host, it is known to require specific configuration when you add LUNs which can generate extra load on the system administrator side.

CHAPTER 8

Configuring snapshots on VNX Block Storage Arrays

This chapter includes the following topics:

- [Snapshot support of VNX Block storage arrays](#).....116
- [Configuring the Navisphere security file](#).....116
- [Configuring Unisphere CLI on VNXe3200](#).....118

Snapshot support of VNX Block storage arrays

NetWorker provides snapshot support and configurations that are specific to VNX Block (CLARiiON) storage arrays.

[Migrating Legacy PowerSnap Configurations](#) provides details on migrating PowerSnap VNX implementations to NetWorker Snapshot Management (NSM).

Snapshot operations with SnapView software

VNX Block storage arrays run SnapView software that enables you to create a copy of a LUN by using clones or snapshots. You can use a clone or snapshot for data backups. A clone is a complete copy of a LUN and takes time to create. A snapshot is a virtual point-in-time copy of a LUN and takes only seconds to create. A NetWorker snapshot operation uses SnapView to create an exact point-in-time snapshot of the volume, creates a clone copy of the entire volume that NetWorker can recover, and clones to conventional storage media.

The *VNX Series Command Line Interface Reference for Block* documentation provides details.

Prerequisites and support for VNX configurations

Install Unisphere host agent, Navisphere CLI, and Snap CLI on the application host and the snapshot mount host. [Storage array specific prerequisites](#) on page 65 provides details.

NetWorker supports the following backup technologies on VNX Block storage arrays:

- Copy on Write (SnapView COW/snapshot)
- MIRROR (SnapView Clone)
- VNX-SNAP (VNX Snapshot)
- VNXe-SNAP (VNXe/VNXe2 Snapshot)

Note

While configuring a Solaris client (with Veritas) for a VNXe Snap backup, you must disable the Veritas controller on the proxy host. Use the following command to disable the Veritas controller:

```
vxdmpadm -f disable ctlr=c4
```

Configuring the Navisphere security file

The Navisphere security file is required on all nodes that participate in snapshot operations with VNX storage arrays. The security file enables VNX `naviseccli` commands for cloning and other features. If this file does not exist or does not contain the permissions that are required by NetWorker, then the NetWorker backups and restores fails.

If a Navisphere security file does not exist, the NetWorker Client Configuration wizard creates the file under the root user home directory (UNIX) or System account (Windows). The security file can be manually created and modified.

Creating the Navisphere file manually on UNIX systems

To manually configure the Navisphere security file on UNIX systems, type the following command:

```
naviseccli -h VNX_server -addusersecurity -user VNX_array_user -
password VNX_array_user_password -Scope VNX_ARRAY_ACCESS -Address
VNX_server_frameIP
```

VNX_ARRAY_ACCESS is set by the Storage Admin based on VNX_array_user. The options are as follows:

- 0 - global
- 1 - local
- 2 - LDAP

The VNX_server_frameIP can be 1 or more, and is the IP where the LUNs are visible to the array.

Creating the Navisphere file manually on Windows systems

To manually create or modify the Navisphere security file on Microsoft Windows systems, complete the following steps.

Procedure

1. On the application host, enable the `naviseccli` pop-up windows by running the Interactive Service Detection feature:
 - a. Select **Start > Run > services.mcs**.
 - b. Start the Interactive Service Detection service.
2. Download `PSEXEC` from Microsoft SysInternals, and then unzip it in a temporary folder.
3. Open a command prompt as an administrator, and then browse to the folder where you unzipped `PSEXEC.EXE`.
4. Type the following command:

```
PSEXEC -i -s -d CMD
A new command prompt appears.
```

5. To verify that the command prompt belongs to the system user account, type the following command:

```
WHOAMI/USER
```

6. To set the global account for all users on this system, type the following command:

```
naviseccli -User username -Password password -Scope 0 -
AddUserSecurity
```

7. To set credentials for a specific VNX storage array, type the following command:

```
naviseccli -h VNX_server -addusersecurity -user VNX_array_user -
password VNX_array_user_password -Scope VNX_ARRAY_ACCESS -
Address VNX_server_frameIP
```

VNX_ARRAY_ACCESS is set by the Storage Admin based on VNX_array_user. The options can be as follows:

- 0 - global
- 1 - local
- 2 - LDAP

The VNX_server_frame IP can be 1 or more, and is the IP where the LUNs are visible to the array.

Configuring Unisphere CLI on VNXe3200

Follow this procedure to run Unisphere CLI on Windows 2012. Unisphere CLI is the only binary required for NSM VNXe-Snap Block snapshot to work. Currently Unisphere CLI is only available as a 32-bit binary.

Note

NetWorker now supports snapshots of both a single and LUN Group devices.

Procedure

1. On the production host and the mount host, install Unisphere CLI.
The Unisphere CLI product documentation provides information for OS specific installation and requirements.
2. In the **Configure Optional Settings** pane, to include Unisphere CLI in the environment path, do not change the **Unisphere CLI in the Environment path** that is selected by default.
3. To avoid runtime backup errors, in the **Configure Optional Settings** pane, select **Low (The certificate will not be verified)**.
4. On both the client and the proxy or storage node, create a security file by typing the following command:

```
uemcli -d <IP address> -u <username> -p <password> -saveUser
```

5. Test the connection by typing the following command:

```
uemcli -d <IP address> /prot/snap show
```

If after running this command the system displays an error, then Unisphere CLI is not setup correctly.

UEMCLI Windows registry setup

The Unisphere CLI might not add the installed path to its registry entry. The full path to the binary is required and you must specify the Windows registry setting for the UEMCLI system path. To ensure you have a valid path for the Unisphere install location, you must manually add a registry entry for NetWorker.

1. From the command prompt or a Windows shell prompt, type `regedit`.
2. Locate the `HKEY_LOCAL_MACHINE\SOFTWARE\Legato\NetWorker` folder.
3. Right-click, and then select **New > String Value**.
4. Name the value `UEMCLI_directory`, and then press **Enter** on the keyboard.
5. Right-click the `UEMCLI_directory`, and then select **Modify**.
6. In the **Value data** field, type the full path to the already installed Unisphere CLI location.

When you provision a file system from VNXe, indicate which host has access to the LUN, and which host has access to the snapshot. For a single node setup, the source or primary host has access to both the LUNs and snapshots. For a multi node setup, the source host has access to LUNs, and the target or secondary host has access to the snapshot.

CHAPTER 9

Configuring snapshots on RecoverPoint

This chapter includes the following topics:

- [Snapshot support of RecoverPoint](#)..... 122
- [Supported RecoverPoint features](#)..... 123
- [RecoverPoint configuration methods](#)..... 124
- [RecoverPoint snapshot retention](#)..... 124

Snapshot support of RecoverPoint

NetWorker snapshot support, practices, and configurations that are specific to RecoverPoint are described.

[Migrating Legacy PowerSnap Configurations](#) on page 167 provides information on migrating PowerSnap RecoverPoint implementations to NetWorker snapshot management.

Note

Recover Point is designed in such a way that it takes a snapshot at CG level whenever CGs are used. NetWorker Snapshot Manager provide users the option to operate at the Individual LUN level inside the CG of Recover Point and take snapshots. It can be recovered by using all the recovery options like file-by-file and Saveset Recovery for both snapshot and Clone backups and Rollback for snapshot backup.

Recover Point Snapshot Restore by Rollback removes snapshots that are created after the snapshot for the LUNs in the same CG, which was used for rollback operation. Due to this behavior, if any LUNs in CG that are used to create separate savesets in NetWorker and are used in separate policies which are backed up at different time slots, if one of them rollback, there is a chance that the snapshot for other saveset is deleted which are taken after the rollback saveset. Due to this limitation, it is recommended that Recover Point snapshots from NSM should be followed by clone backups to avoid accidental deletion of backup copies in case of a rollback operation for a different saveset. For example, CG1 has 3 LUNs, A,B, and C. Each time Recover Point takes a snapshot, it would be at a CG1 level and not at a individual LUN level. NetWorker, provides the option to create multiple clients with subsets of these three LUNs (LUN A and B can be part of client 1 and LUN C can be part of Client 2).

Whenever a NetWorker Snapshot Manager backup is taken for Client 1, Recover Point takes the backup of CG1 and NetWorker would internally use the snapshots of LUNs A and B only. For Client 2 snapshot Recover Point would take snapshot at CG1 and NetWorker would use snapshot of LUN C.

The user has taken backups of Client A and Client B every 6 hrs, one after other starting at 00 hours, with Client A, and 0600 hours with Client B. There will be 2 snapshots for each client after 24 hours. Now at 1800 hours, there would have been 2 snapshot backups of client A and one snapshot backup of client B. If for some reason, client A data is corrupted just after the first snapshot of Client B (at 0600 hours.) and user chooses to recover using the first snapshot of Client A by means of a rollback, all snapshots taken after 0600 hours will be lost. This is because Recover point would treat all these as CG1 snapshots.

To mitigate this potential risk, it is recommended to take a clone for all the backups, so that the data is not lost.

Snapshot operations with RecoverPoint software

RecoverPoint can provide local replication and remote replication of protected application volumes. During production operations, RecoverPoint tracks every write activity on the protected application host's production volumes and records this activity as specific point-in-time (SPIT) bookmarks. By using these SPITs, RecoverPoint can reconstruct any previous state of the volumes within a specified period, enabling any-point-in-time recovery.

NetWorker can recover the snapshots or clone them to conventional storage media. For each snapshot operation, RecoverPoint records a SPIT bookmark of the snapshot. The *RecoverPoint Administration Guide* provides details.

Prerequisite for RecoverPoint configurations

The NetWorker Snapshot Management (NSM) prerequisite for a storage appliance running RecoverPoint is that you must Configure RecoverPoint on a supported VMAX, VNX, XtremIO, or VPLEX storage array.

The *NetWorker E-LAB Navigator* provides details. See [Storage array specific prerequisites](#) on page 65 for more details.

Restrictions for RecoverPoint configurations

The following restrictions apply to NetWorker Snapshot Management (NSM) operations with RecoverPoint:

- NSM backups cannot support consistency groups that are configured with more than one Remote Replication copy. Trying to back up or restore snapshots by using RP_CRR might fail. To work around this issue, use RP_CDP, and use a mount host that has visibility to the local replica LUNs.
- NSM must back up all protected sources in a consistency group in a single session. For example, a consistency group that protects the file systems - G: \ and L: \ must be backed up using NetWorker in one client resource (Save Set: "G: \ L: \"). Otherwise, if backups of G: \ and L: \ begin simultaneously, one backup fails.

Note

Changes that are made to RecoverPoint Consistency Groups outside of NSM operations, including renaming or modifying the contents of protection sets, can result in backup and restore failures of NSM snapshots. If a protection set has changed, start a new NSM snapshot.

Supported RecoverPoint features

The NetWorker Snapshot Management (NSM) integration with RecoverPoint supports the following capabilities:

- Specific point-in-time (SPIT) snapshot capability, implemented as a RecoverPoint bookmark.
- Backup and restore capability through NSM:
 - Snapshot backup (PIT)
 - PIT mount on mount host
 - Clone (backup of PIT)
 - Snapshot restore
 - Snapshot management
 - Rollback
 - Clone restore

Snapshot management policy

NetWorker Snapshot Management (NSM) for RecoverPoint creates RecoverPoint bookmarks and makes the bookmarks available to NetWorker by saving them into the

media database as snapshot save sets (snapsets). The backup administrator uses the NMC Client Configuration Wizard to manage volumes that are protected by RecoverPoint.

Have the following information ready before you configure a NetWorker client with the configuration wizard:

- The RecoverPoint appliance hostname or IP address.
- The username and password for the RecoverPoint appliance and permissions to create bookmarks.
- The mount host that is attached to the RecoverPoint replication storage.
- The backup option for the backup, which is either continuous data protection (CDP), which informs NSM that the mount host has access to the local replication storage, or continuous remote replication (CRR) which informs NSM that the mount host has access to the remote replication storage.

Snapshot backups

NSM for RecoverPoint creates PIT copies of the data during scheduled backups by associating snapshot save sets with RecoverPoint bookmarks.

During the snapshot process, the NSM extracts RecoverPoint appliance credentials from the NetWorker Client resource to discover dependent consistency groups and their copies. With this information, NSM requests a bookmark from the RecoverPoint appliance and saves the bookmark information as part of the snapshot backup.

RecoverPoint configuration methods

The NetWorker Management Console (NMC) Client Configuration Wizard supports the creation of NetWorker Snapshot Management (NSM) RecoverPoint configurations, and is the recommended configuration method.

After you create a RecoverPoint configuration, you can modify the configuration through the NMC property windows. This modification enables you to use Application Information variables, as listed in [Application Information Variables](#).

As an alternative to using the wizard, you can manually create an NSM configuration for RecoverPoint by using the `nrsnapadmin -a -rpcreate` command, as follows:

```
nrsnapadmin -a -rpcreate -s networker_server -app recoverpoint_engine  
-u username -p password
```

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

RecoverPoint snapshot retention

You can identify the lifecycle of a RecoverPoint SPIT only by the RecoverPoint Copy Policy that is configured by the bookmark. When the bookmark reaches the end of its Copy Policy retention period, the SPIT becomes invalid and a NetWorker recovery cannot use the SPIT.

Occasionally, NSM snapshot save sets that correspond to invalid bookmarks can be present in the NetWorker media database. This issue can occur when the daily cleanup process has not yet removed the save set references.

When you plan the management of RecoverPoint snapshot operations with NSM, consider the following potential issues:

- A high storage change rate can force a bookmark out of its Copy Policy retention period before a backup is completed. This issue can cause snapshot clones to fail. For example, RecoverPoint software can track database changes to a production LUN many times each second with an update to its journal volume for each change. When the journal volume reaches its capacity, RecoverPoint automatically discards the oldest journal updates, including bookmarks created for NSM. For a large NSM clone, the time that is needed to back up a SPIT can exceed the duration of the SPIT's record in the journal. If RecoverPoint deleted the bookmark, the backup fails.

A possible solution is to increase the amount of journal space in a consistency group.

- If the NSM snapshot policy retains a snapshot save set longer than the RecoverPoint Copy Policy retention period for its bookmark, the save set remains in the NetWorker media database although it is invalid.
- The restore operation can cause RecoverPoint to remove the oldest bookmarks due to Copy Policy retention periods. Restore from a snapshot or from conventional storage adds updates to the RecoverPoint data changes and can force older bookmarks out of the retention period.
A possible solution is to increase the amount of journal space in a consistency group.

CHAPTER 10

Configuring snapshots in a Cluster Environment


This chapter includes the following topics:

- [NetWorker support of cluster environments](#)..... 128
- [Configuring a cluster environment for snapshots](#)..... 128
- [AIX systems in a cluster environment](#)..... 129
- [ProtectPoint restore and rollback for VCS on Solaris](#)..... 130

NetWorker support of cluster environments

In a cluster configuration, NetWorker creates snapshots from file systems on a virtual node. The cluster's virtual node and physical nodes must be NetWorker clients. The following figure shows an example of the data flow for a snapshot and clone in a cluster environment.

Figure 8 Snapshot and clone in a cluster environment



Failover with snapshots in a cluster environment

If a failover occurs during a snapshot and a necessary resource becomes unavailable, NetWorker stops the snapshot and cleans up to return the snapshot environment to the pre-snapshot state. After the failover is completed, NetWorker retries the snapshot on the active cluster node.

If the application or the cluster node fails over during a recovery, NetWorker stops the operation and does not automatically retry the recovery. If you retry the recovery manually, NetWorker recovers the data to the current active node in the cluster.

Configuring a cluster environment for snapshots

The recommended cluster configuration includes the NetWorker server and the storage node on separate hosts outside of the cluster.

Procedure

1. In a local directory, install NetWorker client software on each physical node of the cluster.

The *NetWorker Cluster Integration Guide* provides details.

2. Although the snapshots are created on the virtual node, each physical node in the cluster must be a cluster-aware NetWorker client. On each physical node run the cluster configuration script:

- Microsoft Windows:

```
NetWorker_install_path\lc_config
```

- UNIX:

```
/NetWorker_install_path/networker.cluster
```

The *NetWorker Cluster Integration Guide* provides details.

3. For each virtual node in the cluster that requires NetWorker snapshot services, configure a NetWorker Client resource. Include the following settings:

- a. In the **Application Information** attribute, specify the shared directory path in the NSR_PS_SHARED_DIR variable.

NetWorker creates the `ssres` file in this directory. The `ssres` file contains transaction logs that NetWorker uses to clean up aborted snapshots and return the snapshot environment to the pre-snapshot state.

This shared directory can be at one of the following locations:

- Storage that the application resource group manages.
- A global file system that is accessible to all the cluster nodes.
[Common Application Information variables](#) provides details.

- b. In the **Remote Access** attribute, specify the system (Windows) or root (UNIX) account and the hostname or cluster IP of each physical node within the cluster. For example:

- Microsoft Windows:

```
system@clus_phys1
system@clus_phys2
```

- UNIX:

```
root@clus_phys1
root@clus_phys2
```

The *NetWorker Administration Guide* provides details on NetWorker in a cluster environment.

AIX systems in a cluster environment

NSM does not support unmanaged file system devices for cluster environments on AIX platforms.

Object Data Manager is not cluster-aware on AIX systems. When you add or remove logical volumes, update each ODM database by using one of the following methods:

- Export and import all modified volume groups on all other nodes.
- Use the `synclvodm` command as the root user to synchronize the device configuration database with the LVM information:

```
synclvodm -v VGName
```

where *VGName* is the name of the volume group to synchronize. The AIX System Management Guide provides details.

ProtectPoint restore and rollback for VCS on Solaris

You can enable a ProtectPoint restore and rollback for VCS on Solaris.

Performing a ProtectPoint VCS restore

Procedure

1. As the root user, on the primary VCS node, perform the following steps:
 - a. List the VCS Service Groups.

```
root:/# hastatus -sum
```

```
-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING        0
A ledma056         RUNNING        0

-- GROUP STATE
-- Group           System          Probed    AutoDisabled    State
B ClusterService  ledma054        Y         N                ONLINE
B ClusterService  ledma056        Y         N                OFFLINE
B oracle_ctl_sg   ledma054        Y         N                ONLINE
B oracle_ctl_sg   ledma056        Y         N                OFFLINE
B oracle_sg       ledma054        Y         N                ONLINE
B oracle_sg       ledma056        Y         N                OFFLINE
B vxfen           ledma054        Y         N                ONLINE
B vxfen           ledma056        Y         N                ONLINE
```

- b. To enable the VCS configuration as Read/Write, type the following command:

```
root:/# haconf -makerw
```

- c. To freeze the VCS service groups by disabling On line/Off line, type the following command:

Note

The following example is for a VCS and Oracle configuration.

```
root:/# hagrps -freeze <oracle_sg> -persistent
```

```
root:/# hagrps -freeze <oracle_ctl_sg> -persistent
```

- d. To confirm the VCS status, type the following command:

```
root:/# hastatus -sum
```

```
-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING        0
A ledma056         RUNNING        0
```

```

-- GROUP STATE
-- Group          System          Probed      AutoDisabled  State
B ClusterService ledma054      Y           N             ONLINE
B ClusterService ledma056      Y           N             OFFLINE
B oracle_ctl_sg   ledma054      Y           N             ONLINE
B oracle_ctl_sg   ledma056      Y           N             OFFLINE
B oracle_sg       ledma054      Y           N             ONLINE
B oracle_sg       ledma056      Y           N             OFFLINE
B vxfen           ledma054      Y           N             ONLINE
B vxfen           ledma056      Y           N             ONLINE

-- GROUPS FROZEN
-- Group
C oracle_ctl_sg
C oracle_sg

-- RESOURCES DISABLED
-- Group          Type          Resource
H oracle_ctl_sg   DiskGroup     oracle_ctl_dg_DG_res1
H oracle_ctl_sg   Mount         oracle_ctl_dg_MNT_res1
H oracle_ctl_sg   Volume        oracle_ctl_dg_VOL_res1
H oracle_sg       DiskGroup     oracle_dg_DG_res1
H oracle_sg       Mount         oracle_dg_MNT_res1
H oracle_sg       Volume        oracle_dg_VOL_res1

```

e. To make the VCS configuration as Read Only, type the following command:

```
root:/# haconf -dump -makero:
```

2. As the Oracle user, on the primary VCS node, perform the following steps:

a. On the Oracle database, run the shutdown and startup mount commands.

- a. oracle:/# sqlplus / as sysdba
- b. SQL > shutdown immediate
- c. SQL > startup mount
- d. SQL > exit

b. Perform the RMAN restore and recovery.

3. As the root user, on the primary VCS node, perform the following steps:

a. Make the VCS configuration Read/Write. Type the following command:

```
root:/# haconf -makerw
```

b. Unfreeze the service groups, and allow On line/off Line. Type the following command:

```
root:/# hagrps -unfreeze <oracle_ctl_sg> -persistent
```

```
root:/# hagrps -unfreeze <oracle_sg> -persistent
```

c. Confirm the VCS status. Type the following command:

```
root:/# hastatus -sum
```

```

-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING       0
A ledma056         RUNNING       0

-- GROUP STATE

```

```

-- Group          System          Probed    AutoDisabled    State
B ClusterService ledma054    Y          N                ONLINE
B ClusterService ledma056    Y          N                OFFLINE
B oracle_ctl_sg  ledma054    Y          N                ONLINE
B oracle_ctl_sg  ledma056    Y          N                OFFLINE
B oracle_sg      ledma054    Y          N                ONLINE
B oracle_sg      ledma056    Y          N                OFFLINE
B vxfen          ledma054    Y          N                ONLINE
B vxfen          ledma056    Y          N                ONLINE

```

Performing a ProtectPoint VCS rollback

Note

A rollback fails if you change the style of the mpio device name. The rollback to the source LUN is successful. However, the fsck and mount fails. In this scenario, manually mount the FS.

Procedure

1. As the root user, on the primary VCS node, perform the following steps:
 - a. List the VCS Service Groups. Type the following command:

```
root:/# hastatus -sum
```

```

-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING       0
A ledma056         RUNNING       0

-- GROUP STATE
-- Group          System          Probed    AutoDisabled    State
B ClusterService ledma054    Y          N                ONLINE
B ClusterService ledma056    Y          N                OFFLINE
B oracle_ctl_sg  ledma054    Y          N                ONLINE
B oracle_ctl_sg  ledma056    Y          N                OFFLINE
B oracle_sg      ledma054    Y          N                ONLINE
B oracle_sg      ledma056    Y          N                OFFLINE
B vxfen          ledma054    Y          N                ONLINE
B vxfen          ledma056    Y          N                ONLINE

```

- b. To enable the VCS configuration as Read/Write, type the following command:

```
root:/# haconf -makerw
```

- c. To freeze the VCS service groups by disabling On line/Off line, type the following command:

Note

The following example is for a VCS and Oracle configuration.

```
root:/# hagrps -freeze <oracle_sg> -persistent
```

```
root:/# hagrps -freeze <oracle_ctl_sg> -persistent
```

- d. To confirm the VCS status, type the following command:

```
root:/# hastatus -sum
```

```
-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING       0
A ledma056         RUNNING       0

-- GROUP STATE
-- Group           System          Probed    AutoDisabled  State
B ClusterService ledma054         Y         N              ONLINE
B ClusterService ledma056         Y         N              OFFLINE
B oracle_ctl_sg   ledma054         Y         N              ONLINE
B oracle_ctl_sg   ledma056         Y         N              OFFLINE
B oracle_sg       ledma054         Y         N              ONLINE
B oracle_sg       ledma056         Y         N              OFFLINE
B vxfen           ledma054         Y         N              ONLINE
B vxfen           ledma056         Y         N              ONLINE

-- GROUPS FROZEN
-- Group
C oracle_ctl_sg
C oracle_sg

-- RESOURCES DISABLED
-- Group           Type          Resource
H oracle_ctl_sg   DiskGroup    oracle_ctl_dg_DG_res1
H oracle_ctl_sg   Mount        oracle_ctl_dg_MNT_res1
H oracle_ctl_sg   Volume       oracle_ctl_dg_VOL_res1
H oracle_sg       DiskGroup    oracle_dg_DG_res1
H oracle_sg       Mount        oracle_dg_MNT_res1
H oracle_sg       Volume       oracle_dg_VOL_res1
```

e. Make the VCS configuration as Read Only. Type the following command:

```
root:/# haconf -dump -makero:
```

2. As an Oracle user, on the primary VCS Node, perform the following steps:

a. On the Oracle database, run the `shutdown` and `startup mount` commands:

- a. `oracle:/# sqlplus / as sysdba`
- b. `SQL > shutdown immediate`
- c. `SQL > startup mount`
- d. `SQL > exit`

b. Perform the RMAN rollback and recovery.

3. As the root user, on the primary VCS node, perform the following steps:

a. To make the VCS configuration Read/Write, type the following command:

```
root:/# haconf -makerw
```

b. To unfreeze the service groups, and allow On line and or off Line, type the following command:

```
root:/# hagr -unfreeze <oracle_ctl_sg> -persistent
```

```
root:/# hagr -unfreeze <oracle_sg> -persistent
```

c. To confirm the VCS status, type the following command:

```
root:/# hastatus -sum
```

```
-- SYSTEM STATE
-- System          State          Frozen
A ledma054         RUNNING       0
A ledma056         RUNNING       0

-- GROUP STATE
-- Group           System          Probed    AutoDisabled  State
B ClusterService ledma054         Y         N              ONLINE
B ClusterService ledma056         Y         N              OFFLINE
B oracle_ctl_sg   ledma054         Y         N              ONLINE
B oracle_ctl_sg   ledma056         Y         N              OFFLINE
B oracle_sg       ledma054         Y         N              ONLINE
B oracle_sg       ledma056         Y         N              OFFLINE
B vxfen           ledma054         Y         N              ONLINE
B vxfen           ledma056         Y         N              ONLINE
```

Note

The service groups will be faulted. But they will come back online in a short time.

CHAPTER 11

Data Management and Recovery

This chapter includes the following topics:

- [Snapshot lifecycle management](#)..... 136
- [Management and recovery of file system snapshot data](#)..... 136
- [Snapshot recovery support and limitations](#)..... 137
- [Restoring from a snapshot with the Recovery Wizard](#)..... 138
- [Restoring a snapshot by rollback](#)..... 141

Snapshot lifecycle management

Snapshot-based operations are fully integrated into NetWorker policies.

The schedule of the workflow controls the number of snapshots taken each day.

A backup action has a minimum retention time which defines when a snapshot can be deleted on the storage array and NetWorker media database to take a new snapshot. This deletion is done as part of the new backup operation. The retention time that is specified in the snapshot backup action controls the snapshot expiration.

An example of how the previous *rollover-only* (“serverless”) policy can be created in the new policy framework, is to define a workflow that:

- Contains a snapshot backup action.
- Followed by a clone action with **Delete source save sets after clone operation** selected.

Management and recovery of file system snapshot data

Only management and recovery operations for file system data are described. For information on database snapshots, such as NetWorker Module for Databases and Applications (NMDA) or NetWorker Module for SAP (NMSAP) database snapshots, refer to the documentation for the NetWorker application module that you are using.

The NetWorker Management Console (NMC) Recovery Wizard and the NetWorker CLI commands provide features that enable you to browse, delete, change snapshot expiration, and recover snapshot data.

Save set IDs and expiration policies

When NSM creates a snapshot, NSM generates a separate save set ID for each snapshot object specified in the Client resource. For example, a single physical snapshot can create save sets for `F:\abc` and `G:\xyz` if they both reside on the same LUN or managed volume. Each save set will have a separate save set ID, even if both save sets belong to the same client and both reside on the same LUN. `F:\abc` and `G:\xyz` cannot be on the same LUN.

During a clone operation to conventional storage media, NetWorker assigns a different clone ID to each cloned snapshot object.

By having two save set IDs, NSM manages the snapshot data separately from the cloned data. Each save set has an independent expiration policy, and when one save set expires, you can still use the other save set to perform a restore.

Browsing snapshot and clone save sets

NetWorker records in the client file index only the files that NSM clones to conventional storage media. Because NetWorker indexes clones, you can browse the files in NMC.

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

Snapshot recovery support and limitations

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

- A snapshot recovery is supported in the following user interfaces:
 - NMC Recovery Wizard
 - `nsrsnapadmin` command utility
 - `nsrsnap_recover` command
- You can restore individual files or complete file systems from snapshot save sets.
- You cannot combine individual files from multiple save sets in a single restore session.
- You can restore data from snapshots that are cloned to conventional storage media by using the NMC Recovery Wizard or other NetWorker methods, as you would for any conventional NetWorker backup. The *NetWorker Administration Guide* provides details.

Note

Recovery from a snapshot that was being rolled over to a secondary media (Disk or Tape), is not possible. Both the rollover and the recovery fail.

Note

From NetWorker 9.2 onwards, for VMAX3 Snapvx/ProtectPoint VMAX3 file system recovery, if the recover LUNS are not part of VMAX default storage group `NsrSnapSG`, file by file recovery is not possible through recovery wizard. The same can be done through NetWorker CLI commands, `nsrsnap_recover`, and `nsrsnapadmin` because the user defined VMAX3 storage group value has to be passed with parameter `NSM_SNAP_SG`. For file by file recovery, currently in recovery wizard, there is no means to provide this parameter before starting the recovery.

Raw partitions and raw devices

The following considerations apply to NSM restores of raw partitions and raw devices:

- NSM does not support mount points for raw file system backups.
- You can recover raw partitions from Microsoft Windows application hosts only to the same drive letter from which NSM backed up the raw partitions. You cannot redirect the recovery to another drive letter.
- Before you perform a snapshot restore of a file system that NSM backed up as a raw device, you must unmount the source file system. After the restore completes, you must run the `fsck` command before you mount the file system.

NetApp restore fails

Snapshot backups with NetWorker fail when you try to restore to a NetApp filer file system directly. Specifically, if the file is a directory or a symbolic link, or contains NT streams, the restore fails for a single file. The NetApp Support site provides additional details: <https://library.netapp.com/ecmdocs/ECMP1196991/html>

Restoring from a snapshot with the Recovery Wizard

You can use the NetWorker Management Console (NMC) Recovery Wizard to restore file system data from a snapshot that is stored on a supported storage array.

Procedure

1. Run NMC, and in the **Enterprise** view, select the NetWorker server name, and then select **Enterprise > Launch Application**.
2. In the NetWorker server's **Recover** view, select **Recover > New Recover**.

The Recovery Wizard appears.

3. On the **Select the Recovery Type** page, select the **Recovery Type**, and then click **Next**.

Complete the **Select the Recovery Hosts** page:

- a. In the **Source Host** field, specify the application host for the production data that was the source for the snapshot you want to restore.
- b. In the **Destination Host** field, specify the application host or an alternative NetWorker client on which you want to restore the snapshot data.

NetWorker now allows you to perform an array level restore (rollback) of a snapshot to an alternate set of devices. Pre-select the device and the devices should be of same size or larger than the original source devices. The devices should be visible to an alternate host. Create the same File systems that were on original source devices. The file systems should be mounted at time of rollback on the alternate host. NSM will unmount the file systems and perform the rollback to these devices and mount them back after the rollback.

On the NMC, when the destination host is not source host, the system automatically rolls back the Smart snapshot to any LUN in the "NsrSnapSourceSG" storage group.

Note

NetWorker uses the following criteria to match devices on the target host with the source devices for VMAXv3 SnapVX rollback to alternate host:

- The device should be of equal or greater size.
 - The device should have no clone or BCV relationships.
 - The total no of source LUNs should match the total number of target LUNs.
-

- c. In the **Available Recovery Types** table, you can either select **Filesystem**, **SmartSnap**, or another supported NetWorker application type that is installed on the client.

The **SmartSnap** option allows you to specify array LUNs World Wide Names (WWNs).

- d. Click **Next**.

4. Complete the **Select a Snapshot** page. You can restore the entire snapshot or you can select the individual directories and individual files from the snapshot:

- a. The **Snapshots** table lists the snapshots on the storage array that are available to the source client. Select the snapshot to restore from, based on the snapshot time and save set volumes.
 - b. Select one of the following types of restores to perform:
 - **Recover save sets**
 - **Rollback snapshot**
-

Note

For SmartSnap, only the Rollback snapshot restore type is supported. If the LUN is a part of a mounted file system or active volume group, you should manually unmount and export it before recovery.

- c. If you selected the **Recover save sets** option, specify the following settings:
 - In the **Select save set** field, select a single save set volume to mount and restore from. The next wizard page will let you browse the directories and files in the mounted save set.
-

Note

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

- In the **Mount save set on** field, select the host on which to mount the save set, ready for the restore operation. The mount host can be one of the following:
 - Destination client that you selected earlier in the wizard.
 - An existing storage node that you can select from the drop-down list.
-

Note

If you use a storage node as the mount host, ensure that the storage node has access to the storage array. For example, you can specify the storage node as the mount host in the properties of the source client, on the APPS and MODULES tab, in the **NSR_DATA_MOVER Application Information** attribute.

Note

Veritas Volume Manager does not support the configuration of production file systems and snapshot file systems that are mounted on the same host. The mount host cannot be the application host.

[Figure 3](#) on page 29 shows an example of the data flow in a restore operation where the NetWorker storage node is the snapshot mount host.

- In the **Recover mode** field, select whether you want to recover individual items from the save set or the full save set:
 - If you select **Browse and recover save set**, when you click **Next**, the wizard mounts the snapshot volume for the save set and displays the **Select Data to Recover** page. The mount operation can take some time.

- If you select **Recover full save set**, when you click **Next**, the wizard displays the **Select the Recovery Options** page.
 - d. If you selected **Rollback snapshot**, a warning is displayed: A `rollback` is a destructive operation. When you click **Next**, the wizard displays the **Perform the Recovery** page.

A rollback restores the entire snapshot to the destination client that you selected in the **Snapshots** table.

[Restoring a snapshot by rollback](#) on page 141 provides details.
 - e. Click **Next**. The result depends on the recover or rollback selections.
 - f. You can select **Advanced options**, and then specify **Debug level** and other recover options.
5. If you selected the **Browse and recover save set** option, complete the **Select Data to Recover** page, otherwise, skip this step.
- a. Specify the location of the items to restore by using the browse tree or typing the full path of the location. Indicate the directories or files for NSM to restore by marking them in the table.

Note

The wizard does not list expired save sets. You can restore existing expired save sets manually by using the `nsrsnapadmin` command utility with the `R` command option or the `nsrsnap_recover` command. [Using nsrsnapadmin for NSM operations](#) and the *NetWorker Command Reference Guide* provide details.

- b. Click **Next**.
6. If you selected any of the **Recover save sets** options, complete the **Select the Recovery Options** page. If you selected the rollback option, skip this step.
- a. In the **File Path for Recovery** field, select, browse, or type a location where NSM restores the files:
 - **Original path**
 - **New destination path**

Note

You cannot repeat the same restore operation to the same destination.

- b. In the **Duplicate file options** field, specify how NSM resolves file name conflicts:
 - **Rename the restored file**—NSM restores the file with a new name that NSM automatically generates.
 - **Do not recover the file**—NSM will not restore the file.
 - **Overwrite the existing file**—NSM replaces the file with the same name.
- c. To specify further options, select **Advanced Options**, and then specify the attributes.

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

Results

The wizard restores the files:

- For a save set restore or file level restore, the data restore path is over the LAN as shown in [Figure 3](#) on page 29.
- For a rollback recovery, the storage array's capabilities perform the restore. [Restoring a snapshot by rollback](#) on page 141 provides details.

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

Restoring a snapshot by rollback

NSM uses the native capabilities of the storage array to perform rollbacks. A rollback is a restore in which a volume on the application host is unmounted and the storage array replaces the entire contents of the unmounted volume with the entire contents of the snapshot volume.

Rollback considerations

Always consider how a rollback can affect any other snapshots or other data on the storage array.

Consider the following limitations and precautions before you perform a rollback:

- NSM supports rollback operations in a clustered environment when the cluster software is disabled on the volumes participating in the operation.
- The file system that you roll back must be the only file system on the application volume.
- The volume must be the only volume in the volume group.
- The file system occupies the entire volume space and no other objects are on the same volume.

- If a rollback fails, the application host's file system may remain unmounted, and you must manually mount the file system.
- VNXe Block array does not support the rollback restore feature.
- XtremIO does not support the rollback restore feature.

NOTICE

Rollbacks overwrite the entire contents of the source LUNs and potentially destroy the existing data.

Note

On Linux and Solaris, if a disk is configured with partitions, you can perform a rollback restore only if you list the entire disk in the `psrollback.res` file. The rollback restore then overwrites the entire disk.

For example, if `/fs1` and `/fs2` are configured with partitions `/dev/sdc1` and `/dev/sdc2` respectively, then enable the rollback restore of `/fs1` by listing the entire disk `/dev/sdc` in the `psrollback.res` file. The rollback restore overwrites the entire disk `/dev/sdc`, so `/fs2` is also restored.

Configurations that override rollback safety checks

By default, NSM performs safety checks to ensure that there are no datasets on the rollback target LUN to the same or alternate locations, other than those for which NSM has snapshots.

Either of the following conditions can override the safety checks:

- The `psrollback.res` file includes the rollback target.
 - You use the force option `-f` with the `nsrsnapadmin` or `nsrsnap_recover` command.
-

Note

The `-f` option is not supported for Database modules.

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

Alternate LUN Rollback

The following considerations apply for rollbacks to an alternate LUN.

File systems

The file system, database, or tablespace name should be the same as the source saveset.

Any additional FS/files/directories must be added to the `psrollback.res` file.

LUNs

The number of alternate host LUNs should be the same as the number of LUNs at the time of backup.

Size:

- For SnapVX, the size of LUNs should be equal to or greater than the size of the backup LUNs.

- For ProtectPoint, the size of the LUNs should be equal to or greater than the size of the static image.

Volume Groups

The volume group name can be different from the source saveset.

Can be any size.

Logical volumes

The logical volume name can be different from the source saveset.

The number of logical volumes on the alternate host can be equal or less than the source saveset.

If the logical volumes on the alternate host are greater than logical volumes in the source saveset, then you should add them to the `psrollback.res` file.

Can be any size or layout.

Host

Alternate LUN Rollback is supported only on a destination host other than the source host.

Array

Alternate LUN Rollback is supported only on Snapshots that are taken on SnapVX and ProtectPoint SnapVX.

SmartSnaps

While Rolling back a Smart snapshot to a destination host, the rollback is done to any destination LUN which was specified in the "NsrSnapSourceSG" Storage Group. By default, the system selects LUNs visible to the destination host. To select any LUN, visible or not visible, specify `SELECT_HOST_VISIBLE_TGTS=false`.

Additional affected file systems that are not a part of the rollback will not be mounted or unmounted.

SmartSnap Alternate Rollback

For a SmartSnap rollback to an alternate host, the LUN's that were selected from NsrSnapSourceSG for a rollback, can be found in the logs. The following sample message shows the LUN's selected for a SmartSnap rollback.

- `nwsnap.raw` messages specific to alternate host rollback.

```
- 0 10/28/2016 07:39:17 AM 5 0 0 3568428832 20739 1477654745
ledme079.lss.emc.com nsrpsd NSR critical LOG [msg #224
vmaxv3_snapvx_snapshot.cpp 973 PSDBG 0]
```

Alternate LUN Rollback device selections

- ```
- 159149 10/28/2016 07:39:17 AM 0 0 5 3568428832 20739
1477654745 ledme079.lss.emc.com nsrpsd NSR info
```

Selecting target device 000196701031:00DC0 for source device 000196701031:00DB9

The SmartSnap restore to alternate feature supports only SnapVX and ProtectPoint with VMAX.

## Example of a destructive rollback

Three file systems, `/fs1`, `/fs2`, and `/fs3`, exist on a LUN, which resides on a storage array standard device. You create a snapshot for the `/fs1` file system. Because `/fs2` and `/fs3` also reside on the LUN, the snapshot includes those file systems. After the snapshot, you create a fourth file system, `/fs4`, on the LUN.

If you perform a rollback of `/fs1`, the snapshot will overwrite the contents of the entire LUN. The rollback will revert the contents of `/fs1`, `/fs2`, and `/fs3`, and it will destroy the new `/fs4` file system. Although NSM safety checks do not normally allow a rollback overwrite such as this, exceptions can occur. The exceptions occur when you roll back with the force option or when `/fs2`, `/fs3`, and `/fs4` are present in the `psrollback.res` file. Either exception will destroy `/fs4` and roll back `/fs2` and `/fs3`.

## Configuring the `psrollback.res` file

The `psrollback.res` file is a resource file for NetWorker snapshot rollback with the following pathname:

- On UNIX

```
/nsr/res/psrollback.res
```

- On Microsoft Windows

```
C:\Program Files\EMC NetWorker\nsr\res
```

Before NetWorker performs a rollback it performs safety checks to verify that the operations will not overwrite any file, directory, partition, or volume that is outside of the save set. NetWorker uses the `psrollback.res` file to provide NetWorker with configuration information for the rollback.

This resource file contains the files, directories, partitions, and volumes to exclude from the rollback safety check. The rollback can overwrite the items that you list in this file.

The resource file includes the following features:

- You can add more files or directories to this file by using the following syntax rules:
  - At least one line per file or directory
  - Pathnames starting with a forward slash (/) are absolute pathnames, for example, `/tmp`
- The file supports the following items:
  - Directory or file pathname
  - File system
  - Block device of a managed or unmanaged raw device, for example, `/dev/vg_01/vol1`
  - The file does not support character devices



**Note**

When you perform a rollback of a partitioned disk on Solaris, the safety check considers all defined partitions. To avoid rollback failure, list unused partitions in the `psrollback.res` file.

**Examples**

You create the following valid entries in the `psrollback.res` file before you perform a rollback of `/fs1/dir` (UNIX) or `C:\fs1` (Microsoft Windows):

- On UNIX:

```
/fs1/dir1
/fs1/dir2/file1
/fs2
```

- On Microsoft Windows:

```
D:\dir1
C:\dir2\file1
C:\fs2
```

**Rollbacks with Veritas Volume Manager**

For rollbacks of Veritas Volume Manager (VxVM) file systems, the application host mounts every file system that is part of the volume group, including file systems not previously mounted.

**Rollbacks with IBM AIX Volume Manager**

NetWorker Snapshot Management (NSM) supports rollbacks with AIX volume manager as follows:

- If you set the Auto On setting to No, NSM supports rollbacks in an HACMP shared volume group environment. This setting prevents AIX from automatically activating the volume group during a system startup.
- NSM does not support rollbacks in an HACMP concurrent volume group environment. Although a rollback can appear to be successful, the concurrent-capable volume group changes into a nonconcurrent volume group.
- NSM does not support rollbacks of file systems with inline logs.

**Configuring the Auto On setting for an HACMP shared volume group**

After a rollback on AIX systems, by default AIX places HACMP shared volume group configurations into a nonsynchronized state.

You can enable rollbacks that maintain a synchronized state.

**Procedure**

1. On the host where the cluster service is online, take the volume group offline by typing the following command:

```
varyoffvg vg_name
```

2. On each HACMP node within the volume group that is offline, perform the following operations:

- a. Export the shared volume group.
- b. Import the shared volume group.
- c. Use the `chvg` command to set the Auto On setting to No with the `-a n` option:

```
chvg -a n -Q y vg_name
```

3. On the host where the cluster service is online, take the volume group online by typing the following command:

```
varyonvg vg_name
```

4. Test for a successful cluster failover by moving the HACMP resource group between hosts.

# CHAPTER 12

## Troubleshooting

You can use the sections in this chapter to identify and resolve issues with NSM configuration and operation. This chapter includes the following topics:

- [NetWorker snapshot backup issues](#) ..... 148
- [NetWorker snapshot restore issues](#) ..... 150

## NetWorker snapshot backup issues

### Snapshot backup on Unity fails

When you perform a snapshot backup with Unity 4.2 or 4.1, the backup fails displaying the following error:

```
"Internal Error: Failed to get status of import operation."
```

#### Workaround

When you configure Unity 4.2 or 4.1 LUN's, you must add snapshot access for the new LUN through UEMCLI. Use the format below:

```
uemcli -d <unity array ip> -u <user name> -p <password> /stor/prov/
luns/lun -id <lun id> set -snapHosts <host id>
```

where *Lun ID* and *Host ID* can be retrieved from the Block and Hosts pages in Unisphere.

### NAS Isilon snapshot mount fails on Linux

For NAS, Isilon snapshot mount fails on Linux if the NAS share does not have permissions for OTHERS.

#### Workaround

Change the configuration on Isilon for NFS share and add the host (where snapshot is mounted/read) to **Root Clients** list.

### Backup on Windows fails with a Delayed Write Failed error

Due to a Microsoft Windows operating system limitation, a snapshot or a rollover can fail with the following message:

```
Delayed Write Failed
```

#### Workaround

Disable the disk caching feature and perform the backup again.

### Backup fails and hangs when NMC user has insufficient privileges

For a snapshot backup, if any of the following NMC user privileges are missing, the backup fails and hangs.

- Configure NetWorker
- Operate NetWorker
- Monitor NetWorker
- Operate Devices and Jukeboxes
- Recover Local Data
- Backup Local Data
- Backup Remote Data

Simultaneously, NMC displays the following error message:

```
nsrd NSR info Snapshot Management Alert: Backup of [/db_logs/newlogs/
NODE0000/LOGSTREAM0000] failed: Backup failed. User does not have
sufficient privileges on the NetWorker server db2-rh71-dpf3
```

#### Workaround

Stop the action manually and retry after providing sufficient privileges to the NMC user.

## Snapshots fail to mount for AIX managed file systems

For AIX managed JFS2 file systems that use inline logs, snapshots can fail to mount on a remote mount host.

#### Workaround

Use the application host as the mount host or use external logs.

## Snapshots fail for Linux Volume Manager on VNX with PowerPath

The use of PowerPath® software is optional for NSM on VNX arrays. An incorrectly configured Linux Volume Manager (LVM) used with PowerPath can result in snapshot failures with the following error:

```
"/dev/sdbd" is not a device that the CLARiiON SCM recognizes as
snappable
```

#### Workaround

Modify the `lvm.conf` file to be able to use NSM.

The *PowerPath for Linux Installation and Administration Guide* provides details.

## Linux Logical Volume Manager snapshots fail with an error

Linux Logical Volume Manager (LVM) snapshots fail with an error as follows:

```
Failed to get status of import operation. Could not run lvm binary
'lvm'
```

#### Workaround

Create a soft link `ln -s /sbin/lvm /bin/lvm` on the proxy or storage node, and then run the policy.

## NetWorker to Media-Clone stops responding and the backup fails

NetWorker to Media-Clone stops responding because scan command `lvmdiskscan` on the host a system stops responding.

#### Workaround

For the backup to succeed, kill the processes, fix the system issue, and then run the backup.

## NetWorker snapshot restore issues

### File-by-file or saveset restore fails

A File-by-file or saveset restore fails when a wrong storage node or a XtremIO initiator name is used during a restore operation.

---

#### Note

XtremIO always mounts the snapshot on the given initiator name irrespective of the storage node.

The first time the restore fails the following error message is displayed:

```
6211 1481325281 0 0 2 2580 2364 0 ledmb071.lss.emc.com nsrsnapagent
NSR info 2 %s 1 0 29 Unable to mount the snapshot.
```

Subsequently, the following error displays in the storage node log:

```
"decrypted data cached: offset 3340 1479626199 0 0 0 4648 6084
1479626194 ledme040 nsrsnapagent NSR
info 2 %s 1 0 202 [msg #233 D:/views/nw/18.1/nsr/storage/ssm/emc_xio/
xioCommunication.cpp 210
PSDBG 5] Command failed with the following error message: {"message":
"vol_already_mapped_by_ig_tg", "error_code":
400}"
```

To fix this issue, manually unmount the snapshot from the host using the XtremIO management software or commands.

### Restore of raw devices fails on Linux with permission issue

If NSM backs up a raw device as a snapshot or clone, and then restores the device, the ownership of the device pathname changes to root. This change prevents nonroot users from using this device pathname.

#### Workaround

Log in as root, and then type the `chown` command to change the owner of the device pathname to the correct user.

### Command `nsrsnap_recover -I` runs but fails to restore a file

When the `nsrsnap_recover -I` command is used with an improper pathname a file is not restored and the following message appears:

```
Completed the restore of invalid-path
```

The NetWorker Console indicates a restored file.

#### Workaround

Run the command with the proper pathname.

### Restore fails with disk signature error

A restore fails with a disk signature error after one restore from primary or secondary Data Domain is already successful.

**Workaround**

Restart the host.

**Workaround to eliminate a restart of the host**

1. Go to device manager, and then uninstall all Symmetrix disks (not just the affected ones).
2. To rescan all devices, use UI Disk Management .
3. Recycle the Virtual Disk service.

**Note**

If the above workaround does not work in a Windows environment, then make the devices that is having issues visible to a Linux host, create a file system on it, and then use it on a Windows host.

**Directed restore files and folder permission issue**

For Isilon, directed restore files and folder have permission issues on a Windows platform.

**Workaround**

1. Create a Common Internet File System (CIFS) access mount on the Isilon device, and then allow the share to have Windows Access Control Lists (ACLs).
2. Mount the CIFS share to a Windows computer.
3. On the mounted share, save the file.
4. Back up the network-attached storage (NAS) device with the CIFS share.
5. Perform a directed recover of the file on the share to a local disk on the Windows computer.

The files and folders are then recovered without any permission issues.

**Snapshot mount might fail because VMAX does not release the lock on Restore FTS LUN**

VMAX sometimes takes time to release the lock on Restore FTS LUNs. If a subsequent restore is done immediately after a restore, the mount operation might fail.

**Workaround**

The customer must wait till the Restore FTS devices lock is released. They can also manually clear the lock on the Restore FTS devices and then try a restore.

**NSM with XtremIO leaves snapshots mounted**

If you use the wrong initiator name by mistake, NetWorker Snapshot Management (NSM) fails with a message indicating an import error or unable to mount the snapshot, and leaves the snapshot mapped and exposed to the hosts configured with the initiator.

**Workaround**

You must manually unmap the snapshot from the initiator using the XtremIO management software GUI or command.





# APPENDIX A

## Application Information Variables

This appendix includes the following topics:

- [Using Application Information variables](#)..... 154
- [Common Application Information variables](#)..... 154
- [Application Information variables for VMAX arrays](#)..... 155
- [Application Information variables for VNX Block arrays](#)..... 158
- [Application Information variables for RecoverPoint appliances](#)..... 159
- [Application Information variables for XtremIO arrays](#)..... 159

## Using Application Information variables

As part of the manual configuration of an application host, some NSM configurations require the use of special variables that provide specific control of snapshot processes.

To implement these controls, the Client resource for the application host, on the Apps and Modules tab, in the Application Information attribute, you can type the variables and their values.

[Configuring the Client resource manually for the application host](#) provides the manual configuration procedure that can include Application Information variables.

## Common Application Information variables

The following table lists Application Information variables that are common to the storage arrays supported for NSM.

**Table 11** Common Application Information variables

| Common variable          | Definition                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NSR_DATA_MOVER           | Specifies the hostname of the snapshot mount host client. The default value is the hostname of the local application host.                                                                                                                                                                                                                                                                                              |
| NSR_POST_SNAPSHOT_SCRIPT | Specifies the pathname of the postprocessing command script. No default value.                                                                                                                                                                                                                                                                                                                                          |
| NSR_PRE_SNAPSHOT_SCRIPT  | Specifies the pathname of the preprocessing command script. No default value.                                                                                                                                                                                                                                                                                                                                           |
| NSR_PS_DEBUG_LEVEL       | Specifies the verbosity level of the logs. Valid values are 0 to 9. The default value is 3.                                                                                                                                                                                                                                                                                                                             |
| NSR_PS_DO_PIT_VALIDATION | Specifies whether NSM validates that it can mount the completed snapshot on the mount host. Valid values are TRUE and FALSE. The default value is FALSE for ProtectPoint, and TRUE for snapshot backups.<br><br>Set to FALSE to prevent the time and expense of the validation. If NSM cannot mount the snapshot, it cannot restore the data.                                                                           |
| NSR_PS_SAVE_PARALLELISM  | Specifies the maximum parallelism, which controls the number of concurrent save streams for each NSM backup. The default value is 16.<br><br>This variable is a throttle to control NSM to run fewer save operations concurrently than usual and not to split what would otherwise be one stream.<br><br>To turn off parallelism so that an NSM backup creates only a single save stream at a time, set the value to 1. |

**Table 11** Common Application Information variables (continued)

| Common variable          | Definition                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NSR_PS_SHARED_DIR        | Specifies the full shared directory pathname. This variable is required in a cluster environment to support a full cleanup after an abort. There is no default value.                                                                                                                                                                                                                                                     |
| NSR_PS_SINGLE_LOG        | Specifies whether NetWorker logs all NSM processes together in the <code>nwsnap.raw</code> file. Valid values are TRUE and FALSE. The default value is TRUE if NSR_PS_DEBUG_LEVEL is 3 or less.<br><br>Set to FALSE to cause logging to individual process-based log files.                                                                                                                                               |
| NSR_SNAP_TYPE            | Specifies the snapshot provider.<br><br>Valid values are protectpoint, symm-dmx, emcclar, and emc_rp. If you do not specify a value, NSM tries each of these values in order.                                                                                                                                                                                                                                             |
| NSR_STRICT_SYNC          | Valid values are TRUE and FALSE. The UNIX default value is TRUE. The Microsoft Windows default value is FALSE.<br><br>If TRUE, NSM forces the lgtosync driver or equivalent OS-level capability to freeze and thaw writes to a disk or volume.                                                                                                                                                                            |
| NSR_NSM_RAW_ARRAY_SNAP   | Specifies that this is a LUNs WWN based snapshot or a SmartSnap. This variable should be used for both backup and restore CLIs.                                                                                                                                                                                                                                                                                           |
| SELECT_HOST_VISIBLE_TGTS | Used for Alternate LUN Rollback Restore of SmartSnap backups. Valid values are TRUE and FALSE. Default value is TRUE.<br><br><b>Note</b><br><br>If the value is TRUE, only LUNs from the "NsrSnapSourceSG" Storage Group that are visible to the destination host are chosen. If the value is FALSE, any LUN from the "NsrSnapSourceSG" Storage Group, whether visible or not visible to the destination host, is chosen. |

## Application Information variables for VMAX arrays

The following table lists Application Information variables that NSM can use for VMAX storage arrays.

### Note

The variables in the following table are not relevant for ProtectPoint, unless specifically noted.

**Table 12** Application Information variables for VMAX arrays

| VMAX variable                | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NSR_PS_SYMM_IP               | Valid values are TRUE and FALSE. The default value is TRUE. FALSE prevents the use of intelligent pairing and causes NSM to use only the <code>symm.res</code> file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| NSR_PS_TERMINATE_SRC_MIRRORS | Valid values are TRUE and FALSE. The default value is FALSE. You can use a source LUN for either snapvx or bcv/clone/vpsnap, not both. If you specify snapvx for a source LUN which has an existing bcv/clone/vpsnap mirror, the backup will fail. Set this attribute to TRUE to make NSM first terminate the bcv/clone/vpsnap mirror relationship with the source LUN, and then perform a SnapVX snapshot.                                                                                                                                                                                                                                                                                        |
| NSM_SNAP_SG                  | Valid value must be a valid VMAX storage group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| SYMM_CLONE_FULL_COPY         | Valid values are TRUE and FALSE: <ul style="list-style-type: none"> <li>TRUE—NSM performs a full data copy of a source LUN.</li> <li>FALSE—NSM places the target in COW (CopyOnWrite) mode and will not perform a full data copy.</li> </ul> The default value is TRUE.<br>Notes<br>Notes <ul style="list-style-type: none"> <li>In a single backup/restore session, NSM can use a BCV as a mirror or a clone, but not both.</li> <li>NSM does not allow a rollback operation for a snapshot taken when this variable is FALSE. An attempted rollback will fail.</li> </ul>                                                                                                                        |
| SYMM_EXISTING_PIT            | Valid values are TRUE and FALSE. The default value is FALSE. Specifies the state of targets for the <code>symm.res</code> file. Set to TRUE to prefer a target LUN that is in SPLIT state with the source LUN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| SYMM_IP_TAKE_UNPAIRED        | Valid values are TRUE and FALSE: <ul style="list-style-type: none"> <li>If set to TRUE, intelligent pairing can reuse old, expired mirrors in the NsrSnapSG group that have a relationship to another LUN. Intelligent pairing terminates the old relationship and then pairs the mirror with the new source LUN. Also, intelligent pairing can pair new, unassociated devices in NsrSnapSG to the source LUN.</li> <li>If set to FALSE, intelligent pairing can select only an available mirror from the devices already paired to the source LUN.</li> </ul> The default value is TRUE. Used by intelligent pairing when NSM cannot use any of the mirrors currently paired with the source LUN. |
| SYMM_ON_DELETE               | The default value is RETAIN_RESOURCE. Specifies the state of the mirror device after a backup. These settings are valid only for BCV, VP Snap, and Clone mirrors with SYMM_CLONE_FULL_COPY=TRUE. For VDEV and Clone                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**Table 12** Application Information variables for VMAX arrays (continued)

| VMAX variable               | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                             | <p>mirrors with SYMM_CLONE_FULL_COPY=FALSE, NSM always terminates the relationship:</p> <ul style="list-style-type: none"> <li>• RETAIN_RESOURCE—NSM resynchronizes the mirror again with the source when it deletes the snapshot.</li> <li>• RELEASE_RESOURCE—NSM leaves the mirror in a split state. This setting is recommended with manual backups or when mirrors are frequently rotated (used with a different source).</li> <li>• START_STATE—NSM leaves the target mirror in the same state (split or synced) as before the backup.</li> </ul>                                                                                                                                                                                                                                                                                              |
| SYMM_RB_OVERRIDE_OTHER_TGTS | <p>Valid values are TRUE and FALSE:</p> <ul style="list-style-type: none"> <li>• FALSE—NSM fails the rollback if any other mirrors are in the synchronized state with the source device.</li> <li>• TRUE—Before a rollback operation, NSM splits all synchronized mirrors and then resynchronizes them on completion of a rollback.</li> </ul> <p>The default value is FALSE.</p> <p>Notes:</p> <ul style="list-style-type: none"> <li>• Ensure that the mirror and the source devices are in a synchronized state when using this variable. The status must not be syncInProg or splitInProg. The InProg status will lead to the loss of the snapshot after a rollback attempt.</li> <li>• The InProg status may not occur. If a sync/split of mirror pairs is manual before a rollback, you must wait until the sync/split completion.</li> </ul> |
| SYMM_RES_USE_POLICY         | <p>The default value is ANY.</p> <ul style="list-style-type: none"> <li>• EXISTING—NSM seeks a resource that is synchronized with the source device. This setting reduces the backup time.</li> <li>• FREE—NSM seeks a resource that is not synchronized with any device. The resource must be in a split or not paired state.</li> <li>• ANY—NSM seeks any existing resource first. If NSM finds none, it uses a FREE resource.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                         |
| SYMM_SNAP_POOL              | <p>The default value is C:\programFiles\EMC networker\Nsr\res\symm.res. Defines the pathname of the symm.res file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| SYMM_SNAP_REMOTE            | <p>Valid values are TRUE and FALSE. The default value is FALSE. Set to TRUE if using SRDF. Set to FALSE if not using SRDF.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| SYMM_SNAP_TECH              | <p>Valid values are SNAP, BCV, CLONE, VPSNAP, R2, and SNAPVX. Defines the type of mirroring to use. If set to R2,</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Table 12** Application Information variables for VMAX arrays (continued)

| VMAX variable | Definition                                                  |
|---------------|-------------------------------------------------------------|
|               | then SYMM_SNAP_REMOTE must be TRUE or the backup will fail. |

## Application Information variables for VNX Block arrays

The following table lists Application Information variables that NetWorker Snapshot Management (NSM) can use for VNX Block storage arrays.

**Table 13** Application Information variables for VNX Block arrays

| VNX Block variable   | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CLAR_ON_DELETE       | <p>Specifies the disposition of the clone LUN when NSM deletes a snapshot:</p> <ul style="list-style-type: none"> <li>• <b>RETAIN_RESOURCE</b>—NSM resynchronizes the clone with its source. This makes the clone LUN available for future snapshot requests.</li> <li>• <b>RELEASE_RESOURCE</b>—NSM does not resynchronize the clone with its source. This makes the clone LUN available for other client operations, if you manually remove it from the clone group of the source LUN. This clone LUN will not be available for future snapshot requests unless you manually add it again to the same clone group.</li> <li>• <b>START_STATE</b>—NSM resynchronizes the clone with its source LUN only if it was in a synchronized state when it was fractured.<br/>In this case, its disposition becomes one of the following: <ul style="list-style-type: none"> <li>▪ – <b>RETAIN_RESOURCE</b> workflow</li> <li>▪ – <b>RELEASE_RESOURCE</b> workflow</li> </ul> </li> </ul> <p>Conventional backups to disk or tape that do not use this snapshot capability are still possible with the NetWorker software, even after the upgrade to NSM with the NetWorker client 8.1 installation. The group configuration determines whether a backup uses NSM features.</p> |
| EMCCLAR_SNAP_SUBTYPE | <p>Mandatory. The default value is COW for copy-on-write backup and recovery workflows. The values you can use are as follows:</p> <ul style="list-style-type: none"> <li>• <b>MIRROR</b>—Use for clone fracture, clone backup, and clone recovery workflows.</li> <li>• <b>VNX-SNAP</b>—Use for VNX snap backup and VNX snap recovery workflows.</li> <li>• <b>VNXe-SNAP</b>—Use for VNXe3200 backup and VNXe3200 snap recovery workflows.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Table 13** Application Information variables for VNX Block arrays (continued)

| VNX Block variable | Definition                                                                                                 |
|--------------------|------------------------------------------------------------------------------------------------------------|
|                    | <ul style="list-style-type: none"> <li>COW—Use for copy-on-write backup and recovery workflows.</li> </ul> |
| FRAME_IP           | Specifies the hostname or IP address of the VNX port to use.                                               |

## Application Information variables for RecoverPoint appliances

The following table lists Application Information variables that NetWorker Snapshot Management (NSM) can use for RecoverPoint appliances.

**Table 14** Application Information variables for RecoverPoint appliances

| RecoverPoint variable | Definition                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NSR_SNAP_TECH         | <p>Specifies the RecoverPoint replication type for a backup or restore. The following are RecoverPoint commands that you can use:</p> <ul style="list-style-type: none"> <li>RP_CDP—Use to notify NSM that local copies will be used to access a bookmark.</li> <li>RP_CRR—Use to notify NSM that remote copies will be used to access a bookmark.</li> </ul> |
| RP_APPLIANCE_NAME     | Specifies the hostname or IP address of the RecoverPoint appliance for NSM to use.                                                                                                                                                                                                                                                                            |

## Application Information variables for XtremIO arrays

The following table lists Application Information variables that NSM can use for XtremIO arrays.

**Table 15** Application Information variables for XtremIO arrays

| XtremIO variable                 | Definition                                                                                               |
|----------------------------------|----------------------------------------------------------------------------------------------------------|
| NSR_SNA_TYPE                     | Specifies the XtremIO replication type for a backup or restore. Use the value <code>emc-xtremio</code> . |
| NSR_XTREMIO_HOSTNAME             | Specifies the hostname or IP address of the XtremIO storage array for NSM to use.                        |
| NSR_XTREMIO_PROXY_INITIATOR_NAME | Specifies the proxy or Mount host initiator name created on the XtremIO array by the user.               |





# APPENDIX B

## Command-Line Operations for Snapshot Management

This appendix includes the following topics:

- [Using CLI commands for snapshot operations](#)..... 162
- [Using nsrsnapadmin for snapshot operations](#)..... 162
- [Example nsrsnapadmin operations](#)..... 163
- [Querying with the mminfo command](#)..... 165

## Using CLI commands for snapshot operations

A summary of CLI commands and examples for NetWorker Snapshot operations are provided in this section.

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

## Using nsrsnapadmin for snapshot operations

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

### Note

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

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

**Table 16** Commands and options supported in nsrsnapadmin interactive mode

| NSM operation                                     | Command and available options                                                                                                    |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| Display snapshot save sets                        | <code>p [-s nsr_server] [-c client] [-v] [path   -S ssid]</code>                                                                 |
| Delete a snapshot save set                        | <code>d [-s nsr_server] [-c client] [-v] [-a] [-y] -S ssid [or -S "ssid ssid ..."]</code>                                        |
| Perform a save set recovery                       | <code>R [-s nsr_server] [-c client] [-M mount_host] [-v] -S ssid [-t destination] [-T recover_host] -m path [-A attr=val]</code> |
| Perform a file-by-file browsing and recovery      | <code>r [-s nsr_server] [-c client] [-M mount_host] [-T recover_host] -S ssid [-A attr=val]</code>                               |
| Perform a rollback                                | <code>B [-s nsr_server] [-c client] [-M mount_host] [-T recover_host] [-Fv] -S ssid [-A attr=val] -m path</code>                 |
| Reset the expiration time for a snapshot save set | <code>e time [-s nsr_server] [-c client] [-v] -S ssid [or -S "ssid ssid ..."]</code>                                             |
| Exit the program                                  | <code>q or quit</code>                                                                                                           |

where:

- `nsr_server` is the hostname of the NetWorker server.
- `client` is the hostname of the application client.
- `mount_host` is the hostname of the mount host.
- `-v` is for verbose logging.

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

## Example nsrsnapadmin operations

After you start the `nsrsnapadmin` utility in interactive mode, at the input prompt you can type a specific command and its options to perform a NetWorker Snapshot Management (NSM) operation.

### Querying snapshot save sets

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

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

where:

- *server* is the hostname of the NetWorker server.
- *client* is the hostname of the client from which NSM backed up the data.
- *path* is the pathname of a particular snapshot save set. Type the pathname to query a single save set only. Otherwise, the output message lists all the save sets.

A message similar to the following appears:

```
nsrsnapadmin> p -s ledma038 -c ledma218
ssid = 3742964283 savetime="February 11, 2013 11:20:10 AM EST"
(1360599610) expiretime="February 11, 2014 11:59:58 PM EST"
(1392181198) sssize=/symm_403_ufs
```

### File-by-file browsing and restore

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

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

where *client* can be a single host IP or a cluster IP (virtual, actual, or public IP).

### Rollback restore

A rollback is a complete restore of all the application source LUNs involved in the snapshot backup. The restore includes all the file systems and the volume groups that reside on these production LUNs. The `nsrsnapadmin` utility supports forced rollback and the safety check features.

To perform a rollback restore, type the following command at the prompt:

```
B -S ssid /source_path
```

For example:

```
B -s server -c client -Fv -M mount_host -S ssid -m source_path
```

where *client* can be a single host IP or a cluster node (cluster IP or public IP).

[Restoring a snapshot by rollback](#) on page 141 provides more information.

## Deleting a snapshot save set

You can type the `nsrsnapadmin`, `nsrmm`, `nsrim`, or `nsrsnapck` command to delete snapshot save sets.

Deleting an NetWorker Snapshot Management (NSM) save set is similar to deleting a standard NetWorker save set. NSM deletes the physical snapshot from the storage array and then deletes all save sets that refer to that physical snapshot from the media database.

For example:

```
command -d -s server -S ssid
```

where:

- *command* is `nsrmm` or `nsrsnapck` if you do not use `nsrsnapadmin`.
- *server* is the hostname of the NetWorker server.
- *ssid* is the snapshot save set ID.

## Modifying the retention period of a snapshot save set

You can type the `e` command at the `nsrsnapadmin` prompt to modify the expiration date of a snapshot. For example:

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

where:

- *time* is the date and time when the snapshot save set expires. Acceptable date formats are as follows:
  - `mm/dd[/yy]`
  - `month_name dd[/yy]`
 Acceptable time formats are as follows:
  - `hh[:mm[:ss]] [meridian] [zone]`
  - `hhmm [meridian] [zone]`
- *server* is the hostname of the NetWorker server.
- *ssid* is the ID of the snapshot save set that you want to modify.
- *client* (optional) is the hostname of the client from which NSM backed up the data.

A message similar to the following appears:

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

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

Notes:

- If you omit the year, the year defaults to the current year.

- If you omit the meridian, NSM uses a 24-hour clock.
- If you omit the time zone (for example, GMT), NSM uses the current time zone.
- If you specify a date mm/dd/yy (for example `-e 09/04/17`), the time defaults to 00:00:00. NSM changes the snapshot save set's browse and retention times to 09/04/17 00:00:00.

If you specify a time hh:mm:ss (for example `-e 20:00:00`), the date defaults to the system date, for example, 09/03/17. NSM changes the snapshot save set's browse and retention times to 09/03/17 20:00:00.

## Querying with the `mminfo` command

You can query a client's snapshot save sets by typing the `mminfo` command. The `-q snap` option lists all snapshot save sets for a particular client.

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

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

where:

- `server` is the hostname of the NetWorker server.
- `clientname` is the hostname of the client from which NSM backed up the data.

Example output:

```
mminfo -s ledma038 -S -a -q "client=ledma218,snap"
volume client date size level name
ledma038.003 ledma218 02/11/13 2 KB full /symm_403_ufs
```

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



# APPENDIX C

## Migrating Legacy PowerSnap Configurations

This appendix includes the following topics:

- [Migrating legacy PowerSnap configurations to NSM](#).....168
- [Deprecated Client resource attributes](#)..... 169
- [Migrating VMAX \(Symmetrix\) arrays](#).....169
- [Migrating VNX \(CLARiiON\) arrays](#)..... 169
- [Migrating RecoverPoint appliances](#)..... 169
- [Starting the nsrpsd process](#).....170
- [Licensing](#).....170

## Migrating legacy PowerSnap configurations to NSM

You can migrate legacy NetWorker PowerSnap Module configurations to NetWorker Snapshot Management. The NetWorker client installation provides all the functionality that is previously handled by the PowerSnap Module.

### Removing PowerSnap on UNIX systems

Before you upgrade to NSM on UNIX systems, remove the existing PowerSnap packages by using the native package management utilities of the operating system.

---

#### Note

Failure to remove previously installed PowerSnap packages causes the NetWorker client installation to fail when performed through a client push installation or the native package management utilities for the operating system.

---

Remove the PowerSnap packages from all computers that participate in the migration:

- Remove the following packages on Linux:
  - lgtopsag-2.5.1.1.x86\_64.rpm
  - lgtopseg-2.5.1.1.x86\_64.rpm
  - lgtopssc-2.5.1.x86\_64.rpm
- Remove the following PowerSnap packages on Solaris:
  - LGTOpsag
  - LGTOpseg
  - LGTOpssc
- Remove the following packages on AIX:
  - LGTOps.psag.rte
  - LGTOps.pseg.rte
  - LGTOps.pssc.rte
- Remove the PowerSnap.pkg package on HP.

### Removing PowerSnap on Microsoft Windows systems

Before you upgrade to NSM on Windows systems, you do not need to uninstall the PowerSnap Module. The NetWorker client installation wizard for Microsoft Windows will uninstall the old PowerSnap Module automatically and replace it with the NSM feature.

Use the upgrade option of the NetWorker client installer.

---

#### Note

Upgrading unsupported Microsoft Windows platforms will both uninstall existing PowerSnap packages and not install the NSM feature of the NetWorker client.

---



## Deprecated Client resource attributes

The Client resource Application Information attribute no longer supports the following variables. The presence of these variables will cause a backup to fail:

- NSR\_IMAGE\_SAVE
- SYMM\_PROVIDER\_DB
- SYMM\_PROXY\_PROVIDER\_DB

## Migrating VMAX (Symmetrix) arrays

Before the upgrade to NetWorker Snapshot Management (NSM), remove any snapshots that were created with the PowerSnap Module from the VMAX array. You can delete the snapshots or clone them to conventional storage media.

Ensure that the operating system, versions, and configuration support NSM. [Components of the NSM network](#) provide details.

Decide whether you continue to use a Symmetrix/VMAX resource file (`symm.res`) or take advantage of NSM intelligent pairing.

---

### Note

The migration procedure does not remove the `symm.res` resource file. The `symm.res` file is optional for NSM, but NSM uses it if present.

VMAX disk groups are no longer required for NSM to operate. If present, NSM ignores them.

---

[Pairing source LUNs to mirror LUNs](#) on page 109 provides details on intelligent pairing.

## Migrating VNX (CLARiiON) arrays

Before the upgrade to NetWorker Snapshot Management (NSM), remove any snapshots that were created with the PowerSnap Module from the VNX array. You can delete the snapshots or clone them to conventional storage media.

Ensure that the operating system, versions, and configuration support NSM. [Components of the NSM network](#) provides details.

---

### Note

The existing VNX (CLARiiON) security files must continue to exist on all nodes that participate in snapshot operations. If you have removed these security files, you can re-create the files through the command line or with the NetWorker Client Configuration wizard. [Configuring the Navisphere security file](#) provides details.

---

## Migrating RecoverPoint appliances

Before the upgrade to NetWorker Snapshot Management (NSM), ensure that the operating system, versions, and configuration support NSM. [Components of the NSM network](#) provides details.

You must create a NetWorker Client resource by using the NMC Client Configuration Wizard. RecoverPoint appliances do not support a nonwizard configuration and existing RecoverPoint Client resources will not work after the upgrade to NSM.

You must configure RecoverPoint credentials in the NetWorker server lockbox because the `nsr_rp_access_config` utility no longer exists, and you cannot use a local credential file.

#### Procedure

1. Ensure that you have available the credentials for username and password for the RecoverPoint appliances that participate in snapshot backups.
2. Open an NMC session to the NetWorker server, and then run the Client Configuration Wizard.
3. Create a Client resource for the application host by selecting NSM and the RecoverPoint option.
4. Compare the new resource and old resource, and then add the required attributes from the old configuration to the new one.
5. On the **Specify the RecoverPoint replication type and Storage Array Options** page, add the RecoverPoint username and password to the lockbox on the NetWorker server.
6. After you finish with the wizard, delete the old configuration.

## Starting the nsrpsd process

In NetWorker, the `nsrpsd` process on the application host starts on demand by NetWorker processes, such as `nsrsnap`. After 30 minutes of inactivity, the `nsrpsd` process terminates. To prevent `nsrpsd` from terminating, create the file `nsrpsd_stay_up` in the `nsr/res` directory.

If you use a version of NMDA or NMSAP on the application host, then the `nsrpsd` process does not automatically start or stop and backups of these applications will fail. You must start or stop the `nsrpsd` process manually, as done in previous PowerSnap releases. In these environments, `nsrpsd` will not terminate after 30 minutes of inactivity.

## Licensing

For the upgrade to NetWorker Snapshot Management (NSM), you do not require any new licenses. NSM honors existing PowerSnap licenses and the NetWorker capacity and traditional licensing models.

[NetWorker snapshot licensing requirements](#) on page 26 provides details.

# GLOSSARY

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

## A

- administrator** Person who normally installs, configures, and maintains software on network computers, and who adds users and defines user privileges.
- application host** Computer that hosts application data to be backed up.

## B

- backup**
1. Duplicate of database or application data, or an entire computer system, stored separately from the original, which can be used to recover the original if it is lost or damaged.
  2. Operation that saves data to a volume for use as a backup.
- backup volume** A volume used to store backup data. NetWorker backup data cannot be stored on an archive volume or a clone volume.

## C

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

**Console server** See [NetWorker Management Console \(NMC\)](#).

## D

**datazone** Group of clients, storage devices, and storage nodes that are administered by a NetWorker server.

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

## F

**file index** See [client file index](#).

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

## G

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

## H

**host** Computer on a network.

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

## L

**logical unit (LUN)** Logical storage space on a storage array that is addressed by SCSI or Fibre Channel protocols. Multiple LUNs can be used to represent a storage volume.

**logical volume manager (LVM)** Software that controls disk resources by mapping data between a logical view of storage space and the actual physical disks.

## M

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

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

## N

|                                            |                                                                                                                                                                                                                                                  |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NetWorker Management Console (NMC)</b>  | Software program that is used to manage NetWorker servers and clients. The NMC server also provides reporting and monitoring capabilities for all NetWorker processes.                                                                           |
| <b>NetWorker server</b>                    | Computer on a network that runs the NetWorker server software, contains the online indexes, and provides backup and restore services to the clients and storage nodes on the same network.                                                       |
| <b>NetWorker Snapshot Management (NSM)</b> | Technology that provides point-in-time snapshot copies of data. NetWorker software backs up data from the snapshot. This allows applications to continue to write data during the backup operation, and ensures that open files are not omitted. |

## P

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

## Q

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

## R

- resource** Software component whose configurable attributes define the operational properties of the NetWorker server or its clients. Clients, devices, schedules, groups, and policies are all NetWorker resources.
- restore** To retrieve individual data files from backup media and copy the files to a client without applying transaction logs.
- retention policy** NetWorker setting that determines the minimum period of time that backup data is retained on a storage volume and available for recovery. After this time is exceeded, the data is eligible to be overwritten.
- rollback restore** Process by which a snapshot is restored to its source or alternate location by using the capability of the storage array. A rollback restore destroys existing data on the target location.

## S

- save set**
1. Group of files or a file system copied to storage media by a backup or snapshot rollover operation.
  2. NetWorker media database record for a specific backup or rollover.
- save set ID (ssid)** Internal identification number assigned to a save set.
- scheduled backup** Type of backup that is configured to start automatically at a specified time for a group of one or more NetWorker clients. A scheduled backup generates a bootstrap save set.
- Single node** A configuration where the application host and mount host are the same host.
- snapshot** Point-in-time, read-only copy of specific data files, volumes, or file systems on an application host. Operations on the application host are momentarily suspended while the snapshot is created on a proxy host. Also called a PIT copy, image, or shadow copy.
- snapshot backup** Snapshot created on a storage array as a backup. Previously called instant backup.
- snapshot mount host** See [mount host](#).
- snapshot restore** Restore from a snapshot backup. Previously called instant restore.
- snapshot save set** Group of files or other data included in a single snapshot. Previously called a snapset.
- source LUN** LUN on the application host, whose production data can be copied by snapshot.
- specific point-in-time (SPIT) copy** RecoverPoint term for PIT copy. See [point-in-time copy \(PIT copy\)](#).
- ssid** See [save set ID \(ssid\)](#) .
- storage array** Integrated collection of subsystem disks, controllers, bus adapters, and software that provides storage services to one or more hosts.

**storage node** Computer that manages physically attached storage devices or libraries, whose backup operations are administered from the controlling NetWorker server. Typically a “remote” storage node that resides on a host other than the NetWorker server.

## T

**target client** NetWorker client on which data is to be restored. This may be the same as the original source client from which the data was backed up, or it may be a different client.

## V

- volume**
1. Unit of physical storage medium, such as a disk or magnetic tape, to which backup data is written.
  2. Identifiable unit of data storage that may reside on one or more computer disks.

**volume group** Logical group that contains sets of disks, as configured by a volume manager.

