EMC® NetWorker®

Version 9.0.x

Snapshot Management Integration Guide

302-001-777

REV 05

EMC²

Copyright © 2001-2016 EMC Corporation. All rights reserved. Published in the USA.

Published June 2016

EMC 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. EMC Corporation 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 EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (https://support.emc.com).

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

CONTENTS

Figures 9

Tables 11

Preface 13

Chapter 1 Overview of NetWorker Snapshot Features 17

NetWorker snapshot management product description ... 18
- Snapshot operations ... 18
- Types of snapshot backups ... 19
- Types of snapshot recoveries ... 20
- NetWorker clone support ... 20
- Backup configuration methods ... 20
- Restore methods ... 21
- Monitoring and reporting snapshot operations ... 21
- Internationalization support ... 21

Components of the snapshot environment ... 22
- Application host ... 22
- FC and iSCSI environments ... 22
- Storage arrays ... 22
- NetWorker server ... 23
- NetWorker storage node ... 23
- Snapshot mount host ... 23
- Backup storage media ... 23
- NetWorker application modules ... 23
- Third-party volume managers ... 23

NetWorker snapshot licensing requirements ... 24

Example NetWorker snapshot environments ... 24
- Example of a snapshot and clone to storage media ... 25
- Example of a restore from a snapshot backup ... 27

Chapter 2 Data Protection Policies 29

Strategies for storage array snapshot backups ... 30
- Default data protection policies ... 30
- Road map for configuring a new data protection policy ... 31
- Creating a policy ... 32
- Creating a workflow in a new policy ... 34
- Protection groups for snapshot backups ... 35
- Actions supported in snapshot backups ... 39
- Supported actions in snapshot workflows ... 40
- Visual representation of snapshot workflows ... 55

Chapter 3 Software Configuration 57

Backup group resource migration ... 58

Roadmap for snapshot configurations ... 60

Snapshot configuration prerequisites ... 60

Storage Array specific prerequisites 61
Application host prerequisites 61
Mount host prerequisites 62
Storage node prerequisites 63
Configuring the user privileges 63
Configuring snapshot backups with the client wizard 64
Configuring snapshot backups manually 68
Configuring the Client resource manually for the application host 68
Configuring the Client resource manually for a mount host 70
Configuring the Application Information variables 70
Configuring preprocessing and postprocessing scripts 71

Chapter 4 Configuring ProtectPoint on VMAX 73

Overview 74
ProtectPoint on VMAX prerequisites 74
Enabling vDisk on a Data Domain system 74
Provisioning protection devices on Data Domain systems 75
Completing the VMAX system configuration 76
Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard 78
Considerations for ProtectPoint device and NetWorker ProtectPoint enabled pools 80
VMAX3 SRDF/S support 80
Rollbacks in the SRDF/S environment 81
Configuring Data Domain NsrSnapSG device groups for intelligent pairing 81
Intelligent Pairing vDisk selection decision tree 81
Intelligent Pairing allocates vDisk for mount, validate, and restore 82

Chapter 5 Configuring ProtectPoint on RecoverPoint with XtremIO 83

Overview 84
Basic backup workflow 84
Basic restore workflow 84
ProtectPoint for RecoverPoint on XtremIO prerequisites 85
Enabling vdisk on the Data Domain 86
Provisioning protection devices on Data Domain systems 87
Provisioning protection devices on Data Domain systems 88
Configuring RecoverPoint and XtremIO storage 89
Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard 90
Considerations for ProtectPoint device and NetWorker ProtectPoint enabled pools 92
Configuring Data Domain NsrSnapSG device groups for intelligent pairing 92
Intelligent Pairing vDisk selection decision tree 93
Intelligent Pairing allocates vDisk for mount, validate, and restore 93

Chapter 6 Configuring snapshots on XtremIO arrays 95

Snapshot support for XtremIO 96
Snapshot operation with XtremIO REST API 96
Prerequisite for XtremIO configurations 96
Supported XtremIO features 96
Snapshot management policy with XtremIO 96
Snapshot backups with XtremIO 96
XtremIO configuration methods 97

Chapter 7 Configuring snapshots on VMAX Storage Arrays 99

Snapshot support of VMAX storage arrays 100
Snapshot operations with TimeFinder software 100
Prerequisites and support for VMAX configurations 100
Types of supported mirror devices 100
Pairing source LUNs to mirror LUNs 101
Intelligent pairing 101
Configuring NsrSnapSG storage groups for intelligent pairing 102
Manual pairing LUNs with the symm.res file 102
Configuring the symm.res file 103
VMAX SRDF/S support 104
Rollbacks in the SRDF/S environment 104
Solutions Enabler Client and Server mode configuration 105
Configuring the Solutions Enabler in Client/Server mode on a SAP client 105
Configuring the Solutions Enabler in Client/Server mode on a Mount Host 105
Known limitation for VMAX 106

Chapter 8 Configuring snapshots on VNX Block Storage Arrays 107

Snapshot support of VNX Block storage arrays 108
Snapshot operations with SnapView software 108
Prerequisites and support for VNX configurations 108
Configuring the Navisphere security file 108
Creating the Navisphere file manually on UNIX systems 108
Creating the Navisphere file manually on Windows systems 109
Configuring Unisphere CLI on VNXe3200 109
UEMCLI Windows registry setup 110

Chapter 9 Configuring snapshots on RecoverPoint 111

Snapshot support of RecoverPoint 112
Snapshot operations with RecoverPoint software 112
Prerequisite for RecoverPoint configurations 112
Restrictions for RecoverPoint configurations 112
Supported RecoverPoint features 112
Snapshot management policy 113
RecoverPoint configuration methods 113
RecoverPoint snapshot retention 114

Chapter 10 Configuring snapshots in a Cluster Environment 115

NetWorker support of cluster environments 116
Failover with snapshots in a cluster environment 116
Configuring a cluster environment for snapshots 116
AIX systems in a cluster environment 117
ProtectPoint restore and rollback for VCS on Solaris 118
Performing a ProtectPoint VCS Restore 118
Performing a ProtectPoint VCS rollback 120

Chapter 11 Data Management and Recovery 123

Snapshot lifecycle management 124
Management and recovery of file system snapshot data 124

Save set IDs and expiration policies 124
Browsing snapshot and clone save sets 124
Snapshot recovery support and limitations 124
Raw partitions and raw devices 125
Restoring from a snapshot with the Recovery Wizard 125
Restoring a snapshot by rollback 128
Rollback considerations 128
Configuring the psrollback.res file 129
Rollbacks with Veritas Volume Manager 130
Rollbacks with IBM AIX Volume Manager 130

Chapter 12 Troubleshooting 133

NetWorker snapshot backup issues 134
Backup on Windows fails with a Delayed Write Failed error 134
Snapshots fail to mount for AIX managed file systems 134
Snapshots fail for Linux Volume Manager on VNX with PowerPath .. 134
Linux LVM snapshots fail with an error 134
NetWorker to Media-Clone stops responding and the backup fails . 134
NetWorker snapshot restore issues 135
Restore of raw devices fails on Linux with permission issue 135
Command nsrsnap_recover -I completes but fails to restore a file .. 135
Restore fails with disk signature error 135

Appendix A Application Information Variables 137

Using Application Information variables 138
Common Application Information variables 138
Application Information variables for VMAX arrays 139
Application Information variables for VNX Block arrays 141
Application Information variables for RecoverPoint appliances 142
Application Information variables for XtremIO arrays 143

Appendix B Command-Line Operations for Snapshot Management 145

Using CLI commands for snapshot operations 146
Using nsrsnapadmin for snapshot operations 146
Example nsrsnapadmin operations 147
Querying snapshot save sets 147
File-by-file browsing and restore 147
Rollback restore 147
Deleting a snapshot save set 148
Modifying the retention period of a snapshot save set 148
Querying with the mminfo command 149

Appendix C Migrating Legacy PowerSnap Configurations 151

Migrating legacy PowerSnap configurations to NSM 152
Removing PowerSnap on UNIX systems 152
Removing PowerSnap on Microsoft Windows systems 152
Deprecated Client resource attributes 153
Migrating VMAX (Symmetrix) arrays 153
Migrating VNX (CLARiiON) arrays 153
Migrating RecoverPoint appliances 153
Starting the nsrpsd process 154

Licensing 154

Glossary **155**

FIGURES

1 Snapshot and clone with the storage node as the mount host 25
2 Snapshot and clone with the application host as the mount host 25
3 Restore from a snapshot with the storage node as the mount host 27
4 Platinum policy configuration 30
5 Gold policy configuration 31
6 Silver policy configuration 31
7 Bronze policy configuration 31
8 Data protection policy example 32
9 All possible workflow actions for a snapshot backup 40
10 Sample snapshot workflow 55
11 Snapshot and clone in a cluster environment 116

TABLES

1 Revision history ... 13
2 Style conventions ... 15
3 Save set criteria ... 38
4 Schedule icons ... 41
5 Schedule icons ... 44
6 Backup type icons ... 47
7 Schedule icons ... 51
8 Migration of Group attributes ... 58
9 vdisk object hierarchy mapping ... 75
10 vdisk object hierarchy mapping ... 87
11 vdisk object hierarchy mapping ... 88
12 Common Application Information variables ... 138
13 Application Information variables for VMAX arrays ... 139
14 Application Information variables for VNX Block arrays ... 141
15 Application Information variables for RecoverPoint appliances ... 142
16 Application Information variables for XtremIO arrays ... 143
17 Commands and options supported in nsrsnapadmin interactive mode ... 146

Preface

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

Contact your EMC technical support professional if a product does not function correctly or does not function as described in this document.

Note

This document was accurate at publication time. Go to EMC Online Support (https://support.emc.com) to ensure that you are using the latest version of this document.

Purpose

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

Audience

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

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

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
05	June 29, 2016	Updated this document for EMC NetWorker 9.0.1. with the following: • Added "Configuring ProtectPoint on RecoverPoint with XtremIO" chapter. • Added "Configuring snapshots on XtremIO arrays" chapter. • Updated the "Configuring snapshot backups with the client wizard" topic to reflect support for additional storage arrays. • Added "Application Information variables for XtremIO arrays" topic.
04	April 13, 2016	Updated the "Strategies for storage array snapshot backups" section with the following changes:

Table 1 Revision history (continued)

Revision	Date	Description
		• Revised the order in which to create data protection policy resources. • Added the section "Default data protection policies."
03	January 25, 2016	Minor update to this document for EMC NetWorker 9.0.
02	December 14, 2015	Minor updates to this document for EMC NetWorker 9.0.
01	September 24, 2015	First release of this document for EMC NetWorker 9.0.

Related documentation

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

- *EMC NetWorker Online Software Compatibility Guide*
  Provides a list of client, server, and storage node operating systems supported by the EMC information protection software versions. You can access the Online Software Compatibility Guide on the EMC Online Support site at https://support.emc.com. From the Support by Product pages, search for NetWorker using "Find a Product", and then select the Install, License, and Configure link.
- *EMC NetWorker Administration Guide*
  Describes how to configure and maintain the NetWorker software.
- *EMC NetWorker Network Data Management Protocol (NDMP) User Guide*
  Describes how to use the NetWorker software to provide data protection for NDMP filers.
- *EMC NetWorker Cluster Integration Guide*
  Contains information related to configuring NetWorker software on cluster servers and clients.
- *EMC 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.
- *EMC NetWorker Updating from a Previous Release Guide*
  Describes how to update the NetWorker software from a previously installed release.
- *EMC NetWorker Release Notes*
  Contains information on new features and changes, fixed problems, known limitations, environment and system requirements for the latest NetWorker software release.
- *EMC NetWorker Command Reference Guide*
  Provides reference information for NetWorker commands and options.
- *EMC 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.
- *EMC NetWorker Performance Optimization Planning Guide*
  Contains basic performance tuning information for NetWorker.

- *EMC NetWorker Server Disaster Recovery and Availability Best Practices Guide*
  Describes how to design and plan for a NetWorker disaster recovery. However, it does not provide detailed disaster recovery instructions. The Disaster Recovery section of the NetWorker Procedure Generator (NPG) provides step-by-step disaster recovery instructions.
- *EMC 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 EMC storage arrays.
- *EMC 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.
- *EMC NetWorker VMware Integration Guide*
  Provides planning and configuration information on the use of VMware in a NetWorker environment.
- *EMC NetWorker Error Message Guide*
  Provides information on common NetWorker error messages.
- *EMC NetWorker Licensing Guide*
  Provides information about licensing NetWorker products and features.
- *EMC NetWorker REST API Getting Started Guide*
  Describes how to configure and use the NetWorker REST API to create programmatic interfaces to the NetWorker server.
- *EMC NetWorker REST API Reference Guide*
  Provides the NetWorker REST API specification used to create programmatic interfaces to the NetWorker server.
- EMC 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.
- EMC NetWorker User Online Help
  The NetWorker User program is the Windows client interface. Describes how to use the NetWorker User program which is the Windows client interface 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

EMC uses the following conventions 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

EMC uses the following type style conventions in this document:

Table 2 Style conventions

Bold	Used for names of interface elements, such as names of buttons, fields, tab names, and menu paths (what the user specifically selects or clicks)

Table 2 Style conventions (continued)

Italic	Used for full titles of publications that are referenced in text
`Monospace`	Used for: • System code • System output, such as an error message or script • Pathnames, file names, prompts, and syntax • Commands and options
Monospace italic	Used for variables
`Monospace bold`	Used for user input
[ ]	Square brackets enclose optional values
|	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate non-essential information that is omitted from the example

Where to get help

EMC support, product, and licensing information can be obtained as follows:

Product information

For documentation, release notes, software updates, or information about EMC products, go to EMC Online Support at https://support.emc.com.

Technical support

Go to EMC Online Support and click Service Center. Several options for contacting EMC Technical Support appear on the site. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Online communities

Go to EMC Community Network at https://community.emc.com for peer contacts, conversations, and content on product support and solutions. Interactively engage online with customers, partners, and certified professionals for all EMC products.

Your comments

Your suggestions help to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to DPAD.Doc.Feedback@emc.com.

CHAPTER 1

Overview of NetWorker Snapshot Features

This chapter includes the following topics:

- NetWorker snapshot management product description 18
- Components of the snapshot environment 22
- NetWorker snapshot licensing requirements 24
- Example NetWorker snapshot environments 24

NetWorker snapshot management product description

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

The NetWorker extended client installation provides all the functionality that the EMC NetWorker PowerSnap™ Module previously handled. The *EMC NetWorker Installation Guide* provides more details. Migrating Legacy PowerSnap Configurations on page 151 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 you use with NetWorker Snapshot Management.

The *EMC NetWorker Online Software Compatibility Guide* provides details on versions that NetWorker supports, including volume managers, NetWorker modules, and cluster environments.

The *EMC NetWorker Release Notes* provides details on NetWorker Snapshot Management features, support, fixed problems, and known problems with workarounds.

Snapshot operations

NetWorker Snapshot Management supports the application host, which is a NetWorker client that writes production data to volumes on a supported EMC 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 a local LUN or a LUN on a remote array or remote appliance.

NetWorker supports the following storage array and storage appliance configurations:

- EMC ProtectPoint—EMC VMAX3™ or EMC XtremIO® to Data Domain vdisk snapshot operations.
- EMC VMAX® arrays—EMC TimeFinder® Clone, VDEV, BCV, VP Snap™, SnapVX, and Symmetrix Remote Data Facility (EMC SRDF®) operations.
- EMC VNX® and EMC VNXe® Block arrays—EMC SnapView™ Copy-on-write (COW/Snapshot), Mirror (clone), and VNX Snap operations.
- EMC 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 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 on page 24 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 retains the snapshot on the storage array only. 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 retains the snapshot on a Data Domain device only. 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 in this document refers to both storage array only snapshots, as well as 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 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 EMC NetWorker Module for SAP (NMSAP) has configuration parameters that allow you to exclude the snapshot of specific files. The *EMC 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 *EMC NetWorker Online Software Compatibility Guide* provides details.

Types of snapshot recoveries

The type of recoveries that you can perform for snapshot-based data depends on the location of the data and certain other factors:

- Restore from a snapshot—NetWorker mounts the snapshots to a mount host, browses, and then 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 the 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 clone support

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

- DD Boost, Advanced File Type Device (AFTD), Tape—You can clone any types of snapshots to these types of conventional media.
- DD ProtectPoint—You can clone VMAX3 SnapVX and RecoverPoint/XtremIO snapshots to a DD ProtectPoint device.

Note

DD ProtectPoint is not applicable for NAS.

Cloning supports full, incremental, and cumulative incremental cloning. Cloning permits restores to leverage traditional NetWorker restore methods.

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

- For file system backups—NetWorker records the content of the snapshot in the client file index (CFI) only during a clone operation to conventional media.
- For application backups—NetWorker records the content of the backups (only the application files being protected) in the CFI for both snapshot and clone operations.

The *EMC 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 using the NetWorker Management Console (NMC) interface. All the supported storage arrays support the following configuration methods:

- NetWorker Client Configuration Wizard—EMC recommends 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 verification of configuration dependencies.

- NMC Client Properties windows—Provides a manual, nonwizard user interface that you can use to create or to 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 on page 137.

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

Restore methods

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

- The NMC Recovery Wizard—EMC recommends that you use the wizard interface to restore data from the snapshots and conventional storage media.
- The `nsrsnapadmin` command utility—An interactive Command Line Interface (CLI) tool for various snapshot-related operations, including restore from a snapshot. Using nsrsnapadmin for snapshot operations on page 146 provides details.
- The `nsrsnap_recover` command—Another CLI method to restore data from a snapshot or conventional storage media.

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

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

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, 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 24 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 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.

Install the following software on each application host meets the following requirements:

- NetWorker client 9.0.x.
- NetWorker extended client 9.0.x.
- NMDA/NMSAP 9.0.x if you are protecting DB2, Oracle, 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). NSM supports iSCI for VMAX, VNX, and XtremIO. NetWorker snapshots do not support FCoE environments.

Note

NSM supports the VMware guest operating environments for VNX when those VNX volumes are using iSCSI. NSM does not support RDM volumes with VNX. For 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:

- EMC VMAX (Symmetrix) storage array
- EMC VNX Block (CLARiiON) storage array
- EMC VNXe Block
- EMC RecoverPoint storage appliance
- EMC XtremIO native snapshot

The *EMC NetWorker Online Software Compatibility Guide* provides support details about the latest qualified 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, and tape. Snapshot management requires a storage node for all 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 cloning to conventional storage media.

This 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 considers 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, DD Boost devices, and tape.

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 EMC RecoverPoint.

The *EMC NetWorker Online Software Compatibility Guide* provides details on supported versions.

The *EMC NetWorker Module for Databases and Applications Administration Guide* and the *EMC 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 on separate hosts only, such as the application host and a different mount host.
- If multiple LUNs with the same disk signature or 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 *EMC NetWorker Online Software Compatibility Guide* provides support details.

NetWorker snapshot licensing requirements

Two types of licensing can enable snapshot management:

- NetWorker capacity licenses
- Traditional licenses

In both cases, the NetWorker software reports on capacities that are consumed for the standard (nonsnapshot) backups and the snapshot backups.

If NetWorker detects valid older PowerSnap licenses, then NetWorker honors these licenses.

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

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

The *EMC NetWorker Licensing Guide* and an EMC sales representative can provide details about both 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 18 describes basic snapshot concepts.

Example of a snapshot and clone to storage media

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

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



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



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/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 image 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 25.
   - A remote mount host with the NetWorker client and extended client software installed.

Note

If the NetWorker Client resource specifies the Client Direct and DD Boost options, then on-client data deduplication processing occurs on the mount host during clone operations to conventional media.

4. NetWorker manages the snapshot according to the options in the Client resource 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 *EMC 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 and description of processes illustrate 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



1. NetWorker selects the snapshots that contain the data you want to restore. NetWorker then mounts the snapshot on the mount host.
2. If you are restoring file system data, you browse for 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 as shown in the preceding figure.
   - 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.

CHAPTER 2

Data Protection Policies

This chapter includes the following topics:

- Strategies for storage array snapshot backups 30

Strategies for storage array snapshot backups

Multiple data protection policy strategies are available to help you optimize how NSM perform snapshot backups.

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

- Probe
- Check connectivity
- Snapshot backup
- Clone

Actions supported in snapshot backups on page 39 provides more details.

Default data protection policies

NetWorker provides you with preconfigured data protection polices resources that you can use immediately to protect your environment, modify to suit your environment, or use an example to create new resource configurations. To use these policy resources, you must add clients to the appropriate group resource.

Each protection policy provides an example of the EMC best practices that you should follow when you design your data protection solution:

- Separate file system backups from application database backups, to provide ease of access at recovery time.
- Stagger the start times for file system backup from the application database backups, to prevent disk contention on the target hosts.

The default data protection policy resources mimic the requirements of a service provider, with different policies to that are designed to provide protection based on service level agreements.

Platinum policy

The Platinum policy provides you with an example of a data protection policy for an environment that contains EMC storage arrays or 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. The policy contains two workflows, one to protect Hyper-V hosts and one to protect VMware hosts. Each workflow contains a backup action followed by a clone action.

Figure 5 Gold policy configuration



Silver policy

The Silver policy provides an example of a data protection policy for an environment that contains non-virtualized machines and requires backup data redundancy. The policy contains two workflows, one to protect hosts file systems and one to protect database applications. Each workflow contains a backup action followed by a clone action.

Figure 6 Silver policy configuration



Bronze policy

The Bronze policy provides an example of a data protection policy for an environment that contains non-virtualized machines. The policy contains two workflows, one to protect hosts file systems and one to protect database applications. Each workflow contains a backup action.

Figure 7 Bronze policy configuration



Road map for configuring a new data protection policy

Procedure

1. Create a policy.

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

2. Within the policy, create a workflow 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.

   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 actions for the workflow.
5. To define the backup data that you want to protect, configure Client resources, and then assign the client resources to a protection group.

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

Figure 8 Data protection policy example



Creating a policy

Procedure

1. On 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 128.

Note

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

4. In the **Comment** box, 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 all the actions in the policy complete, select **On Completion**.
   - To send a notification with information about each failed workflow and action after all the actions in the policy complete, select **On Failure**.
6. In the **Send notification** attribute when you select the **On Completion** or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` action to write the notifications to a log file or configure an email notification.

   The default notification action is to log the information to the `policy_notifications.log` file. The `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows, by default. You can use the `smtpmail` application on Windows or the default mailer program on Linux to send email messages.

   For example:

   - To log notifications to a file named `policy_notifications.log`, type the following command:

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

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

     ```
     mail -s subject recipient
     ```

   - 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.

Creating a workflow in a new policy

A policy must contain one or more workflows.

Procedure

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

   The maximum number of characters for the name of the group is 64.
4. In the **Comment** box, type a description for the workflow. The maximum number of characters for the **Comment** field is 128.
5. 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 determine when to send the notification, select **Set at policy level**.
   - To send notifications with information about each successful and failed workflow and action, after all the actions in the workflow complete, select **On Completion**.
   - To send notifications with information about each failed workflow and action, after all the actions in the workflow complete, select **On Failure**.
6. In the **Send notification** attribute when you select the **On Completion** or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` action to write the notifications to a log file or configure an email notification.

   The default notification action is to log the information to the `policy_notifications.log` file. The `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows, by default. You can use the `smtpmail` application on Windows or the default mailer program on Linux to send email messages.

   For example:

   - To log notifications to a file named `policy_notifications.log`, type the following command:

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

     ```
     mail -s subject recipient
     ```
   - 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.

7. In the **Running** group box, define when and how often the workflow runs.
   a. To ensure that the actions contained in the workflow run when the policy or workflow is started, 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 is started, clear this option.
   b. To ensure that the workflow starts 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 running at the time that is specified in the **Start time** attribute, clear this option.
   c. To define the time to start the actions in the workflow, in the **Start Time** attribute, use the spin boxes.

      The default value is 9:00 P.M.
   d. To define how frequently to repeat the actions that are defined in the workflow over a 24 hour period, In the **Interval** attribute, use the spin boxes.

      The default 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 define the last time to start a workflow in a defined interval period, use the spin boxes .
   e. To define 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, when you set the **Start Time** to 7:00 PM, the **Interval** to 1 hour, and the **Interval end time** to 11:00 P.M., then the workflow automatically starts every hour beginning at 7:00 P.M. and the last start time is 11:00 PM.
8. 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

Create a protection group for a snapshot backup which identifies the client resources to backup.

Snapshot backups supports two types of groups:

- Basic client group—Defines a static list of Client resources to back up.
- Dynamic client group—Specifies a dynamic list of Client resources to back up. A dynamic client group automatically generates a list of Client resources that use client tag which matches the client tag that is specified for the group.

Creating a client group

Basic client groups define a static list of Client resources for a traditional backup, check connectivity, or probe action.

Before you begin

Create the Client resources for the data to include in a protection group.

Procedure

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

   The **Create Group** dialog box appears, starting with the **General** tab.
3. In the **Name** attribute, type a name for the group.

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

   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. Select the workflow in which to assign the group from the **Policy-Workflow** list.

   Note

   You can also assign the group to a workflow when you create or edit a workflow.
7. (Optional) On the **Restricted Datazones** tab, to specify the Restricted Datazone (RDZ) for the group, select the RDZ from the list.
8. Click **OK**.

After you finish

Create Client resources. The Client Configuration wizard and **General** tab on the **Client Properties** dialog box properties page provide you with the ability to assign clients to a protection group.

Creating a dynamic client group

Dynamic client groups prevent you from having to edit 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 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 attribute value.

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

Procedure

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

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

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

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

   Note

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

4. From the **Group Type** list, select **Dynamic Clients**.
5. In the **Comment** field, type a description of the group.
6. Select the workflow in which to assign the group from the **Policy-Workflow** list.

   Note

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

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

After you finish

Create Client resources. The Client Configuration wizard and **General** tab on the **Client Properties** dialog box properties page provide you with the ability to assign clients to a protection group and define one or more tags.

Creating a save set group

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

Before you begin

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

Procedure

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

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

3. In the **Name** box, type a name for the group.
4. From the **Group Type** list, select **Save Set ID List**.
5. In the **Comment** box, type a description of the group.
6. (Optional) To associate the group with a workflow, from the **Workflow (Policy)** list, select the workflow .

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

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

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

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

Creating a query group

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

Procedure

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

   The **Create Group** dialog box appears, starting with the **General** tab.
3. In the **Name** box, type a name for the group.
4. From the **Group Type** list, select **Save Set Query**.
5. In the **Comment** box, type a description of the group.
6. (Optional) To associate the group with a workflow, from the **Workflow (Policy)** list, select the workflow.

   You can also assign the group to a workflow when you create or edit a workflow.
7. Specify one or more of the save set criteria in the following table.

   Note

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

Table 3 Save set criteria

Criteria	Description
Date and time range	Specify the start date and time range for the save sets. To specify the current date and time as the end date for the range, select **Up to now**. To specify a different date and time as the end date for the range, select **Up to**, and then select the date and time from the lists.
Backup level	In the **Filter save sets by level** section, next to the backup level for the save set, select the checkbox: • **full** • **cumulative incr** • **logs** • **incremental** • **manual**
Limit the number of clones	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. **Note** The default is set to 0, and cannot be changed for NAS or Block.

Table 3 Save set criteria (continued)

Criteria	Description
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	Next to the media pool on which the save set is stored in the **Pools** list, select the checkbox. **Note** You cannot select Pools for NAS.
Name	In the **Filter save sets by name** box, specify the name of the save set. **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

Snapshot workflow supports the following actions types:

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

Only a backup action can follow a probe action.

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

Snapshot backup
A shapshot backup action performs a snapshot of data on supported snapshot hosts.

Clone

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

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

Supported actions in snapshot workflows

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

The following sections provide details on supported actions that can follow the lead action and other actions in a workflow.

All possible workflow actions for a snapshot backup

The following figure illustrates all the possible workflow actions that are associated with a snapshot backup.

Figure 9 All possible workflow actions for a snapshot backup



Note

You can configure the 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 on a per client basis. 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 and Clone actions, then after the client1 backup completes, NetWorker will clone 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 connectivity between clients and the NetWorker server, usually before another action such as a backup occurs.

Before you begin

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

Procedure

1. In the expanded left pane, select the workflow, and then perform one of the following tasks in the right pane to start the Policy Action wizard:
   - If this 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 **Specify the Action Information** page appears.
2. In the **Name** field, type the name of the action.

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

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, any action that occurs after a disabled action will not start, even if the succeeding options are enabled.

5. From the **Action Type** list, select **Check Connectivity**.
6. When you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is grayed out.
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, select the action that should precede this action from the **Previous** box.
   - If the action should run concurrently with an action, select the concurrent action from the **Previous** box, and then select the **Concurrent** checkbox.
8. Select whether to use 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. Click the icon on each day to specify whether to check connectivity with the client.

   The following table provides details on the icons.

   Table 4 Schedule icons

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

   To check connectivity every day, select **Execute** from the list, and then click **Make All**.

10. Click **Next**.

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

    The **Specify the Advanced Options** page appears.
13. (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.
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 only 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, specify the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no limit.
17. From the **Hard Limit** list, specify the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no limit.
18. (Optional) Configure overrides for the task that is scheduled on a specific day.

    To change the month on which to schedule the override, use the navigation buttons and the month list box. To change 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, 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**, then use the drop downs. Click **Add Rules based override**.
      - To perform the action task on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.
      - In the **Override** field, type an override.

        Note

        To remove an override, delete the entry from the **Override** field.
19. Click **Next**.

    The **Action Configuration Summary** page appears.
20. Review the settings that you specified for the action, and then click **Configure**.

After you finish

(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 clients that use the probe. Create a client Probe resource on the NetWorker server, and then 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 workflow, and then perform one of the following tasks in the right pane to start the Policy Action wizard:
   - If this 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 **Specify the Action Information** page appears.
2. In the **Name** field, type the name of the action.

   The maximum number of characters for the action name is 64.
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, any action that occurs after a disabled action will not start, even if the succeeding options are enabled.

5. From the **Action Type** list, select **Probe**.
6. When you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is grayed out.

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, select the action that should precede this action from the **Previous** box.
   - If the action should run concurrently with an action, select the concurrent action from the **Previous** box, and then select the **Concurrent** checkbox.
8. Select whether to use 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. Click the icon on each day to specify whether to probe the client.

   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.

   To perform a probe every day, select **Execute** from the list, and then click **Make All**.

10. Click **Next**.

    The **Specify the Probe Options** page appears.
11. Choose whether to start the subsequent backup action only after all probes succeed by selecting or clearing the **Start backup only after all probes succeed** checkbox.
    - To start the backup only if all the probes associated with Client resources in the assigned group succeed, select the checkbox.
    - To start the backup if any one of the probes are associated with a Client resource in the assigned group succeed, clear the checkbox.
12. Click **Next**.

    The **Specify the Advanced Options** page appears.
13. In the **Retries** box, 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 will not retry a failed backup or probe action.

    Note

    The **Retries** option only 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.

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

Note

The **Retry Delay** option only 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 is allowed to fail to communicate back to the server.

    If the job fails to respond within the timeout value, the server considers the job a failure. If a job fails, NetWorker retries the job immediately. This ensures that no time is lost due to failures.

    Increase the timeout value if a backup consistently aborts due to inactivity. Inactivity timeouts may 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 only 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 **Parallelism** field, specify the maximum number of concurrent operations for the action.

    Note

    The **Parallelism** value should not exceed 25.

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 only 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. Leave the default selections for the Notification group box. NetWorker does not support notifications for probe actions and ignores the values that are defined in the attributes.
19. From the **Soft Limit** list, specify the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no limit.
20. From the **Hard Limit** list, specify the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no limit.

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

    To change the month on which to schedule the override, use the navigation buttons and the month list box. To change 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, 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**, then use the drop downs. Click **Add Rules based override**.
      - To perform the action task on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.
      - In the **Override** field, type an override.

        Note

        To remove an override, delete the entry from the **Override** field.

22. Click **Next**.

    The **Action Configuration Summary** page appears.

23. 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 workflow, and then perform one of the following tasks in the right pane to start the Policy Action wizard:
   - If this 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 **Specify the Action Information** page appears.
2. From the **Action Type** list, select **Backup**.
3. From the secondary action list, select **Snapshot**.
4. When you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is grayed out.
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, select the action that should precede this action from the **Previous** box.
   - If the action should run concurrently with an action, select the concurrent action from the **Previous** box, and then select the **Concurrent** checkbox.

6. Select whether to use 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. Click the icon on each day to specify the type of backup to perform.

   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.
	Incr	Perform an incremental backup on this day. Incremental backups include files that have changed since the last backup of any type (full or incremental).
	Cumulative Incr	Perform a cumulative incremental backup. Cumulative incremental backups include files that have changed since the last full backup.
	Logs Only	Perform a backup of only database transaction logs. **Note** This is not applicable for NAS.
	Synthetic Full	Perform a synthetic backup on this day. A synthetic full backup includes all data that changed since the last full backup and subsequent incrementals to create a synthetic full backup. **Note** Synthetic Full backups are not supported.
	Skip	Do not perform a backup on this day.

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

   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.

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.

    When the minimum retention time passes, a snapshot action that is in progress can remove the snapshot from the storage device to ensure that there is sufficient disk space for the new snapshot.

13. Click **Next**.

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

14. In the **Retries** box, 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 will not retry a failed backup or probe action.

    Note

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

    Note

    The **Retry Delay** option only 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 is allowed to fail to communicate back to the server.

    If the job fails to respond within the timeout value, the server considers the job a failure. If a job fails, NetWorker retries the job immediately. This ensures that no time is lost due to failures.

    Increase the timeout value if a backup consistently aborts due to inactivity. Inactivity timeouts may 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 only 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.

17. In the **Parallelism** field, specify the maximum number of concurrent operations for the action.

    Note

    The **Parallelism** value should not exceed 25.

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

    The default notification action is to log the information to the `policy_notifications.log` file. The `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows, by default. You can use the `smtpmail` application on Windows or the default mailer program on Linux to send email messages.

    For example:

    - To log notifications to a file named `policy_notifications.log`, type the following command:

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

    - On Linux, to send a notification email, 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, specify the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no limit.
22. From the **Hard Limit** list, specify the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no limit.
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 workflow, and then perform one of the following tasks in the right pane to start the Policy Action wizard:
   - If this 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 **Specify the Action Information** page appears.
2. In the **Name** field, type the name of the action.

   The maximum number of characters for the action name is 64.
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, any action that occurs after a disabled action will not start, even if the succeeding options are enabled.
5. From the **Action Type** list, select **Clone**.
6. When you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is grayed out.
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, select the action that should precede this action from the **Previous** box.
   - If the action should run concurrently with an action, select the concurrent action from the **Previous** box, and then select the **Concurrent** checkbox.
8. Select whether to use 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. Click the icon on each day to specify whether to perform cloning.

   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.

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

10. Click **Next**.

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

11. In the **Data Movement** group box, define the volumes and devices to which NetWorker sends the clone 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**, select the option to instruct NetWorker to remove the source save set information from the client file index, and to mark the save set as recyclable in the media database during a Server expiration maintenance action. Clear this option to allow the source save sets to expire based on the defined retention time.

    c. From the **Destination Pool** list, select the target media pool for the cloned save sets.

    Note

    When you clone SnapVX snapshots or ProtectPoint snapshots, you can clone to media devices, for example Data Domain or to a ProtectPoint device. If you select a pool that contains ProtectPoint devices and media devices, then NetWorker clones the SnapVX and ProtectPoint save sets to the ProtectPoint devices. This is not applicable for NAS devices.

    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** group box, 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—Use the **Time** section to define the time range in which NetWorker should inspect, when searching for eligible save sets to clone in the media database. Use the spin boxes to specify the start of the time range and the end of the time range. The **Time** filter list includes three options, which 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 filter criteria.

- **Accept**—The clone save set list includes save sets whose save time is within the time range that is specified by the spin boxes and meet all the other defined filter criteria.
- **Reject**—The clone save set list does not include save sets whose save time is within the time range that is specified by the spin boxes and meet all the other defined filter criteria.

b. Save Set filter—Use the **Save Set** section to instruct NetWorker to include or exclude ProtectPoint and Snapshot save sets, when searching for eligible save sets to clone in the media database. The **Save Set** filter list includes three options, which define how NetWorker determines save set eligibility, based on the save set criteria:

- **Do Not Filter**—NetWorker inspects the save sets in the media database to create a clone save set list that meets the filter criteria.
- **Accept**—The clone save set list includes eligible ProtectPoint or Snapshot save sets, when you also enable the ProtectPoint or Snapshot checkboxes.
- **Reject**—The clone save set list does not include eligible ProtectPoint and Snapshot save sets when you also enable the ProtectPoint and Snapshot checkboxes.

c. Clients filter—Use the **Client** section to define a list of clients to include or exclude, when NetWorker searches for eligible save sets to clone in the media database. The **Client** list includes three options, which define how NetWorker determines save set eligibility, based on the client criteria:

- **Do Not Filter**—NetWorker inspects save sets that are associated with the clients in the media database, to create a clone save set list that meets the 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—Use the **Levels** section to define 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 three options, which define how NetWorker determines save set eligibility, based on the level criteria:

- **Do Not Filter**—NetWorker inspects save sets regardless of level in the media database, to create a clone save set list that meets all the 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.

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.

    Note

    The **Parallelism** value should not exceed 25.

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 only 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:
    - Select **Set at policy level** to use the notification configuration that is defined in the Policy resource to send the notification.
    - Select **On Completion** to send a notification on completion of the action.
    - Select **On Failure** to send a notification only if the action fails to complete.

18. In the **Send notification** attribute when you select the **On Completion** or **On failure** option, the **Command** box appears. Use this box to configure how NetWorker sends the notifications. You can use the `nsrlog` action to write the notifications to a log file or configure an email notification.

    The default notification action is to log the information to the `policy_notifications.log` file. The `policy_notifications.log` file is located in the `/nsr/logs` directory on Linux and the `C:\Program Files\EMC NetWorker\nsr\logs` folder on Windows, by default. You can use the `smtpmail` application on Windows or the default mailer program on Linux to send email messages.

    For example:

    - To log notifications to a file named `policy_notifications.log`, type the following command:

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

    - On Linux, to send a notification email, 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.

19. From the **Soft Limit** list, specify the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no limit.
20. From the **Hard Limit** list, specify the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no limit.
21. (Optional) Configure overrides for the task that is scheduled on a specific day.

    To change the month on which to schedule the override, use the navigation buttons and the month list box. To change 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, 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**, then use the drop downs. Click **Add Rules based override**.
      - To perform the action task on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.
      - In the **Override** field, type an override.

        Note

        To remove an override, delete the entry from the **Override** field.

22. Click **Next**.

    The **Action Configuration Summary** page appears.
23. 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

When you create actions for a workflow, a map provides a visual representation of the actions in the second right pane of the **Protection** window of the **Administration** interface.

Figure 10 Sample snapshot workflow

The oval icon at the beginning of the visual representation specifies the group to which the workflow applies, the rounded rectangle icons identify actions, and the parallelogram icons identify the destination pool for the action.

- Adjust the display of the visual representation by right-clicking and selecting one of the following options:
  - **Zoom In**—Use to increase the size of the visual representation.
  - **Zoom Out**—Use to decrease the size of the visual representation.
  - **Zoom Area**—Use to limit the display to a single section of the visual representation.
  - **Fit Content**—Use to fit the visual representation to the window area.
  - **Reset**—Use to reset the visual representation to the default settings.
  - **Overview**—To view a separate dialog box with a high-level view of the visual representation and a legend of the icons.
- View and edit the properties for the group, action, or destination pool by right-clicking the icon for the item and selecting **Properties**.
- Create a group, action, or destination pool by right-clicking the icon for the item and selecting **New**.

CHAPTER 3

Software Configuration

This chapter includes the following topics:

- Backup group resource migration 58
- Roadmap for snapshot configurations 60
- Snapshot configuration prerequisites 60
- Configuring the user privileges 63
- Configuring snapshot backups with the client wizard 64
- Configuring snapshot backups manually 68
- Configuring the Application Information variables 70
- Configuring preprocessing and postprocessing scripts 71

Backup group resource migration

During the migration process, NetWorker creates resources to replace each Group resource, and then migrates configuration attributes from the 8.2.x and earlier resources to the new NetWorker 9.0.1 resources.

Resource migration for group resources when Snapshot is enabled

This section summarizes the Group attribute values that migrate to 9.0.1 resources attributes, when the group is Snapshot enabled.

Table 8 Migration of Group attributes

9.0.1 Resource type	9.0.1 Resource name	Migration process overview Attribute values that are migrated from Group resource	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 backup files ystem and NMM data.	Not applicable
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-9.0.1 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.	• Autostart • Start Time • Next Start • Interval • Restart Window • End Time attribute value is set to Start Time+(Interval*(n-1)) When the Probe backup group attribute was enabled, the following values are also migrated: • Probe Interval—To the Interval attribute • Probe Start Time—To the Start Time attribute • Probe End Time—To the End Time attribute

Table 8 Migration of Group attributes (continued)

9.0.1 Resource type	9.0.1 Resource name	Migration process overview Attribute values that are migrated from Group resource	Attribute values that are migrated from Group resource
Probe	Probe	The Probe action resource appears when the Probe based group attribute was enabled in the pre-9.0.1 migrated group.	Not applicable
Action—Snapshot backup	Backup	The Snapshot Backup action appears for a Group resource that has the Snapshot attribute enabled.	• Parallelism • Retries • Retry delay • Success Threshold • Option attributes: ▪ 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 *EMC 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 and earlier snapshot policy resource that was configured for **Every nth snapshot** value is not migrated to NetWorker 9.0.1.
3. A NetWorker 8.2.x and earlier snapshot policy resource that was configured for **nth snapshot** is not migrated to NetWorker 9.0.1, and a clone action is not created.

   For example, a NetWorker 8.2.2 snapshot policy of 6-6-Week-3 which resolves to take 6 snapshots, retain 6, expire snapshots every week, and backup the first and third snapshot. The last value **Backup Snapshots** or **3** will not be migrated. Instead, a clone action is created to backup every snapshot.
4. Each backup action under snapshot policy is followed by a clone action. Clone actions have filtering options. Check whether filtering meets the backup requirements of snapshot management.

Roadmap for snapshot configurations

Plan NetWorker snapshot configurations with a high-level road map that outlines the sequence of configuration tasks that you need to perform.

1. Ensure that you meet the configuration prerequisites.
   Snapshot configuration prerequisites on page 60 provides details.
2. Configure the user privileges on the application host and the storage node.
   Configuring the user privileges on page 63 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 on page 64
   - Configuring snapshot backups manually on page 68
4. Configure any necessary Application Information variables.
   Configuring the Application Information variables on page 70 provides details.
5. Configure any necessary preprocessing or postprocessing scripts.
   Configuring preprocessing and postprocessing scripts on page 71 provides details.
6. Based on the array or appliance that you will use for snapshot backups, follow the configuration instructions in the appropriate chapter:
   - Configuring snapshots on VMAX Storage Arrays on page 99
   - Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard on page 78
   - Configuring snapshots on VNX Block Storage Arrays on page 107
   - Configuring snapshots on RecoverPoint on page 111
   - Configuring snapshots in a Cluster Environment on page 115
7. Test the configuration.

Snapshot configuration prerequisites

Verify the basic compatibility of all systems used for NetWorker snapshot operations. Components of the snapshot environment on page 22 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:
  - EMC Solutions Enabler on the application host and (if cloning snapshots) the NetWorker Server.
- VNX Block storage array:
  - EMC Unisphere® host agent, also known as EMC 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 if the LUNs are visible on the application host or the mount host.
  - EMC 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.
  - EMC 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.
  - EMC 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.
  - EMC Naviseccli or UEMCLI on the application host and the mount host.—This CLI is responsible for LUN discover and snapshot sync/split operations. The *EMC NetWorker Online Software Compatibility Guide* provides more details.
  - EMC UEMCLI on the mount host and application host.—This CLI is used for later VNX software releases and replaces other CLI. Check the EMC array specific requirements.
- RecoverPoint appliance:
  - 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 EMC Online Support, provides details on system and software requirements.

Application host prerequisites

Ensure that you meet the following prerequisites before you configure the application host for snapshot operations:

- The application host has installed the supported NetWorker client and extended client software.
- The application host has the support NMDA/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.
- You have synchronized the application host system clock with the mount host system clock.

- You have 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, XtremIO only support RDM via iSCSI.
  - A RecoverPoint, XtremIO and VNX appliance also requires 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 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.

Ensure that you perform the following prerequisites before you configure the mount host for snapshot operations:

- 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 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 requires 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-Z). NSM uses a drive letter when it mounts target volumes for restore 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 from 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 9.0.1 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 Group 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, click **User Groups** in the resources tree.
3. In the **User Groups** table, right-click the group that you want to modify and select **Properties**.
4. In the **Users** attribute, type the following information to specify the user as root, administrator, or System on the application host and the mount host:
   - 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 you meet the necessary prerequisites. Snapshot configuration prerequisites on page 60 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, right-click **Clients** in the navigation tree, and select **New Client Wizard**.
3. Complete the **Specify the Client Information** wizard page:
   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. If you want to add the current client to an existing group of clients that use the same workflow of snapshot, clone, or backup actions, then select the group from the **Group** list.

      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. Complete the information on the **Select the Backup Configuration Type** page:
   a. In the **Available Applications** table, select **Filesystem** or the supported NetWorker application type that is installed on the client.
   b. Select the checkbox that is labeled **Enable NetWorker Snapshot Management on the selected application**.
   c. Click **Next**.
5. Complete the **Select the Snapshot Management Options** page:
   a. Select the type of **Storage Array** or storage appliance that the client uses for primary storage and where the snapshots are created. Only arrays that are compatible with NetWorker and the client operating system display. The options are:

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

b. If you want to use the application host as the snapshot mount host, select **Use the current client as the mount host**.

Provided the selected array can use 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 on page 138 provides details.

---

c. Select the hostname of the NetWorker storage node from the drop-down list.

d. Click **Next**.

The storage array you selected will determine the next wizard screen.

6. If you selected the **EMC VMAX/Symmetrix** storage array option, the **Select the VMAX Mirror Policy** page displays. Perform the following steps:

a. Select the **VMAX Mirror Policy**:

- Select **Local Operation** if both the source device and the target mirror devices reside on the same VMAX array.
- Select **Remote (SRDF) Operation** if the source device and the target mirror devices reside on different VMAX arrays connected by a Symmetrix Remote Data Facility (SRDF).

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 using TimeFinder clone functionality, select **CLONE**.
- To create snapshots using TimeFinder/Snap (COW) functionality, select **VDEV**.

Note

VMAX version 3 and later does not support VDEV functionality.

- To create snapshots using TimeFinder VP Snap functionality, select **VP Snap**.
- To create snapshots using 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 100 provides details on mirror operations.

c. Click **Next** and you are taken to the **Select the NetWorker Client Properties** page. Go to Step 12.

7. If you selected the **EMC ProtectPoint for VMAX3** storage array option, the **Select ProtectPoint Destination** page displays. Perform the following step:

   a. In the **ProtectPoint Remote Operation** area:

      - Select **Local Operation** if both the source device and the target mirror devices reside on the same ProtectPoint array.
      - Select **Remote (SRDF) Operation** if the source device and the target mirror devices reside on different ProtectPoint arrays connected by a Symmetrix Remote Data Facility (SRDF).

   b. Click **Next** and you are taken to the **Select the NetWorker Client Properties** page. Go to Step 12.

8. If you selected the **EMC ProtectPoint for RecoverPoint** storage array option, the **Specify the RecoverPoint Storage Array Options** page displays. Perform the following step:

   a. Select or enter the **RecoverPoint Applicance Hostname / IP** from the drop-down list 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** and you are taken to the **Select the NetWorker Client Properties** page. Go to Step 12.

9. If you selected the **EMC 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 SnapView copy-on-write functionality, select **Copy on Write (COW/Snapshot)**.
      - To create the snapshots by using SnapView clone functionality, select **Mirror (SnapView Clone)**.
      - To create the snapshots by using Redirect on Write (ROW) functionality, select **VNX-SNAP (VNX Snapshot)** or **VNXe-SNAP (VNXe/VNXe2 Snapshot)**. Snapshot operations with SnapView software on page 108 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** and you are taken to the **Select the NetWorker Client Properties** page. Go to Step 12.

10. If you selected the **RecoverPoint** storage array option, the **Specify the RecoverPoint replication type and Storage Array Options** page displays. Perform the following steps:

a. Select the **Replication Type**:

- **CDP** (Continuous Data Protection)
- **CRR** (Continuous Remote Replication)

Snapshot operations with RecoverPoint software on page 112 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** and you are taken to the **Select the NetWorker Client Properties** page. Go to Step 12.

11. If you selected the **EMC XtremIO** storage array option, the **Specify the XtremIO storage process's credentials** page displays. Perform the following steps:

    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, the **Select the NetWorker Client Properties** page displays. Perform the following steps:

    a. Optionally, 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. If this process is not available on the client, then the storage node performs the clone operation.

    b. Click **Next**.

13. For the file system, complete the **Select Snapable Filesystem Objects** page:

Note

If you are using NMDA/NMSAP, instead of this page, there are application specific wizard pages displayed. Go through those 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 there are no compatible file systems, then an error appears.

On UNIX systems, the browse tree displays only those file systems 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, 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 could take some time.

c. Click **Next**.

14. Complete the **Client Configuration Summary** page:

a. Review the attributes and values listed in the summary.

To modify a setting, click **Back** or click the link in the step panel, and then make the revisions.

b. Optionally, 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.

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 right-click the Client resource and select **Check Snapshot Configuration**.

c. To accept and create the configuration, click **Create**.

15. Complete the **Check Results** page:

a. Ensure that the client configuration successfully completes.

b. Click **Finish**.

Configuring snapshot backups manually

EMC recommends that you use the 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 137.

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 by using the NMC Client Properties windows. You can manually modify a

RecoverPoint appliance or XtremIO configuration, but you cannot create the configuration with the new NMC Client Properties window, because the wizard is used to enter username and password information into the NetWorker server lockbox.

You can use the NMC Client Properties window to specify uncommon directives or options that the wizard interface does not support, such as the variables listed in Application Information Variables on page 137.

Procedure

1. Ensure that the necessary prerequisites are met.

   Application host prerequisites on page 61 provides details.

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

   Note

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

4. On the **General** tab, specify the following fields:
   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, ensure that you specify each item on a separate line with a fully qualified pathname. Pathnames are case sensitive.

      Note

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

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

5. Select 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 on page 137 provides details.

   Note

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

6. Put NMC in Diagnostic mode. Go to the **View** menu and select **Diagnostic mode**.

7. If you plan to clone snapshots to conventional storage media, select the **Globals (2 of 2)** tab and 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, right-click the Client resource in NMC and select **Check Snapshot Configuration**. The validation can take some time. A pop-up message lists detected problems.

   The *EMC 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 NSM with NMDA and NMSAP.

Configuring the Client resource manually for a mount host

If you did not use the Client Configuration Wizard but 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 the necessary prerequisites.

   Mount host prerequisites on page 62 provides details.

4. If you use the mount host only for snapshot restores and clones, then 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**, right-click the mount host in the right panel, and select **Modify Client Properties** to view its **Properties**.
   d. On the **General** tab, clear the selection for **Scheduled Backup**.
   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 by using the NetWorker Client Properties windows. Application Information Variables on page 137 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

EMC Technical Support does not support the content of the user defined scripts. You can write the scripts for a particular operation that you may must perform as long as the script returns the correct exit code to NSM.

Use the following guidelines for the scripts:

Procedure

1. Use preprocessing scripts and postprocessing scripts for operations such as application quiescing, shutdown, or startup. The scripts can produce output such as log files, but the scripts must return an exit status of 0, which means that the script did not fail and the backup can run. Any other exit code for a preprocessing script causes the backup to fail.
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**, and click **Advanced** on the **Security** tab.
   - 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. Also, 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 using 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 completes, verify the log files 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 section 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 74
- ProtectPoint on VMAX prerequisites 74
- Enabling vDisk on a Data Domain system 74
- Provisioning protection devices on Data Domain systems 75
- Completing the VMAX system configuration 76
- Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard 78
- Considerations for ProtectPoint device and NetWorker ProtectPoint enabled pools 80
- VMAX3 SRDF/S support 80
- Configuring Data Domain NsrSnapSG device groups for intelligent pairing 81

Overview

The EMC® ProtectPoint™ solution integrates primary storage on an EMC VMAX3 array and protection storage for backups on an EMC 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.

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. When you create a snapshot backup or clone action in a workflow, select the destination pool that you used when you labeled the device.

ProtectPoint on VMAX prerequisites

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

- Use a VMAX array.
- Ensure that the appropriate version of Solutions Enabler 8.2.0 or later is installed on the application host, mount host, storage node, mount host, and in some configurations on the NetWorker Server.

Configuring ProtectPoint

1. Enable vDisk on DD.
2. Provision protection vDisk devices on Data Domain.
3. Complete the configuration of the protection devices on the VMAX.
4. (Optional) Create restore vDisk devices on Data Domain if restoring directly from Data Domain.
5. Create a NetWorker device and label it under a NetWorker pool using the wizard.
6. Follow the steps in the Data Protection Policies on page 78 section to create snapshot policies, and then set the destination pool to a pool containing a ProtectPoint device.
7. Follow the steps in the Configuring snapshot backups with the client wizard on page 64 or Configuring snapshot backups manually on page 68 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 *EMC*

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. Run the following command to verify that the vDisk license is enabled:

   ```
   # license show
   ```

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

   If the DD vDisk license is disabled, run 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. Run the following command to enable the vDisk service subsystem:

   ```
   # vdisk enable
   ```

   ```
   DD VDISK enabled
   ```

4. Run the following command to verify that the vDisk service is enabled:

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

Procedure

1. Run the following vdisk commands to create the vdisk device pool and device group:

   `vdisk pool create pool_name user username`

   `vdisk device-group create device_group_name pool pool_name`
   The *EMC 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. Run the following vdisk command to create the backup and restore devices with the same geometry as the VMAX source LUNs, and provision the devices accordingly:

   `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 FTS software, and provision the restore LUNs to the restore host.

   The latest *EMC 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. Create a VMAX storage group named NSRSNAPSG for the restore devices. Or you can create your own VMAX storage group and specify the client resource with the NSM_SNAP_SG attribute. This group name is case-insensitive:

   ```
   # symsg -sid SymmID create NsrSnapSG
   ```

   ```
   # symsg -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 *EMC 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 00028:0029 -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 complete and check the status:

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

```
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
```

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

You can complete the following steps to create or modify a ProtectPoint device with the Device Configuration Wizard:

Procedure

1. In the NMC **Enterprise** view, select the NetWorker server name and 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 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 may vary according to the type of configuration chosen.

4. On the **Select the Device Type** page, select **ProtectPoint** device type, and click **Next**.
5. On the **Data Domain Preconfiguration Checklist** page, review the requirements.

   Configure the Data Domain system for ProtectPoint, and 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 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 signifies a vdisk pool on the Data Domain.

   You see the table populated with the NetWorker Device Name and Disk Pool information. Click **Next**.

8. On the **Configure Pool Information** page, specify the following settings:

   - Select the **Configure Media Pools for Devices** option.
   - Specify the **Pool Type** that targets clients to the devices. **Backup** for backups or **Backup Clone** for cloning or staging operations, or create a pool.

Note

If you have 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.

- Select the **Label and Mount device after creation** option.

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 you can mark it as a **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 the **Gather Usage Information** selection.

    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 launcher link for the Data Domain interface.

11. Review the **Review the Device Configuration Settings** page and then click **Configure**.

    NetWorker configures, mounts, and labels the ProtectPoint device for the specified pool.

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

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 non-ProtectPoint devices.
- NetWorker preserves the Data Domain credentials (hostname, username, and password), which were provided during the ProtectPoint device creation wizard in the device resource. NSM retrieves and uses this configuration information for all Data Domain communication.
- NetWorker only supports one mounted Protection Point device in a media pool.
- When the source snapshot is a ProtectPoint snapshot, then one of three clone operations that are listed 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 type of clone 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 performs a clone to an alternate media device.
  - Data Domain to Data Domain (CCR/MFR) Replication—This type of clone operation occurs when the source snapshot is a ProtectPoint snapshot and 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.

VMAX3 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 ProtectPoint on a remote VMAX array with SRDF/S functionality, 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.

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

- NetWorker automatically determines the state of the SRDF/S link at runtime. NetWorker does not require you to manually configure environment variables or application variables.

- If there is no SRDF/S link at the beginning of an operation, then 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 RDF 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 81 provides specific 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 procedure refers to the replication state, not the physical connection between the R1 and R2 devices on the separate VMAX arrays.

1. Transitions the link to split between the R1 and R2 devices.
2. Rolls back the data from the DD restore device to the R2 device.
3. Synchronizes the R2 device to the R1 device by using reverse synchronization.
4. Transitions the link to the synchronized state.
5. Leaves the RDF link in a synchronized state, after the rollback completes.

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 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 diskgroup 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 section includes the following topics:

- Overview 84
- Basic backup workflow 84
- Basic restore workflow 84
- ProtectPoint for RecoverPoint on XtremIO prerequisites 85
- Enabling vdisk on the Data Domain 86
- Provisioning protection devices on Data Domain systems 87
- Provisioning protection devices on Data Domain systems 88
- Configuring RecoverPoint and XtremIO storage 89
- Configuring NetWorker ProtectPoint, RecoverPoint and VMAX devices and pool with the wizard 90
- Considerations for ProtectPoint device and NetWorker ProtectPoint enabled pools 92
- Configuring Data Domain NsrSnapSG device groups for intelligent pairing 92

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 of 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 copies the changed data to a Data Domain storage device.
4. The Data Domain system creates and 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 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.
- 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 intended for XtremIO connectivity and all the XtremIO ports 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 than16 paths per device.
  - Enable MPIO if two or more paths are zoned to a Windows AR host.

Configuring ProtectPoint with NetWorker

1. Enable vDisk on DD.
2. Provision protection vDisk devices on Data Domain.
3. Create restore vDisk devices on Data Domain.
4. Create a NetWorker device and label it under a NetWorker pool using the wizard.
5. Create snapshot policies, follow the steps in the Data Protection Policies on page 78 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 64 or Configuring snapshot backups manually on page 68 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 *EMC 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. Run the following command to verify that the vDisk license is enabled:

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

   If the DD vDisk license is disabled, run 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. Run the following command to enable the vDisk service subsystem:

```
# vdisk enable
```

   The DD vDisk is enabled

4. Run the following command to verify that the vDisk service is enabled:

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

Procedure

1. Run the following vdisk commands to create the vdisk device pool and device group:

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

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

   The *EMC 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. Run the following vdisk command to create the backup and restore devices with the same geometry as the XtremIO source LUNs, and provision the devices accordingly:

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

   Note

   By default NSM's Intelligent Pairing will look for a Data Domain device group of name NSRSnapSG. It's 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. 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. Add XtremIO initiators to the group:

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

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

Procedure

1. Run the following vdisk commands to create the vdisk device pool and device group:

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

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

The *EMC 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. Run the following vdisk command to create the backup and restore devices with the same geometry as the VMAX source LUNs, and provision the devices accordingly:

```
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
   ```

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 make the Volumes available to the AR host.
2. Complete the following tasks on the RecoverPoint cluster.
   a. Register the XtremIO Management Server on the RecoverPoint cluster.
   b. Register the Data Domain system on the RecoverPoint cluster.
   c. Create a RecoverPoint consistency group that contains all the XtremIO Volumes created for the ProtectPoint environment.
3. Complete the following tasks on the Data Domain system:
   a. Login to an SSH session on the Data Domain system.
   b. Enable block services for ProtectPoint if it is not already enabled.
   c. Enable DD Boost if it is not already enabled.
   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, create a storage unit on the ProtectPoint MTree.

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 required for ProtectPoint with RecoverPoint backups.

g. Create a block-services for ProtectPoint device-group for restores that resides in the restore pool 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 complete the following steps to create or modify a ProtectPoint device with the Device Configuration Wizard:

Procedure

1. In the NMC **Enterprise** view, select the NetWorker server name and 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 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 may vary according to the type of configuration chosen.

4. On the **Select the Device Type** page, select **ProtectPoint** device type, and click **Next**.
5. On the **Data Domain Preconfiguration Checklist** page, review the requirements.

   Configure the Data Domain system for ProtectPoint, and 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 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 signifies a vdisk pool on the Data Domain.

   You see the table populated with the NetWorker Device Name and Disk Pool information. Click **Next**.

8. On the **Configure Pool Information** page, specify the following settings:

   - Select the **Configure Media Pools for Devices** option.
   - Specify the **Pool Type** that targets clients to the devices. **Backup** for backups or **Backup Clone** for cloning or staging operations, or create a pool.

     Note

     If you have 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.

   - Select the **Label and Mount device after creation** option.

   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 you can mark it as a **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 the **Gather Usage Information** selection.

    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 launcher link for the Data Domain interface.

11. Review the **Review the Device Configuration Settings** page and then click **Configure**.

    NetWorker configures, mounts, and labels the ProtectPoint device for the specified pool.

12. The final **Device Configuration Results** page is informational only. To exit the wizard, click **Finish**.

13. In the NMC **Devices** view, verify that NetWorker has labeled and mounted the device, ready for use. This view also lists the volume name for the device.

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 non-ProtectPoint devices.
- NetWorker preserves the Data Domain credentials (hostname, username, and password), which were provided during the ProtectPoint device creation wizard in the device resource. NSM retrieves and uses this configuration information for all Data Domain communication.
- NetWorker only supports one mounted Protection Point device in a media pool.
- When the source snapshot is a ProtectPoint snapshot, then one of three clone operations that are listed 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 type of clone 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 performs a clone to an alternate media device.
  - Data Domain to Data Domain (CCR/MFR) Replication—This type of clone operation occurs when the source snapshot is a ProtectPoint snapshot and 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 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 diskgroup 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 section includes the following topics:

- Snapshot support for XtremIO 96
- Snapshot operation with XtremIO REST API 96
- Prerequisite for XtremIO configurations 96
- Supported XtremIO features 96
- Snapshot management policy with XtremIO 96
- Snapshot backups with XtremIO 96
- XtremIO configuration methods 97

Snapshot support for XtremIO

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

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 NSM prerequisites for XtremIO array are as follows:

- XtremIO Storage Array version 4.0.1 or higher is required for RESTfull 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 supporter

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 RecoverPoint.

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 imitator 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.

CHAPTER 7

Configuring snapshots on VMAX Storage Arrays

This chapter includes the following topics:

- Snapshot support of VMAX storage arrays 100
- Pairing source LUNs to mirror LUNs 101
- VMAX SRDF/S support 104
- Solutions Enabler Client and Server mode configuration 105
- Known limitation for VMAX 106

Snapshot support of VMAX storage arrays

This chapter describes NetWorker snapshot support, practices, and configurations that are specific to VMAX (Symmetrix) storage arrays.

Migrating Legacy PowerSnap Configurations on page 151 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 will manage all NetWorker TimeFinder client and server operations.

The TimeFinder software maintains multiple, host-independent copies of production data by 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 EMC Solutions Enabler is installed on the application host and the mount host. Storage array prerequisites on page 61 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 105 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 created with the application variable SYMM_CLONE_FULL_COPY=FALSE.
  - Rollback operations fail to 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 on page 139 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 source and target pairs.
  - When you perform a rollback with SnapVX snapshots, the snapshot will not be deleted from the NetWorker media database and be 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: **symsg -sid *vmax_id* create NsrSnapSG**

To add a device to this group that NetWorker can use as a mirror, run the following command: **symsg -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 will create 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, then 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, then 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 81 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 procedure refers to the replication state, not the physical connection between the R1 and R2 devices on the separate VMAX arrays.

1. Transitions the link to split between the R1 and R2 devices.
2. Rolls back the data from the DD restore device to the R2 device.
3. Synchronizes the R2 device to the R1 device by using reverse synchronization.
4. Transitions the link to the synchronized state.
5. Leaves the RDF link in a synchronized state, after the rollback completes.

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 *EMC Solutions Enabler Installation Guide* provides more information on SYMAPI Client/Server configuration. The following diagram displays a Symmetrix array in the client server system.

Configuring the Solutions Enabler in Client/Server mode on a SAP client

Procedure

- Modify the `netcnfg` file `/usr/emc/API/SYMAPI/config/netcfng`. Type the following line:

  ```
  Proxyservername - TCPIP Proxyservername ProxyserverIP 2707
  ANYNSM_SERVER - TCPIP Proxyservername ProxyserverIP 2707 ANY
  ```

- Type the following lines in the `/nsr/nsrrc` files:

  `SYMCLI_CONNECT = Proxyservername` (as defined in the `netcfng` file)

  ```
  SYMCLI_CONNECT_TYPE = REMOTE
  ```

  ```
  export SYMCLI_CONNECT
  ```

  ```
  export SYMCLI_CONNECT_TYPE
  ```

Configuring the Solutions Enabler in Client/Server mode on a Mount Host

Procedure

- Modify the `netcnfg` file `/usr/emc/API/SYMAPI/config/netcfng`.

  ```
  NSM_SERVER - TCPIP Proxyservername ProxyserverIP 2707 ANY
  ```

- Enable the Storsrvd service. Type the following command:

  ```
  stordaemon start storsrvd
  ```

  You must add the storsrvd service to the statup procedure of the mount host.

  During the Zoning and or Masking process, you must grant the mount host access to the luns of the application hosts. However, the luns must not be mounted, they must only visible in order for the snapshots to run.

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 108
- Configuring the Navisphere security file 108
- Configuring Unisphere CLI on VNXe3200 109

Snapshot support of VNX Block storage arrays

This chapter describes NetWorker snapshot support, practices, and configurations that are specific to VNX Block (CLARiiON) storage arrays.

Migrating Legacy PowerSnap Configurations on page 151 provides information on migrating PowerSnap VNX implementations to NetWorker snapshot management.

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 that NetWorker can recover and can clone to conventional storage media.

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

Prerequisites and support for VNX configurations

Install EMC Unisphere host agent and EMC SnapCli on the application host and the snapshot mount host. Storage Array specific prerequisites on page 61 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)

Configuring the Navisphere security file

The Navisphere security file is required on all nodes that participate in snapshot operations with VNX 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 will create 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, run 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_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 unzip it in a temporary folder.
3. Open a command prompt as an administrator and browse to the folder where you unzipped `PSEXEC.EXE`.
4. Run the following command:

   ```
   PSEXEC -i -s -d CMD
   ```

   A new command prompt appears.

5. Type the following command to verify that the command prompt belongs to the system user account:

   ```
   WHOAMI/USER
   ```

6. Run the following command to set the global account for all users on this system:

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

7. Run the following command to set credentials for a specific VNX storage array:

   ```
   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_frameIP 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 get Unisphere CLI running 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.

Procedure

1. Install Unisphere CLI on both the production and mount host.

   The Unisphere CLI product provides information for OS specific details and requirements.

2. On the Configure Optional Settings pane, to include Unisphere CLI in the environment path, leave the **Unisphere CLI in the Environment path** that is selected by default.
3. On the **Configure Optional Settings** pane, select the **Low (The certificate will not be verified)** security option to avoid runtime backup errors.
4. On both the client and proxy or storage node, create a security file in the command line by running the following command:

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

5. Test the connection by running 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

In order to setup UEMCLI, specify the Windows registry setting for the UEMCLI system path. It has been observed on occasion where the Unisphere CLI did not add the installed path to its registry entry. In order for NSM to work with UEMCLI, the full path to the binary is required. To ensure you have a valid path Unisphere install location, manually add a registry entry for NetWorker.

1. From the command prompt or a Windows shell prompt, type **`regedit`**.
2. Browse to the `HKEY_LOCAL_MACHINE\SOFTWARE\Legato\NetWorker` folder.
3. Right-click and select **New** › **String Value**.
4. Name the value **`UEMCLI_directory`**, and press **Enter** on the keyboard.
5. Right-click the UEMCLI_directory, and 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 while 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 112
- Supported RecoverPoint features 112
- RecoverPoint configuration methods 113
- RecoverPoint snapshot retention 114

Snapshot support of RecoverPoint

This chapter describes NetWorker snapshot support, practices, and configurations that are specific to RecoverPoint.

Migrating Legacy PowerSnap Configurations on page 151 provides information on migrating PowerSnap RecoverPoint implementations to NetWorker snapshot management.

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 from the snapshots or clone them to conventional storage media. For each snapshot operation, RecoverPoint records a SPIT bookmark of the snapshot.

The *EMC RecoverPoint Administration Guide* provides details.

Prerequisite for RecoverPoint configurations

The NSM prerequisite for RecoverPoint appliances is as follows:

- Configure RecoverPoint on a supported VMAX, VNX, XtremIO, or VPLEX storage array.

The *EMC NetWorker Online Software Compatibility Guide* provides details. See Storage Array specific prerequisites on page 61 for more details.

Restrictions for RecoverPoint configurations

The following restrictions apply to 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 may 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 two 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 content 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 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

Snapshot management policy

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 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 with permissions to create bookmarks.
- The mount host attached to the RecoverPoint replication storage.
- The backup option for the backup, which is either CDP, which informs NSM that the mount host has access to the Local Replication storage, or 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 NMC Client Configuration Wizard supports the creation of NSM RecoverPoint configurations, and this is the recommended method.

However, after you have created a RecoverPoint configuration, you can modify the configuration through the NMC property windows. This enables you to use Application Information variables, as listed in Application Information Variables on page 137.

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

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

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

RecoverPoint snapshot retention

You can determine the lifecycle of a RecoverPoint SPIT only by the RecoverPoint Copy Policy 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 can occur when the daily cleanup process has not yet removed the save set references.

When planning the management of RecoverPoint snapshot operations with NSM, consider the following potential problems:

- A high storage change rate can force a bookmark out of its Copy Policy retention period before a backup completes. This action can cause snapshot clones to fail. For example, RecoverPoint software can track database changes to a production LUN many times per 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.

  The 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 act of restoring data 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.
  The 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 116
- Configuring a cluster environment for snapshots 116
- AIX systems in a cluster environment 117
- ProtectPoint restore and rollback for VCS on Solaris 118

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 11 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 aborts the snapshot and cleans up to return the snapshot environment to the pre-snapshot state. After the failover completes, NetWorker retries the snapshot on the active cluster node.

If the application or the cluster node fails over during a recovery, NetWorker aborts 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. Install NetWorker client software in a local directory on each physical node of the cluster.

   The *EMC NetWorker Cluster Integration Guide* provides details.

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

- Microsoft Windows:

  NetWorker_install_path\lc_config

- UNIX:

  /NetWorker_install_path/networker.cluster

  The *EMC NetWorker Cluster Integration Guide* provides details.

3. Configure a NetWorker Client resource for each cluster virtual node that requires NetWorker snapshot services. 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 on page 138 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 *EMC 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

Follow the steps below to enable a ProtectPoint restore and rollback for VCS on Solaris.

Performing a ProtectPoint VCS Restore

Procedure

1. As the root user, perform the following steps on the primary VCS node:
   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. Enable the VCS configuration as Read/Write.

      root:/# haconf -makerw

   c. Freeze the VCS service groups by disabling On line/Off line. Type the following command:

      Note

      This is an example of a VCS and Oracle configuration.

      root:/# hagrp -freeze <oracle_sg> -persistent

      root:/# hagrp -freeze <oracle_ctl_sg> -persistent

   d. Confirm the VCS status, by typing 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 the oracle user, perform the following steps on the primary VCS node:

a. Run the `shutdown` and `startup mount` commands on the Oracle database.

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, perform the following steps on the primary VCS node:

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:/# hagrp -unfreeze <oracle_ctl_sg> -persistent

root:/# hagrp -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, perform the following steps on the primary VCS node:
   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. Enable the VCS configuration as Read/Write.

      root:/# haconf -makerw

   c. Freeze the VCS service groups by disabling On line/Off line. Type the following command:

      Note

      This is an example of a VCS and Oracle configuration.

      root:/# hagrp -freeze <oracle_sg> -persistent

      root:/# hagrp -freeze <oracle_ctl_sg> -persistent

   d. Confirm the VCS status, by typing 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, perform the following steps on the primary VCS Node:

a. Shutdown and startup mount the Oracle database:

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, perform the following steps on the primary VCS node:

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

root:/# haconf –makerw

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

root:/# hagrp -unfreeze <oracle_ctl_sg> -persistent

root:/# hagrp -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
```

Note

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

CHAPTER 11

Data Management and Recovery

This chapter includes the following topics:

- Snapshot lifecycle management....124
- Management and recovery of file system snapshot data....124
- Snapshot recovery support and limitations.... 124
- Restoring from a snapshot with the Recovery Wizard.... 125
- Restoring a snapshot by rollback....128

Snapshot lifecycle management

Snapshot-based operations are fully integrated into NetWorker policies.

The schedule of the workflow controls the number of snapshots taken per 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 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

The chapter describes management and recovery operations for file system data only. For information on database snapshots, such as NMDA or NMSAP database snapshots, refer to the documentation for the NetWorker application module that you are using.

The 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

The NetWorker client file index records 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 will mount 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 supports the following user interfaces:
  - NMC Recovery Wizard GUI
  - `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 *EMC NetWorker Administration Guide* provides details.

Note

Recovering from a snapshot being rolled over to a secondary media (Disk or Tape), is not possible. This results in both rollover and recovery failures.

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.

Restoring from a snapshot with the Recovery Wizard

You can use the NMC Recovery Wizard GUI 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 select **Enterprise** › **Launch Application**.
2. In the NetWorker server's **Recover** view, select **Recover** › **New Recover** to launch the Recovery Wizard.
3. Complete the **Select Client to Recover** page:
   a. In the **Source Client** field, specify the application host whose production data was the source for the snapshot you want to restore.
   b. In the **Destination Client** field, specify the application host or an alternative NetWorker client on which you want to restore the snapshot data.
   c. In the **Available Recover Types**, select **Snapshot Management**.
   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**

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 **NSR_DATA_MOVER Application Information** attribute on the **APPS and MODULES** tab in the properties of the source client.

  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 27 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.
  - **Browse and recover save set**. When you click **Next**, the wizard mounts the snapshot volume for the save set and open the **Select Data to Recover** page. The mount operation can take some time.
  - **Recover full save set**. When you click **Next**, the wizard goes to 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 goes to 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 128 provides details.

e. Click **Next**. The result depends on the recover or rollback selections.

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 on page 146 and the *EMC 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 specify the attributes.

   d. Click **Next**.

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

   a. In the **Recovery Name** field, type a name for the recovery.

   b. In the **Recovery Start Time** field, specify the following attributes:

      - **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 27.
- For a rollback recovery, the storage array's capabilities perform the restore. Restoring a snapshot by rollback on page 128 provides details.

The *EMC 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 an unmounted snapshot volume.

Rollback considerations

Always consider how a rollback can affect any other snapshots or other data on the storage array. You can perform a rollback only one time because the rollback deletes the snapshot and its save set record. The rollback could delete the original snapshot and all subsequent snapshots that existed with it before the rollback.

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 `psrollback.res`. 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 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 *EMC NetWorker Command Reference Guide* and man pages provide details on the `nsrsnapadmin` and `nsrsnap_recover` commands.

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. Some time 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 revert /fs2 and /fs3.

Configuring the `psrollback.res` file

The `psrollback.res` file is an NetWorker snapshot rollback resource file 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:
  - There must be 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

NSM supports rollbacks with AIX volume manager as follows:

- NSM supports rollbacks in a HACMP shared volume group environment, provided you set the Auto On function to No. This setting prevents AIX from automatically activating the volume group during a system startup.
- NSM does not support rollbacks in a 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 a HACMP shared volume group

After a rollback on AIX systems, AIX places HACMP shared volume group configurations into a nonsynchronized state by default.

To enable rollbacks that maintain a synchronized state:

Procedure

1. On the host where the cluster service is online, take the volume group offline with 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 with 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 134
- NetWorker snapshot restore issues 135

NetWorker snapshot backup issues

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.

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 EMC PowerPath® software is optional for NSM on VNX arrays. An incorrectly configured Linux Volume Manager (LVM) used with PowerPath can result in snapshot failures and 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 *EMC PowerPath for Linux Installation and Administration Guide* provides details.

Linux LVM snapshots fail with an error

Linux 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

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
As root, use the `chown` command to change the owner of the device pathname to the correct user.

Command nsrsnap_recover -I completes 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 DD is already successful.

Workaround 1
Reboot the host.

Workaround to eliminate a reboot of the host

1. Go to device manager and 'uninstall' *ALL* SYMMETRIX disks (not just the impacted ones).
2. Use UI Disk Management to rescan all devices.
3. Recycle the Virtual Disk service.

APPENDIX A

Application Information Variables

This appendix includes the following topics:

- Using Application Information variables 138
- Common Application Information variables 138
- Application Information variables for VMAX arrays 139
- Application Information variables for VNX Block arrays 141
- Application Information variables for RecoverPoint appliances 142
- Application Information variables for XtremIO arrays 143

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, type the variables and their values in the Application Information attribute of the Apps and Modules tab of the Client resource for the application host.

Configuring the Client resource manually for the application host on page 68 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 12 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. There is no default value.
NSR_PRE_SNAPSHOT_SCRIPT	Specifies the pathname of the preprocessing command script. There is 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. 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 per NSM backup. The default value is 16. 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. To turn off parallelism so that an NSM backup only creates a single save stream at a time, set the value to 1.
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.

Table 12 Common Application Information variables (continued)

Common variable	Definition
NSR_PS_SINGLE_LOG	Specifies whether NetWorker logs all NSM processes together in the `nwsnap.raw` file. Valid values are TRUE and FALSE. The default value is TRUE if NSR_PS_DEBUG_LEVEL is 3 or less. Set to FALSE to cause logging to individual process-based log files.
NSR_SNAP_TYPE	Specifies the snapshot provider. 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. If TRUE, NSM forces the lgtosync driver or equivalent OS-level capability to freeze and thaw writes to a disk or volume.

Application Information variables for VMAX arrays

The following table lists Application Information variables that NSM can use for VMAX storage arrays.

Note

None of the variables in the following table are relevant for ProtectPoint, except the ones that are specifically listed.

Table 13 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 `symm.res` file.
SYMM_CLONE_FULL_COPY	Valid values are TRUE and FALSE. The default value is TRUE. • TRUE—NSM performs a full data copy of a source LUN. • FALSE—NSM places the target in COW (CopyOnWrite) mode and will not perform a full data copy. Notes: • In a single backup/restore session, NSM can use a BCV as a mirror or a clone, but not both. • NSM does not allow a rollback operation for a snapshot taken when this variable is FALSE. An attempted rollback will fail.

Table 13 Application Information variables for VMAX arrays (continued)

VMAX variable	Definition
SYMM_EXISTING_PIT	Valid values are TRUE and FALSE. The default value is FALSE. Specifies the state of targets for the symm.res 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. The default value is TRUE. Used by intelligent pairing when NSM cannot use any of the mirrors currently paired with the source LUN. • 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. • If set to FALSE, intelligent pairing can select only an available mirror from the devices already paired to 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 mirrors with SYMM_CLONE_FULL_COPY=FALSE, NSM always terminates the relationship: • RETAIN_RESOURCE—NSM resynchronizes the mirror again with the source when it deletes the snapshot. • 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). • START_STATE—NSM leaves the target mirror in the same state (split or synced) as before the backup.
SYMM_RB_OVERRIDE_OTHER_TGTS	Valid values are TRUE and FALSE. The default value is FALSE. • FALSE—NSM fails the rollback if any other mirrors are in the synchronized state with the source device. • TRUE—Before a rollback operation, NSM splits all synchronized mirrors and then resynchronizes them on completion of a rollback. Notes • 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.

Table 13 Application Information variables for VMAX arrays (continued)

VMAX variable	Definition
	• 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.
SYMM_RES_USE_POLICY	The default value is ANY. • EXISTING—NSM seeks a resource that is synchronized with the source device. This setting reduces the backup time. • FREE—NSM seeks a resource that is not synchronized with any device. The resource must be in a split or not paired state. • ANY—NSM seeks any existing resource first. If NSM finds none, it uses a FREE resource.
SYMM_SNAP_POOL	The default value is `C:\programFiles\EMC networker \Nsr\res\symm.res`. Defines the pathname of the `symm.res` file.
SYMM_SNAP_REMOTE	Valid values are TRUE and FALSE. The default value is FALSE. Set to TRUE if using SRDF. Set to FALSE if not using SRDF.
SYMM_SNAP_TECH	Valid values are SNAP, BCV, CLONE, VPSNAP, R2, and SNAPVX. Defines the type of mirroring to use. • If set to R2, then SYMM_SNAP_REMOTE must be TRUE or the backup will fail.
NSR_PS_TERMINATE_SRC_MIRRORS	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, this will cause the backup to fail. Values are TRUE and FALSE. The default value is FALSE. Set this attribute to TRUE to make NSM to first terminate the bcv/clone/vpsnap mirror relationship with the source LUN, and then proceed to do a snapvx snapshot.
NSM_SNAP_SG	Valid value must be a valid VMax storage group.

Application Information variables for VNX Block arrays

The following table lists Application Information variables that NSM can use for VNX Block storage arrays.

Table 14 Application Information variables for VNX Block arrays

VNX Block variable	Definition
CLAR_ON_DELETE	Specifies the disposition of the clone LUN when NSM deletes a snapshot:

Table 14 Application Information variables for VNX Block arrays (continued)

VNX Block variable	Definition

	• RETAIN_RESOURCE—NSM resynchronizes the clone with its source. This makes the clone LUN available for future snapshot requests.
	• RELEASE_RESOURCE—NSM does not resynchronize the clone with its source. This makes the clone LUN available for other client operations, provided 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.
	• START_STATE—NSM resynchronizes the clone with its source LUN only if it was in a synchronized state when it was fractured. In this case, its disposition becomes one of the following:
	  ▪ – RETAIN_RESOURCE workflow
	  ▪ – RELEASE_RESOURCE workflow
	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.

EMCCLAR_SNAP_SUBTYPE	Mandatory. The default value is COW for copy-on-write backup and recovery workflows. • Set to MIRROR for clone fracture and for clone backup and clone recovery workflows. • Set to VNX-SNAP for VNX snap backup and VNX snap recovery workflows. • Set VNXe-SNAP for VNXe3200 backup and VNXe3200 snap recovery workflows.
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 NSM can use for RecoverPoint appliances.

Table 15 Application Information variables for RecoverPoint appliances

RecoverPoint variable	Definition
NSR_SNAP_TECH	Specifies the RecoverPoint replication type for a backup or restore. Use the value RP_CDP to notify NSM that local copies will be used to access a bookmark, or RP_CRR to notify NSM that remote copies will be used to access a bookmark.
RP_APPLIANCE_NAME	Specifies the hostname or IP address of the RecoverPoint appliance for NSM to use.

Table 15 Application Information variables for RecoverPoint appliances (continued)

Application Information variables for XtremIO arrays

The following table lists Application Information variables that NSM can use for XtremIO arrays.

Table 16 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 **`emc-xtremio`**.
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 146
- Using nsrsnapadmin for snapshot operations 146
- Example nsrsnapadmin operations 147
- Querying with the mminfo command 149

Using CLI commands for snapshot operations

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

The *EMC 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, type `nsrsnapadmin` at the CLI prompt. When you receive an input prompt, you can type a specific command and its available options to perform the NetWorker options listed in the following table.

Table 17 Commands and options supported in nsrsnapadmin interactive mode

NSM operation	Command and available options
Display snapshot save sets	**p [-s *nsr_server*] [-c *client*] [-v] [*path* | -S *ssid*]**
Delete a snapshot save set	**d [-s *nsr_server*] [-c *client*] [-v] [-a] [-y] -S *ssid* [or -S "*ssid ssid* ..."]**
Perform a save set recovery	**R [-s *nsr_server*] [-c *client*] [-M *mount_host*] [-v] -S *ssid* [-t *destination*] [-T *recover_host*] -m *path* [-A attr=*val*]**
Perform a file-by-file browsing and recovery	**r [-s *nsr_server*] [-c *client*] [-M *mount_host*] [-T *recover_host*] -S *ssid* [-A attr=*val*]**
Perform a rollback	**B [-s *nsr_server*] [-c *client*] [-M *mount_host*] [-Fv] -S *ssid* [-A attr=*val*] -m *path***
Reset the expiration time for a snapshot save set	**e *time* [-s *nsr_server*] [-c *client*] [-v] -S *ssid* [or -S "*ssid ssid* ..."]**
Exit the program	**q or quit**

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

Example nsrsnapadmin operations

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

Querying snapshot save sets

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

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

where:

- *server* is the hostname of the NetWorker server.
- *client* is the hostname of the 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)   ssname=/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 128 provides more information.

Deleting a snapshot save set

You can use the `nsrsnapadmin`, `nsrmm`, `nsrim`, or `nsrsnapck` command to delete snapshot save sets.

Deleting an 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

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

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

where:

- *time* is the date and time when the snapshot save set will expire.
  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 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 as in `-e 09/04/14`, the time defaults to 00:00:00. NSM changes the snapshot save set's browse and retention times to 09/04/14 00:00:00.

If you specify a time hh:mm:ss as in `-e 20:00:00`, the date defaults to the system time, for example, 09/03/14. NSM changes the snapshot save set's browse and retention times to 09/03/14 20:00:00.

Querying with the `mminfo` command

Another way to query a client's snapshot save sets is with 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, type the following at the command prompt:

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

where

- *server* is the hostname of the NetWorker server.
- *clientname* is the hostname of the 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 *EMC 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 152
- Deprecated Client resource attributes 153
- Migrating VMAX (Symmetrix) arrays 153
- Migrating VNX (CLARiiON) arrays 153
- Migrating RecoverPoint appliances 153
- Starting the nsrpsd process 154
- Licensing 154

Migrating legacy PowerSnap configurations to NSM

This appendix provides examples of how to migrate legacy NetWorker PowerSnap Module configurations to NetWorker snapshot management. The NetWorker client installation provides all the functionality 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 PowerSnap packages on Solaris:
  - LGTOpsag
  - LGTOpseg
  - LGTOpssc
- Remove the following packages on AIX:
  - LGTOps.psag.rte
  - LGTOps.pseg.rte
  - LGTOps.pssc.rte
- Remove the following packages on HP:
  - PowerSnap.pkg

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 NSM, remove any snapshots 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 supports NSM. Components of the NSM network on page 22 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 101 provides details on intelligent pairing.

Migrating VNX (CLARiiON) arrays

Before the upgrade to NSM, remove any snapshots 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 supports NSM. Components of the NSM network on page 22 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 on page 108 provides details.

Migrating RecoverPoint appliances

Ensure that the operating system, versions, and configuration supports NSM. Components of the NSM network on page 22 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.

Also, 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 credentials for username and password available for the RecoverPoint appliances that participate in snapshot backups.
2. Open an NMC session to the NetWorker server and 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 add the required attributes from the old configuration to the new one.
5. In the **Specify the RecoverPoint replication type and Storage Array Options** screen, add the RecoverPoint username and password to the lockbox on the NetWorker server.
6. When finished 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 `nsrpsd` does not automatically start or stop and backups of these applications will fail. You must start or stop `nsrpsd` manually, as done in previous PowerSnap releases. In these environments, `nsrpsd` will not self-terminate after 30 minutes of no activity.

Licensing

You do not require any new licenses. NSM will honor existing PowerSnap licenses and the NetWorker capacity and traditional licensing models.

NetWorker snapshot licensing requirements on page 24 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.

media Physical storage, such as a disk file system or magnetic tape, to which backup data is written. **See** volume.

media database Database that contains indexed entries of storage volume location and the life cycle status of all data and volumes managed by the NetWorker server.

mirror
1. Exact duplicate of data or another object, such as a disk.
2. To write duplicate data to more than one device.

mount To make a volume physically available for use, such as the placement of a removable disk volume or tape into a drive for reading or writing.

mount host Host in a network that is used to mount storage array snapshot volumes to perform snapshot restore and rollover operations.

N

NetWorker Management Console (NMC) Software program that is used to manage NetWorker servers and clients. The NMC server also provides reporting and monitoring capabilities for all NetWorker processes.

NetWorker server Computer on a network that runs the NetWorker server software, contains the online indexes, and provides backup and restore services to the clients and storage nodes on the same network.

NetWorker Snapshot Management (NSM) EMC technology that provides point-in-time snapshot copies of data. NetWorker software backs up data from the snapshot. This allows applications to continue to write data during the backup operation, and ensures that open files are not omitted.

P

point-in-time copy (PIT copy) Fully usable copy of a defined collection of data, such as a consistent file system, database, or volume that contains an image of the data as it appeared at a specific point in time. A PIT copy is also called a snapshot or shadow copy.

policy Set of defined rules for client backups that can be applied to multiple groups. Groups have dataset, schedule, browse, and retention policies.

pool
1. NetWorker sorting feature that assigns specific backup data to be stored on specified media volumes.
2. Collection of NetWorker backup volumes to which specific data has been backed up.

Q

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

R

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 location by using the capability of the storage array. A rollback restore destroys existing data on the source location.

S

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

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

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

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.