EMC® NetWorker®

Version 9.0.x

EMC Data Domain® Boost Integration Guide

302-001-769

REV 06

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 7

Tables 9

Preface 11

Chapter 1 DD Boost Features and Environment 15

DD Boost integration features 16
- DD Boost and data deduplication 16
- Client Direct data handling 16
- DD Boost device storage 16
- DD Boost Synthetic Full backups 17
- DD Boost in-flight encryption 17
- DD Boost clone operations 17
- NMC management of DD Boost operations 18
- Feature not supported by the integration 18

Data Domain network environment 18
- FC and IP network connectivity 18
- Data Domain storage system 18
- NetWorker client 19
- NetWorker server 19
- NMC server 19
- NetWorker storage node 19
- NetWorker application modules 19

Licensing in Data Domain systems 19
- Traditional licensing for Data Domain systems 20

Chapter 2 Planning and Practices 21

DD Boost storage characteristics 22
- DD Boost storage structures and limits 22
- DD Boost volume sharing 23
- DD Boost performance 23
- Memory requirements for DD Boost 23
- Devices in mixed device environments 23
- DD Boost devices and media pools 24
- Reclaiming expired storage space 24
- Removing a DD Boost device 24
- DD Boost devices on Extended Retention systems 24

Network requirements 25
- Ethernet IP support 25
- FC support 26
- Firewall requirements 27

Deduplication efficiency 28
- Retention periods 28
- Types of data backed up 28
- Change rate 28
- Frequency of full backups 29

Host naming guidelines 29
Example name formats 30
IP addresses 30
Example topologies 30
Client Direct deduplication environment 30
Disaster recovery environment 31
Cascaded replication environment 33
Shared datazones environment 34
Dedicated storage node environment 35

Chapter 3 Data Protection Policies 37

Performing clone and replicate operations 38
Overview of data protection policies 38
Default data protection policies 39
Strategies for traditional backups 41
Road map for configuring a new data protection policy 41
Creating a policy 42
Creating a workflow in a new policy 43
Protection groups for traditional backups 45
Supported actions in traditional backup workflows 47
Actions sequences in traditional backup workflows 48
Visual representation of workflows 64
Cloning with Data Domain (DD Boost) 65
Clone formats 65
CCR requirements 66
Cloning by pools 67
DD Boost clone and replication support 67
Clone formats 67
Native Data Domain replication considerations 68
Automated Multi-streaming (AMS) 69
Configuring the CCR environment 69
Strategies for cloning 71
Road map for configuring a new cloning data protection policy 71
Road map to add a clone workflow to an existing policy 83
Clone reports 93
Cloning with nsrclone 93

Chapter 4 Software Configuration 95

DD Boost configuration road map 96
Configuring the Data Domain system for DD Boost 96
Configuring NetWorker for DD Boost devices 99
DD Boost device performance considerations 99
Methods to configure DD Boost devices 100
Configuring DD Boost devices with the wizard 100
Configuring DD Boost devices manually 103
Adding a host Data Domain system to NMC Enterprise view 104
Configuring a DD Boost device manually 104
Creating a volume label template for DD Boost devices 108
Creating pools to target DD Boost devices 109
Labeling and mounting devices on the Data Domain device 111
Configuring clients to back up to DD Boost devices 111
Configuring a backup client with the wizard 111
Configuring a backup client with NMC property windows 112
Deactivating and removing DD Boost devices 115

Converting a device to read-only 115
Disabling a device 115
Deleting an AFTD or DD Boost device 115

Chapter 5 **Restoring Data** 117

Restoring DD Boost deduplicated data 118
Restore requirements for deduplicated data 118
Supported NetWorker restore procedures 118
Restoring by Client Direct over IP from an FC-enabled device 119
Disaster recovery 119
Causes of disaster 119
Potential losses 120
Disaster recovery requirements 120
Disaster recovery scenarios 120

Chapter 6 **Monitoring, Reporting, and Troubleshooting** 121

Monitoring Data Domain events, statistics, and logs 122
Viewing the statistics, logs, and alerts 122
Viewing backup statistics in NMC 123
Viewing backup alerts (SNMP traps) in NMC 123
Configuring SNMP for an NMC managed Data Domain system 124
Generating reports 125
Configuring a report 125
Types of backup reports 127
Data Domain statistic reports 129
Advanced reporting 129
Replacing a failed or old storage node 129
Troubleshooting 131
Name resolution issues 131
Network connection issues 131
Device access errors 132
Backup fails for older NetWorker application modules 132
Multiple recovery fails on AIX clients with less than 2 GB RAM 133
Backing up streams from NetWorker to Apollo DD is rejected 133

Appendix A **DD Boost Conversion and Upgrade** 135

Upgrading from legacy DD Boost configurations 136
Converting DD Boost devices from IP to FC connectivity 136
Redirecting backups from other devices to DD Boost 136
Migrating data into Data Domain systems 137
Migration versus native Data Domain replication 138
Migration methods 138
Migrating legacy save sets to DD Boost devices 139
Migration scenarios 140
Migration to DD Boost from conventional tape or AFTD 140
Migration to DD Boost from deduplicated VTL or AFTD storage 142

Glossary 145

Index 151

FIGURES

1 Client Direct backup versus traditional storage node backup 31
2 CCR for disaster recovery 32
3 CCR cascaded to multiple Data Domain systems 34
4 Data Domain system shared across two NetWorker datazones 35
5 Single datazone with dedicated storage nodes and one high-bandwidth link 36
6 Data Protection Policy 39
7 Platinum policy configuration 40
8 Gold policy configuration 40
9 Silver policy configuration 40
10 Bronze policy configuration 41
11 Data protection policy example 42
12 All possible workflow actions for a traditional backup 48
13 Workflow path from a traditional backup action 48
14 Visual representation of a workflow 65
15 Workflow path from a clone action 77
16 Visual representation of a clone workflow 82
17 Example of a policy with separate workflows for backup and cloning 83
18 Example of the device name and the access information for a DD Boost device 106
19 NetWorker Administration window displaying DD Boost devices 122
20 Data Domain alerts to monitor 125
21 Report configuration 127
22 Migration from conventional storage to DD Boost devices on a different storage node 141
23 Migration from conventional storage to DD Boost devices on the same storage node 142
24 Migration from VTL to DD Boost devices on a different storage node 143
25 Migration from VTL to DD Boost devices on the same storage node 143

TABLES

1 Revision history 11
2 Style conventions 14
3 Firewall ports for DD Boost 27
4 Schedule icons 49
5 Schedule icons 52
6 Schedule icons 55
7 Schedule icons 60
8 Save set criteria 73
9 Schedule icons 78
10 Save set criteria 85
11 Schedule icons 89
12 Data Domain report configuration parameters 126
13 Data Domain basic reports 127
14 Data Domain statistics drill-down report 129

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 DD Boost devices 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
06	June 29, 2016	Updated this document for EMC NetWorker 9.0.1. with the following: • Improved performance of NetWorker Clone Controlled Replication (CCR) for Data Domain because of Automated Multi-Streaming. • DD Boost Fibre Channel supports Solaris 10 and 11 for the File System Client.
05	April 13, 2016	Updated the "Data Protection Polices" chapter with the following changes: • Modified the "Strategies for traditional backup" section to revise the order in which you create the data protection policy resources.

Table 1 Revision history (continued)

Revision	Date	Description
		• Added the section "Default data protection policies."
04	January 18, 2016	Added information in the "Configuring a backup client with NMC property windows" section to include details about the Backup disk devices attribute.
03	December 14, 2015	Minor updates to this document for EMC NetWorker 9.0.
02	November 12, 2015	Minor updates to this document for EMC NetWorker 9.0.
01	September 24, 2015	Minor updates to 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.
- *EMC Data Domain Boost Compatibility Guide*
  Provides compatibility information for EMC and third party applications, and Fibre Channel hardware solutions, that interoperate with EMC Data Domain Boost (DD Boost) technology.

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)
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

DD Boost Features and Environment

This chapter includes the following topics:

- DD Boost integration features 16
- Data Domain network environment 18
- Licensing in Data Domain systems 19

DD Boost integration features

The EMC® NetWorker® integration with EMC Data Domain® systems uses DD Boost deduplication devices, which offer the following integration features.

DD Boost and data deduplication

NetWorker client software uses DD Boost components to integrate with DD Boost logical storage devices on Data Domain systems, and perform data deduplication backups. Data deduplication is a type of data compression that removes duplicate information to dramatically reduce the amount of backup data sent to the storage devices. The reduction in data that is sent to the devices reduces the bandwidth that is required for the data transport.

DD Boost can run up to 60 concurrent sessions or save streams on each DD Boost device for backup and recovery. This high concurrence reduces the number of required devices and reduces the impact on the performance and maintenance of the Data Domain system. The resulting performance gain provides an advantage over conventional advanced file type device (AFTD) or virtual tape library (VTL) interfaces that do not handle these high session rates.

During recovery to a NetWorker client, the Data Domain system converts the stored data to its original non-deduplicated state.

Client Direct data handling

The Client Direct feature enables clients that have a direct network or a DD Boost over Fibre Channel (DFC) connection to the Data Domain system to send and receive data directly to Data Domain AFTD and DD Boost devices. Client Direct supports multiple concurrent backup and restore operations that bypass the NetWorker storage node, which eliminates a potential bottleneck. The storage node manages the devices that the clients use but does not handle the backup data. Client Direct was previously known as Direct File Access (DFA).

When the connection is available, NetWorker enables Client Direct by default and also uses Client Direct to recover duplicated backups that NetWorker had performed with a storage node.

DD Boost device storage

A Data Domain system stores deduplicated NetWorker backup data on DD Boost storage devices on the Data Domain system. The NetWorker server, storage nodes, and Client Direct clients can all access the DD Boost devices. The DD Boost devices appear as folders that reside in storage unit (SU) partitions.

Secure multi-tenancy

NetWorker supports DD Boost devices in secure multi-tenancy (SMT) storage on DD OS 5.5 and later systems. SMT enables service providers to isolate tenant users on a Data Domain system. A global storage administrator assigns or creates a tenant unit (TU) for each tenant user. Tenant users, for example, backup administrators, must use a DD Boost username and password to create the secure storage units (SUs) that the DD Boost devices use to store data.

Retention tier storage

The Data Domain Extended Retention software option, available with DD OS 5.2 and later, extends the Data Domain storage structure with internal tiers, for long-term retention of deduplicated backup data.

You can use Data Domain operations to migrate the data from the active tier to the retention tier. You can create DD Boost devices on a Data Domain Extended Retention system that runs DD OS 5.0 or later software. The archive tier does not require additional capacity licensing.

This model of a Data Domain system can support mixed environments that may include DD Boost devices, VTL, and CIFS or NFS AFTD disk configurations. On an Extended Retention-enabled Data Domain system, the NetWorker software interacts with the active tier only and is not aware of any migration activity between the internal storage tiers.

DD Boost Synthetic Full backups

The NetWorker Synthetic Full backup feature is an efficient way to create full backups by combining existing full and incremental backups. This feature integrates the NetWorker Synthetic Full backup feature and the DD OS 5.3 and later virtual-synthetics feature. NetWorker creates the Synthetic full backups directly on the DD Boost devices. Both the NetWorker software and the Data Domain system are configured by default to enable DD Boost synthetic full backups.

The *EMC NetWorker Administration Guide* provides more details about the Synthetic Full feature.

DD Boost in-flight encryption

NetWorker enables DD Boost clients to have in-flight data encryption with a Data Domain system that uses DDOS 5.5 or later over a WAN connection. To use this feature, configure the Data Domain system to use medium- or high-strength TLS encryption. The configuration is transparent to NetWorker.

The Data Domain documentation provides more information about DD Boost in-flight encryption.

DD Boost clone operations

You can create a clone of backup data that is stored on DD Boost devices, and then copy the clone data to remote offsite storage for added protection and efficient disaster recovery. To use the NetWorker clone feature, you add a clone action to a workflow in a data protection policy. The clone action generates information that NetWorker stores in the client file index and media database to enable data recovery. The retention policy assigned to the clone action defines the length of time that NetWorker retains this data. All data movement for NetWorker clone operations must use Fibre Channel (DFC) or IP network connectivity.

There are two types of clone operations:

- NetWorker clone-controlled replication (CCR or enhanced cloning) operations, which replicates data from a DD Boost device to another DD Boost device at a different location. CCR preserves the deduplicated data format and minimizes bandwidth usage between the Data Domain systems.
- Clone to native format operations, which clone data from DD Boost storage to conventional storage media, for example, disk or tape. This operation reverts the data

to the native non-deduplicated format, to enable recovery from a conventional disk device or tape device.

NMC management of DD Boost operations

Use NMC to efficiently configure, monitor, and generate reports for DD Boost devices. The NMC server and the NetWorker server must have network access to each managed Data Domain system.

The NMC Device Configuration Wizard simplifies the configuration of DD Boost storage devices, backup clients, storage pools, volume label operations, and save set clone operations.

Feature not supported by the integration

Native Data Domain directory replication (MTree replication) does not fully support DD Boost devices, which are rendered as read-only directories. Native Data Domain replication considerations on page 68 provides details.

Data Domain network environment

A NetWorker with Data Domain network environment can include the following components.

FC and IP network connectivity

DD Boost devices support data transport over Fibre Channel (FC) and Ethernet IP (IPv6 and IPv4) network connections for backup and recovery operations with Data Domain systems. DD Boost over FC (DFC) connectivity requires DD OS 5.3 or later.

NetWorker supports DD Boost devices on IPv6 networks. DD OS 5.5.1 supports IPv6 network usage with DHCP, DNS, and DDNS internet services. The *EMC Data Domain Operating System Administration Guide* provides configuration details.

When you use DFC for data transport, review the following information:

- The NetWorker server also requires IP connections to communicate with all the hosts that are involved in DD Boost operations and for data transport during recovery and clone-controlled replication operations.
- Client Direct backup with DFC is not supported for 32-bit Linux NetWorker clients that are installed on 64-bit Linux systems. The backup reverts to a storage node backup.
- For DFC to work the encryption strength for the client "*" on the DD OS, set the option to **None**. By default the setting is set to **High**.

Data Domain storage system

A Data Domain system can store deduplicated backup or clone data on DD Boost devices, and supports mixed environments that may include DD Boost devices, VTLs, and CIFS or NFS AFTD disk configurations. The Data Domain system may require additional licenses for the DD Boost functionality.

The *EMC NetWorker Online Software Compatibility Guide* provides compatibility information.

NetWorker client

A NetWorker client is a supported host whose data requires protection. The NetWorker client software includes an integrated DD Boost plugin. The NMC server, NetWorker server, and NetWorker storage nodes are also NetWorker clients.

NetWorker clients that use Client Direct deduplication must have direct network access to the Data Domain system, which stores the data. NetWorker enables Client Direct by default in the properties of the Client resource.

Client Direct with FC connectivity to DD Boost devices requires NetWorker client 8.1 or later.

The *EMC NetWorker Online Software Compatibility Guide* provides information on supported releases.

NetWorker server

The NetWorker server provides services to back up and recover the data of any NetWorker host in a datazone. The NetWorker server also acts as a storage node and can control multiple remote storage nodes.

NMC server

The NetWorker Management Console (NMC) server or Console server is a Java-based web application and database server. The NMC server manages all NetWorker servers and clients. The NMC server also provides reporting and monitoring capabilities for all NetWorker servers and clients in the environment. NMC server relies on the EMC NetWorker Authentication Service for user account authentication.

NetWorker storage node

NetWorker storage nodes manage DD Boost and other storage devices. When a client does not use Client Direct, the NetWorker storage node deduplicates the backup data, then sends the deduplicated data to the DD Boost devices.

Install the same version, including the service pack of the NetWorker storage node software on each host in the datazone that stores backup or clone data on DD Boost devices.

NetWorker application modules

NetWorker supports Client Direct deduplication backup and recovery on clients with supported NetWorker application modules, for example, NetWorker Module for Databases and Applications (NMDA), NetWorker Module for Microsoft Applications (NMM), and NetWorker Module for SAP (NMSAP). The clients must have direct network access or Fibre Channel access to the Data Domain system. The release notes for the application module provide details.

Licensing in Data Domain systems

There are three types of licensing models that enable the NetWorker server to interact with a Data Domain system:

- The EMC Licensing Solution with capacity entitlement, introduced in NetWorker 9.0, which uses an EMC License Server and a license file.

- Traditional Licensing, which uses individual enabler codes to license features. NetWorker requires only a single enabler to support multiple interfaces and multiple network identities for Data Domain systems.
- Capacity Licensing from NetWorker 8.2.x and previous releases, which licenses the datazone by using capacity-based enabler codes.

The *EMC NetWorker Licensing Guide* provides licensing details.

Traditional licensing for Data Domain systems

If using traditional licensing, a new installation of the NetWorker server software enables you to evaluate all the features for 30 days, including the Data Domain features, without the use of an enabler (license key). To extend this evaluation period by 15 additional days, type the word grace in the Auth code field of the NetWorker server evaluation license before the end of the 30-day period. After the evaluation period ends, you cannot perform a backup unless you install permanent license enabler codes.

Data Domain licenses

Enable either the Data Domain system or the Data Domain system with Extended Retention software by using the following licenses for DD Boost operations:

- DDBOOST license
- REPLICATION license, if you use CCR

To generate a list of the enabled licenses on the Data Domain system, type the `license show` command. The output displays DDBOOST when a DD Boost license is installed, and REPLICATION when a replication license is installed. Configuring the Data Domain system for DD Boost on page 96 provides details.

For license inquiries, go to the Data Domain portal for service and support at https://support.emc.com.

CHAPTER 2

Planning and Practices

This chapter includes the following topics:

- DD Boost storage characteristics 22
- Network requirements 25
- Deduplication efficiency 28
- Host naming guidelines 29
- Example topologies 30

DD Boost storage characteristics

NetWorker integrates with Data Domain systems by storing backup data on DD Boost devices.

The *EMC NetWorker Online Software Compatibility Guide* provides information about supported releases.

DD Boost storage structures and limits

DD Boost devices use a folder structure on the Data Domain system that has the following characteristics:

- The Data Domain storage consists of separate logical partitions called storage units (SUs) or managed trees (MTrees).
- By default, the NetWorker Device Configuration Wizard creates one SU for each NetWorker datazone. The wizard names the SU after the short hostname of the NetWorker server.
- You can define up to 99 active SUs on a Data Domain system. For DD OS earlier than 5.3, running more than 14 active SUs can impair performance.
- DD Boost devices appears as subfolders within the SU folder.
- You can associate each DD Boost device with only one NetWorker storage volume. However, a single NetWorker volume can share multiple DD Boost devices, which in some environments can improve performance.
- You must use DD Boost credentials to create the SUs and the DD Boost devices, and to enable secure multi-tenancy (SMT) access to the DD Boost devices.
- Avoid changing the user of a device; the new user will not have permission to the files and directories created by the previous user and cannot re-label the volume; create a device for the new user.
- For SMT, NetWorker supports up to 512 DD Boost devices on each SU. Otherwise, there is no limit on the number of DD Boost devices that you can create. For best performance, use fewer devices and more backup sessions on each device.
- Data Domain MTree quotas manage the usable capacity of SUs. When an MTree reaches the defined quota, any running DD Boost backup operation terminates.

SMT structure

For SMT storage, a global storage administrator isolates DD Boost users, for example backup administrators, by assigning them to tenant units (TUs). A TU cannot span Data Domain systems. You can assign a DD Boost user to only one TU, but you can assign multiple DD Boost users to the same TU. Each DD Boost user can create SUs within the assigned TU only. Security is enforced at the TU level by the DD Boost user assignment on the Data Domain system, and at the SU level by the DD Boost credentials.

For example:

Tenant 1: bob, tu1, su1

Tenant 2: joe, tu2, su2

Tenant 3: sue, tu2, su3

DD Boost volume sharing

Multiple DD Boost devices, specified by different names or aliases, can concurrently share a single NetWorker storage volume.

Each DD Boost device operates with a single NetWorker storage volume and you must specify each device by the device access pathname.

- Each DD Boost device operates with a single NetWorker storage volume and you must specify each device by its device access pathname.
- You can create multiple devices with the same device access pathname, provided that you assign each device a different name, as an alias. You can use the different device aliases, for example, to manage different client hosts that share the same volume.

Configuring a DD Boost device manually on page 104 provides details on device access information.

DD Boost performance

DD Boost devices use multiple concurrent `nsrmmd` (media mover) processes per device and each `nsrmmd` process uses multiple concurrent save sessions (streams or threads). This reduces the performance and maintenance impacts on the Data Domain system.

Balance the session load among the available DD Boost devices so that new sessions attach to devices with the least load. To enable best performance, you can adjust the **Target Sessions**, **Max Sessions**, and **Max nsrmmd Count** attributes assigned to the Device resource on the NetWorker server.

Configuring a DD Boost device manually on page 104 provides details on session settings.

The Data Domain documentation provides additional details on save sessions and performance.

Memory requirements for DD Boost

The physical memory requirement for a NetWorker storage node and a Client Direct client depends on the peak usage of the DD Boost devices.

Ensure the following memory requirements:

- A storage node host that manages DD Boost devices with other typical devices and typical services must have a minimum of 8 GB of RAM.
- A DD Boost client requires a minimum of 4 GB of RAM at the time of backup to ensure the best performance for Client Direct backups.

Devices in mixed device environments

A Data Domain system can support mixed NetWorker environments, which can include DD Boost devices, VTLs, and CIFS or NFS AFTD disk configurations. Each different storage environment must use a different interface connection. Use a NIC for IP data transport and an FC port for SAN data transport. You can use the NetWorker clone process over an IP connection to migrate legacy data that is stored in these traditional storage environments to DD Boost devices.

DD Boost devices and media pools

Use media pools to send client backups or clones of the backup data to specific storage devices. You must configure pools for DD Boost devices with the following important consideration.

A pool that contains the DD Boost devices must not contain any other type of storage media and must not contain devices on more than one Data Domain system.

This practice ensures an optimal backup window and deduplication ratio with the least amount of interference. When you store each backup on a single Data Domain system, you simplify recovery and Data Domain file replication.

Periodically review and remove unused pools that are no longer relevant to the storage environment.

Reclaiming expired storage space

When a backup on a Data Domain system reaches the retention expiry date, all references to the data become invalid. However, unlike conventional storage systems, the Data Domain system does not immediately free up disk space on the device because other existing backups or other clients may continue to have deduplication references to the same data.

For example, the deletion of 1 GB of data, even of unique data from a NetWorker perspective, does not mean that 1 GB of space is immediately available on the Data Domain system.

The DD OS `filesys show space` or `df` command shows the data that is eligible for deletion under the **Cleanable GiB** column.

The Data Domain system periodically reviews deduplication references and storage space and performs a cleanup. By default, these cleanups occur weekly.

Note

If a DD Boost device becomes full during a backup, the backup immediately fails. The device does not pause or wait for space to become available.

The DD OS `filesys clean` command describes all the available options that you can use to reclaim and recycle expired NetWorker save sets and other redundant or expired data.

Removing a DD Boost device

You must use special procedures to remove DD Boost devices.

Deleting an AFTD or DD Boost device on page 115 provides details.

DD Boost devices on Extended Retention systems

You can use SU folders or MTrees and DD Boost devices on Data Domain Extended Retention systems, similar to those on standard Data Domain systems, with the considerations described in this section.

The Data Domain documentation provides details on Data Domain Extended Retention features.

Active and extended retention tier structure

On Data Domain Extended Retention systems, you can store data in two tiers of SUs, an active tier and an extended retention tier.

- Data Domain places all incoming data first in the active file system tier, which is for short-term data storage and is similar to standard Data Domain systems.
  You can use the active tier for client backups, provided that you apply appropriate data movement and data retention policies. It is recommended that you create separate SUs for backup operations.
- You can move data from the active tier to the archive tier, based on data movement policies that you apply at the SU level.

DD OS 5.0 and 5.1 and later support up to 14 active and 99 defined SUs with the Extended Retention software feature. NetWorker does not limit the number of DD Boost devices that you can create on the system.

Data movement between tiers

Each SU has a single data movement policy that applies to all the devices that the SU manages within the corresponding NetWorker datazone.

You can create a Data Domain SU data movement policy to specify when the data moves from devices in the active tier to devices in the archive tier. Typically, you would not move data to the archive tier for less than 30 days retention. The policy and movements are internal to the Data Domain Extended Retention system, and the NetWorker software has no awareness of the operations.

You can assign alternative data movement policies to the client data by using additional SUs, created by NMC or `nsradmin`, for moving data from one tier to another. For example, you can store data to different archive DD Boost devices in separate SUs with different archive policies. Also, you can move data within the same Data Domain Extended Retention system by using CCR.

To use CCR on the same Data Domain system that includes the Extended Retention software feature, you must replicate between two different SUs. You can apply different retention policies to manage the data efficiently.

Network requirements

DD Boost devices support data transport over both Ethernet IP networks and FC SAN environments for both data backup and data recovery operations.

The NetWorker server requires Ethernet IP connections to control all hosts involved in the DD Boost operations.

Ethernet IP support

DD Boost devices do not distinguish between different TCP/IP network types (LAN, WAN, or MAN) and can successfully operate in a network where packet loss is strictly 0% and latency is less than 20ms. Variations of IP network connections can improve data throughput, depending on the Data Domain system model.

EMC recommends a minimum of two separate IP network connections to the Data Domain system. One is used for administration and the other is used for data backup.

Aggregated multiple connections can further improve data throughput for the Data Domain system. For example, you can use multiple 1 GbE connections for dedicated storage nodes and storage devices. Connections for 10 GbE are also available and you can use these instead of or with 1 GbE links.

You can configure two basic IP interfaces:

- Dedicated 1 GbE data connection from the storage node directly to the Data Domain system. This connection provides a private, high-bandwidth data connection and avoids the latency and complexity of a shared IP connection. You also require a separate conventional IP connection for administration and NetWorker Console access.
  The *EMC Data Domain Operating System Administration Guide* provides details on network support.
- Two or more NICs on the Data Domain system with 1 GbE or 10 GbE connections, which are aggregated together by using the Data Domain `ifgroup` command. This grouping provides increased data capacity and can improve resiliency. The Data Domain system provides automatic Advanced Load Balancing and Link Failover for NIC connections.

  Note

  Do not use aggregated connections for replication operations.

  The *EMC Data DomainBoost for Open Storage Administration Guide* describes the benefits, limitations, and examples of using ifgroups, which apply to NetWorker.

FC support

NetWorker supports data backup and recovery operations to DD Boost devices over Fibre Channel (DFC or FC) connections, which are configured as a SAN, as follows:

- The NetWorker storage nodes and all Client Direct clients must have FC SAN network access to the Data Domain systems that have FC-enabled DD Boost devices.
- The environment must have an Ethernet IP network. The NetWorker server uses IP connections to communicate with the clients, storage nodes, and the Data Domain system. DD Boost devices that are involved in CCR operations must have IP connectivity for the data transport.
- FC-enabled NetWorker clients must run on a supported Windows, Linux, HP-UX, AIX, or Solaris operating system. HP-UX systems must use NetWorker 8.2 SP1 or NetWorker 9.0 clients and storage nodes. AIX systems must use NetWorker 9.0 clients and storage nodes. Supported Solaris versions on a client are Solaris 10 and 11 on SPARC with x86 architectures. Solaris uses SCSI generic device driver, sgen. The sgen driver is included in the Solaris installation. The *EMC NetWorker Online Software Compatibility Guide* and the *EMC Data Domain Boost Compatibility Guide* provide details.

  Note

  On AIX, DD Boost-over-FC requires a device driver. The AIX Installation Chapter in the *EMC NetWorker Installation Guide* provides more details.

- The NetWorker clients and NetWorker storage nodes must run NetWorker 8.1 or later software.
- The Data Domain system must run DD OS 5.3 or later.
- All hosts that use FC must have an HBA card with at least 4 Gbps bandwidth capacity and must devote an initiator port on each card to FC for DD Boost devices. You should configure Access groups. The *EMC DD OS Administration Guide* provides details.
- FC-enabled DD Boost devices support Client Direct backup and restore over FC, provided that you have enabled the clients with FC connections and settings.

- Data Domain systems support the coexistence of FC-enabled DD Boost devices together with VTL devices on the same Data Domain system. However, the FC-enabled DD Boost devices must not share an initiator with VTL devices on the same Data Domain system.
- You can convert an existing DD Boost device from IP to FC connectivity and settings without losing the stored data. You can restore the data to FC-enabled Client Direct clients through their FC connection, and to IP-only clients through the storage node. Converting DD Boost devices from IP to FC connectivity on page 136 provides details.
- The NetWorker server can migrate legacy backup data stored on a VTL or tape device to an FC-enabled DD Boost device. You can create a clone pool for this migration. Considerations for migrating legacy save sets on page 137 provides details.
- CCR is supported between FC-enabled DD Boost devices provided that there is IP connectivity between the Data Domain systems. CCR is not supported over a Fibre Channel network.

The *EMC NetWorker Online Software Compatibility Guide* provides the latest details of supported versions.

Firewall requirements

Regardless of the network connections that are used, communication through a firewall requires the use of specific ports and specific protocols to perform backup, monitoring, and replication operations across sites.

The following table lists the required firewall ports, which you must open between the Data Domain system, the NetWorker server, and the NMC server.

Table 3 Firewall ports for DD Boost

Port	Use
TCP 111	NFS portmapper
UDP 111	NFS portmapper
TCP 161	For the NMC server to query for alerts and statistics
TCP 162	SNMPTRAP for the NMC server to monitor status and events
TCP 2049	NFS
TCP 2051	Replication, if you use CCR to replicate between Data Domain systems
TCP *xxxx*	1. Select a random port for NFS mountd. The default port is 2052. 2. On the Data Domain system, type the following command from SE mode: **`# nfs set mountd-port xxxx`** 3. Restart the Data Domain system.

The Data Domain system provides functionality to review the network configuration and network capabilities and provides SSH Telnet to help diagnose issues.

Deduplication efficiency

The deduplication ratio measures the efficiency of reduction in storage space that results from the data deduplication and compression technology. Ratios of 20:1 are broadly achievable and reductions of even 5:1 are extremely valuable.

Several factors can contribute to the deduplication ratio:

- Retention periods
- Types of data backed up
- Change rates
- Frequency of full backups
- Use of encryption and compression

For the best use of storage space, consider the factors in the following sections, along with the periodic clearing of expired storage space, and the removal of unused pools.

Retention periods

The deduplication ratio increases with longer data retention periods. The longer you retain the stored save sets, the greater the chance that identical data will exist on the storage that NetWorker uses to deduplicate each subsequent backup, and the greater is the efficiency of the storage usage.

When you define longer retention periods, the data remains on the Data Domain device for a longer period of time. This enables NetWorker to use the retained data to deduplicate subsequent backups, and results in a more efficient use of storage.

Types of data backed up

Some types of data, for example, text documents, slide presentations, spreadsheets, email, source code, and most database types, contain redundant data and are good deduplication candidates.

Some other types of data, for example, audio, video, and scanned images already consist of compressed data. Typically, the first full deduplication backup of these data types yields low reductions, but subsequent backups generally produce high deduplication ratios if the data has a low change rate.

Change rate

Data with a low change rate changes little between backups, produces high deduplication ratios, and is a good candidate for deduplication. Deduplication removes data that is already in storage and only stores new data.

When a new save set is deduplicated, the number of unique blocks within the save set can vary widely depending on the data type, and often there is little that can be deduplicated. Yet because the Data Domain system compresses the data blocks, there is typically a 2:1 to 3:1 (50%–75%) data reduction.

The storage savings increase with each subsequent backup of the save set because a deduplication backup writes to disk only the data blocks that are unique to the backup. In conventional business operations, the data change rate is typically low and unique data may represent only 1%–2% of the data present in each additional backup set. The remainder of the backup is deduplicated against the data already stored on the system.

Frequency of full backups

Frequent full backups result in high deduplication ratios, but also increase the data processing operations on the NetWorker storage node or Client Direct client. For example, compare daily full deduplication backups with weekly full and added daily incremental deduplication backups. Both of these schedules require essentially the same amount of storage space and the same network bandwidth between the storage node and the Data Domain system. The backups send only unique data to storage, even for full backups.

A daily full backup schedule, however, sends a greater amount of data from the client to the storage node for processing than the weekly full with daily incremental schedule.

Host naming guidelines

The network environment has an impact on hostname resolution methods and you must follow the manufacturer recommendations. Use the local hosts file to help diagnose and resolve naming issues. You can use the `net hosts add` command on the Data Domain system to add hosts to the `/etc/hosts` file.

Use the following guidelines to create consistent, easy-to-identify hostnames that improve the configuration, report generation, and troubleshooting experience in the DD Boost environment:

- Create names that are unique across all NetWorker datazones. Use names that identify the network role, for example, administration, backup, cloning, or production. A name can also include a location or a server name.
- Use a single hostname that is associated with each NIC, IP, or FC interface within the same NetWorker datazone.
- Names can include abbreviations for the source or the target to quickly identify whether the network connections are correct. For example, add an abbreviation of the storage node hostname in the Data Domain name and an abbreviation of the Data Domain hostname in the storage node name. Include the names in the Data Domain `/etc/hosts` file.
- Specify all aliases, such as long and short names and IP addresses, for the NetWorker server and the storage nodes in their respective Client resources. Specify the aliases in the **Aliases** attribute on the **Globals (1 of 2)** tab of a Client resource.
- Test to ensure that you can consistently resolve all hostnames in the network from multiple locations in both directions. For example, ensure that you can resolve the short name to IP address, long name to IP address, IP address to short name, and IP address to long name.
- In general, use short, easy-to-identify, descriptive names instead of IP addresses or fully qualified name strings for devices and storage nodes. Long names may not fit into some views. The following examples include a long name and a short name:
  NWDD365-1.burloak.lab.mycorp.com:/NWDZ_Dr1

  NWDD365-1:/NWDZ_Dr1
- Except for hostnames, use standard alphanumeric characters, including dot (.), hyphen (-), and underscore (_), with no spaces and no special characters. Hostnames may not use underscores (_).
- Use consistent formats, in terms of text field length and text case, and include leading zeros in numbers, with a maximum of 50 characters.

- Avoid the use of dates in names where the dates could change or become meaningless in the future.

Example name formats

The following examples provide some name formats.

DD Boost devices

Format: *Data_Domain_system_name-device_name*

For example: dd-tenendo-device01

Folders on Data Domain system

Create DD Boost device names that refer to the NetWorker storage node and indicate whether you use them for either backup or clone operations.

Format: *storage_node_name-operation-device_name*

For example: dzburl-back-dd01

Volume labels for Data Domain

Format: *media_type-label_number*

For example: ddmedia-001

IP addresses

Avoid IP addresses because numbers are more difficult to identify and troubleshoot than descriptive names. However, there are exceptions:

- The Data Domain system requires the use of IP addresses to interface with an ifgroup for Advanced Load Balancing and Link Failover features.
- For CCRs, the hosts file on the source Data Domain system must list the IP address of the target Data Domain system. Otherwise, the CCR will use the same network access that the backup and restore operations use.

The Data Domain documentation provides details.

Example topologies

This section provides some examples of how you can deploy the Data Domain integration in NetWorker backup environments. EMC recommends that you use two interfaces in Ethernet IP networks, 1 GbE for administration and 10 GbE for data. For FC environments, use IP interfaces for administration and clone operations, and a SAN interface for backup operations. Use the following examples to plan your environment.

Client Direct deduplication environment

Client Direct functionality is enabled by default in the Client resource and NetWorker tries to use Client Direct for data backup or recovery operations. If the client does not have a direct network connection to the Data Domain system, then these operations automatically use the traditional storage node workflow.

Client Direct data handling on page 16 describes the Client Direct feature, which leverages client DSP software to send deduplicated data directly from a client to a Data Domain system and bypass the storage node.

The following figure shows an example Client Direct environment.

Figure 1 Client Direct backup versus traditional storage node backup



Client Direct deduplication provides the following advantages:

- Deduplication on the client host dramatically reduces the bandwidth that is required for the backup data transfer.
- You can share a single storage volume among multiple DD Boost devices and among multiple backup hosts and storage nodes. You can improve performance and maintainability by configuring multiple hosts and multiple sessions for each device, instead of creating multiple devices.
- Client Direct deduplication offers an alternative to an environment that uses dedicated NetWorker storage nodes, as described in Dedicated storage node environment on page 35. The dedicated storage node environment requires additional licensing and configuration, and the backup clients or the applications on the clients may not support a dedicated storage node.

Disaster recovery environment

To perform a disaster recovery you can use CCR to copy individual save sets or complete volumes from one Data Domain system to another at a geographically distant location. Each cloned replication, or optimized clone, is a complete and independent copy of the source deduplicated data. NetWorker policies manage both the source or primary data and the clone or secondary data. For additional protection, NetWorker can clone some or all the stored data from the secondary system to tape storage.

You must configure, enable, and manage both the primary and secondary Data Domain systems within a single NetWorker datazone. Configure target devices on the secondary Data Domain system. You can use either a single storage node or separate storage nodes for the local and remote Data Domain system within the datazone.

Note

NetWorker does not support CCR across datazones or to Data Domain devices that are not managed by NetWorker.

The following figure illustrates an example of a disaster recovery environment. The NetWorker server requires two Data Domain Storage System Enablers, one for the primary Data Domain system and one for the remote secondary Data Domain system. DD in the figure signifies Data Domain. The following actions occur in this example:

Figure 2 CCR for disaster recovery



1. The NetWorker server starts the backup of the client groups within the datazone.
2. Two storage nodes in the datazone write the backup data to media pools, which target specific DD Boost devices on the primary system. The pool that is associated with the data protection policy defines which storage devices receive the data.
3. The storage nodes communicate with the primary Data Domain system over a dedicated 10 GbE network connection, and store deduplicated backup data on the devices.

Note

An ifgroup configuration of 1 GbE or 10 GbE NICs on the Data Domain system enables multiple storage nodes to use the same network identity. This bandwidth aggregation can improve performance for DD Boost devices. The Data Domain documentation provides details.

---

4. You can use CCR to store optimized clone copies of backups from the primary Data Domain system over a network to a geographically distant secondary Data Domain system for disaster recovery.
5. An additional option enables a further clone to conventional disk or conventional tape media. A NetWorker storage node, which is attached to the secondary Data Domain system, creates an additional NetWorker clone copy of the data for one of the backup groups, which NetWorker stores on conventional media. NetWorker reverts

the data in this copy to the native non-deduplicated format, which is necessary for storage on conventional media.

Cascaded replication environment

A variation of the disaster recovery environment is the cascaded replication environment. Once a deduplicated backup completes successfully, you can use the backup to create multiple clone copies in other locations, either simultaneously from the original deduplicated backup or in sequence from a CCR copy. Each clone replication is a complete and independent copy of the source backup. NetWorker does not limit the number of cascaded clone copies that you can create, provided that the source save set for each clone successfully completes.

As with the previous example, configure, enable, and manage each Data Domain system in a cascaded environment within a single NetWorker datazone. Configure target devices on the Data Domain systems that receive the clone copies.

The figure in this section illustrates an example of a cascaded replication environment with three separate Data Domain systems at three different sites.

- The first site is the primary backup location and is the production site.
- The second site is a local site with good communication links to the production site, typically within the same geographic location as the first site.
- The third site serves as the disaster recovery site, which is located at a geographically distant location. Communication to this more distant site is subject to greater restrictions on bandwidth and latency. This distant site could be in a different country or 250 kilometers (150 miles) or more distant from either of the other two sites.

Note

The NetWorker server requires three Data Domain Storage System Enablers, one for each Data Domain system. DR in the figure signifies disaster recovery.

This example environment operates as follows.

1. The NetWorker server starts the backup of production site client groups within its datazone.
2. The production site storage node assigns the backup data to media pools, which uses specific DD Boost devices on the primary Data Domain system.
3. The storage node communicates with the primary Data Domain system over dedicated 10 GbE network connection, and stores deduplicated backup data on devices DD Device 01 and DD Device 02.
4. After the backup completes successfully, you can use CCR to store optimized clone copies of the backups, which reside on the primary Data Domain system, over the network to Data Domain systems at a local secondary site. You can create these clone copies by using one of the following methods:
   - Sequential method—NetWorker performs only one clone operation at a time, in sequence. This method allows the production system to continue to function without the need to create additional clones for a distant site.
     For example, NetWorker uses the original backup on the primary Data Domain system to create an optimized clone copy on a local secondary Data Domain system. Once this process completes, NetWorker uses this copy to create an additional optimized clone copy on the geographically distant Data Domain system.

     Data paths 1a and 1b in the following figure represent this method.

- Concurrent method—NetWorker may be able to perform the clone operations simultaneously. This method impacts the production system and requires more replication bandwidth.

  Note

  The concurrent method depends on many factors, and you would must validate and test the performance at the individual sites.

  For example, NetWorker uses the original backup on the primary Data Domain system as the source to create simultaneous clones on two target Data Domain systems.

  Data paths 2a and 2b in the following figure represent this method.

Figure 3 CCR cascaded to multiple Data Domain systems



Shared datazones environment

You can store backups from two separate datazones on a single Data Domain system. In this configuration, consider dividing the stream counts and the memory resources to manage the two datazones as separate entities. Do not let one datazone impact the performance of the other datazone. The total number of streams and devices cannot exceed the total capacity of the Data Domain system.

The figure in this section illustrates a dedicated 10 GbE network connection shared by three storage nodes in two NetWorker datazones. Two storage nodes belong to the DZ-A datazone, and one storage node belongs to the DZ-B datazone.

1. The two NetWorker servers begin the backups within their respective datazones.

2. The three storage nodes write the deduplicated backup data to DD Boost storage devices on the Data Domain system over the 10GbE connection. The pool that is associated with the data protection policy defines which storage devices receive the data.

   Note

   You cannot share a DD Boost device across datazones.

3. You can perform an additional backup to tape storage operation, either directly from a storage node or by a NetWorker clone operation from the Data Domain system.

Figure 4 Data Domain system shared across two NetWorker datazones



Dedicated storage node environment

NetWorker supports deduplication backups for high-volume clients that are also a dedicated storage node. For example, you can configure a client host that runs NetWorker Module for Databases and Applications (NMDA) as a dedicated storage node.

This environment can coexist with data protection policy configurations that use shared NetWorker storage nodes in the same datazone. However, because this is a private network, the connection and the devices that the storage node uses are not available to other NetWorker clients.

Note

The Client Direct (DFA) feature can provide similar benefits without the need for storage node licenses.

The figure in this section illustrates a mixed environment of shared and dedicated storage nodes.

1. The NetWorker server starts a backup of file system and module data on a dedicated storage node.
2. The storage nodes write the deduplicated backup data to the DD Boost storage devices on the Data Domain system. The pool that is associated with the data protection policy defines which storage devices receive the data.

Note

An ifgroup configuration of 1 GbE or 10 GbE NICs on the Data Domain system enables multiple storage nodes to use the same identify on an IP network. This aggregation of bandwidth can improve performance for DD Boost devices. The Data Domain documentation provides details.

3. A high-volume storage node uses an additional dedicated 10 GbE direct connection.

EMC NetWorker Online Software Compatibility Guide provides information on NetWorker application modules compatible with Data Domain systems.

Figure 5 Single datazone with dedicated storage nodes and one high-bandwidth link



CHAPTER 3

Data Protection Policies

This chapter includes the following topics:

- Performing clone and replicate operations 38
- Overview of data protection policies 38
- Default data protection policies 39
- Strategies for traditional backups 41
- Cloning with Data Domain (DD Boost) 65
- DD Boost clone and replication support 67
- Automated Multi-streaming (AMS) 69
- Configuring the CCR environment 69
- Strategies for cloning 71
- Clone reports 93
- Cloning with nsrclone 93

Performing clone and replicate operations

Data Protection policies provide you with the ability to backup data, which you can then clone and replicate.

Overview of data protection policies

Data protection policy is a concept that provides you with the ability to design a data protection solution for the environment at the data level instead of at the host level. With a data protection policy, each client in the environment is a backup object and not simply a host.

Data protection policies enable you to back up and manage data in a variety of environments, as well as to perform system maintenance tasks on the NetWorker server.

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

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

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

Workflows
Workflows define the start time for a series of actions, the frequency in which the actions run, the order of actions in a sequence, and the protection group to which the workflow applies.

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

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

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

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

Actions
Actions are the key resources in a workflow for a data protection policy and define a specific task, for example, a backup, clone, or snapshot. NetWorker uses a work list to

the define task. A work list is composed of one or several work items. Work items include client resources, virtual machines, save sets, or tags. You can chain multiple actions together to occur sequentially or concurrently in a workflow. All chained actions use the same work list.

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

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

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

Figure 6 Data Protection Policy



Default data protection policies

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 7 Platinum policy configuration

Gold policy

The Gold policy provides an example of a data protection policy for an environment that contains virtual machines and requires backup data redundancy. 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 8 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 9 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 10 Bronze policy configuration



Strategies for traditional backups

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

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 11 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 traditional backups

Create protection groups for traditional backups, which identifies the Client resources to back up.

You can create two types of protection groups for a traditional backup:

- 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.

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

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

Note

A Client resource can belong to more than one group.

Creating a 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.

Supported actions in traditional backup workflows

Traditional backup workflows can optionally include a probe or check connectivity action before the backup, and a clone action either concurrently with or after the backup.

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.

Traditional backup

A traditional backup is a scheduled backup of the save sets defined for the Client resources in the assigned group. You must specify the destination storage node, destination pool, the schedule (period and activity), and the retention period for the backup.

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.

Actions sequences in traditional backup workflows

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

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

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

All possible workflow actions for a traditional backup

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

Figure 12 All possible workflow actions for a traditional backup



Workflow path from a traditional backup action

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

Figure 13 Workflow path from a traditional backup action



Creating a check connectivity action

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

Before you begin

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

Procedure

1. In the expanded left pane, select the 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.

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

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

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

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

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

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

      Note

      The **Abort action** option 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.

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

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

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

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

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

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

      Note

      The **Abort action** option 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 traditional backup action

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

Before you begin

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

Procedure

1. In the expanded left pane, select the 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 **Backup**.
6. From the secondary action list, select the backup type, for example, **Traditional**.
7. 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.
8. 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.
9. 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**.
10. Click the icon on each day to specify the backup level to perform.

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

Table 6 Schedule icons

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

Table 6 Schedule icons (continued)

Icon	Label	Description
	Note Not supported for NDMP.	
	Skip	Do not perform a backup on this day.

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

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

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

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

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

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

11. Click **Next**.

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

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

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

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

15. From the **Client Override Behavior** box, specify how NetWorker uses certain client configuration attributes that perform the same function as attributes in the Action resource.

    - **Client Can Override**—The values in the Client resource for **Schedule**, **Pool**, **Retention policy**, and the **Storage Node** attributes will take precedence over the values that are defined in the equivalent Action resource attributes.

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

16. Click **Next**.

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

17. 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.

18. 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.

19. 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.

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

    Note

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

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

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

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

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

21. 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.

22. 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.
23. 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.
24. (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.

25. 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.

26. 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.

27. Click **Next**.

    The **Action Configuration Summary** page appears.

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

After you finish

(Optional) Create a clone action to automatically clone the save sets after the backup. A clone action is the only supported action after a backup action in a workflow.

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.

Table 7 Schedule icons (continued)

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.

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

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

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

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

16. From the **Failure Impact** list, specify what to do when a job fails:
    - To continue the workflow when there are job failures, select **Continue**.
    - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action** .

      Note

      The **Abort action** option 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 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.

The following figure illustrates the visual representation of a sample workflow for a traditional backup.

Figure 14 Visual representation of a 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**.

Cloning with Data Domain (DD Boost)

As with other NetWorker devices, you can use Data Domain device types to perform clone operations. You can clone single save sets or the entire Data Domain volume from a Data Domain device. You can also use the Data Domain device as the target device, to receive cloned data.

Cloning works differently for deduplication devices. You can perform clone-controlled replication (CCR), or optimized cloning of data, from one Data Domain system to another. Or you can clone data from a Data Domain device to tape or to any other device type.

Clone formats

Yo can clone data that is stored on a Data Domain device in one of two formats, which depend on the target media device:

- CCR format
- Regular clone format

Clone-controlled replication format

When you clone data to a target Data Domain device, typically at a remote location, the data retains the deduplication format. This format is known as clone-controlled replication (CCR), or as an optimized clone.

CCR uses the native Data Domain replication feature to copy data from one Data Domain system to another.

CCR uses a special Data Domain API and differs from standard directory level replication, which is also supported. The clone is created quickly and uses low bandwidth and low storage capacity.

You can use a clone that is created in this format for data recovery or to create further copies, for example, to traditional disk or tape storage. This method results in minimal impact on production or primary backup and recovery operations.

Immediate cloning

NetWorker supports immediate cloning with CCR. Immediate cloning means that you can clone each save set when the backup completes instead of waiting until the backup completes for all save sets in the action. Cloning operations can complete sooner because they can now run in parallel instead of sequentially. Performance gains are most noticeable when there are many backup save sets in the backup queue or when there are many save sets of different sizes.

You can set up immediate cloning by specifying the clone action as concurrent to the previous backup action in a policy workflow.

Regular clone format

When you clone the data on the Data Domain device to a traditional disk or tape, the clone process reverts the data to the native non-deduplicated format, known as "regular clone" format.

NetWorker requires the data on traditional disk or tape to be in regular clone format to ensure that the data is completely recoverable, without the need of a Data Domain system.

The process that takes data that has been deduplicated and then reverts it to normal or regular data is called rehydration.

CCR requirements

Before you use CCR to clone data, ensure that following requirements are met.

1. Ensure that both the source and target storage nodes are clients of the same NetWorker server.
2. Ensure that the Data Domain systems are properly licensed, including a replication license, which is required to create optimized clones.
3. Ensure that the Client resource for the NetWorker server and both storage nodes specify all of the host identifiers in the **Aliases** attribute.
   - Fully-qualified domain name
   - Short name
   - Aliases
   - IP address

   Note

   If you use an `nsrclone` command to perform an optimized clone from a host that is not the NetWorker server, then you must specify the primary hostname of the NetWorker server by using the -S option. The primary hostname of the NetWorker server is the name that appears in the NMC Enterprise view. Otherwise, a regular clone might be produced instead of an optimized clone.
4. Ensure that a target pool, for example, newclonepool, has been created for Backup Clone type with the Media type required attribute set to Data Domain.
   With this setting, if a Data Domain device is not available for a clone operation in the specified target pool, then NMC displays a "Media waiting" message.
5. Ensure that the source Data Domain device is mounted and available on the source storage node.

If the source device is not mounted, then NetWorker will perform a regular, non-deduplicated clone. However, if the specified target pool is of Backup Clone type with the Media type required attribute set to Data Domain a non-deduplicated clone will not be performed.

6. Ensure that the target Data Domain device is labeled for a clone pool, and mounted on the target storage node. The pool selected for the device label operation, for example, newclonepool, must be of Backup Clone pool type.
7. Verify that the target clone pool, for example, newclonepool, is properly specified or selected:
   - For CLI clone operations, use the `nsrclone -b newclonepool` command.
   - For the clone action, in the Destination pool attribute of the Action resource, select **newclonepool**.
   - For clones of entire volumes, Cloning by pools on page 67 provides details.

Cloning by pools

To copy save sets from Data Domain storage to a Data Domain device, you must specify a pool. This pool is known as a "clone pool." A clone pool must be assigned to a device on the target Data Domain system, where it will be available for use.

There are two main purposes for a clone pool:

- To copy existing deduplicated VTL or CIFS/NFS AFTD save sets to a Data Domain device.
- To copy the existing save sets from one Data Domain device to another Data Domain device, typically at a remote location for disaster recovery purposes.

DD Boost clone and replication support

For additional data protection, you can use the NetWorker clone feature to copy save sets on a DD Boost device to a different location. A clone is a complete and independent copy of the data that you can use for data recovery or to create additional clones. You can clone single save sets or the entire volume of a DD Boost device. A clone retains the original NetWorker browse and retention policies by default.

You can configure clones to run immediately after each save set completes, or you can configure clones to run in an independently defined maintenance window after the entire policy completes in the main backup window.

Clone formats

The type of NetWorker clone you produce depends on the type of storage media you use for the clone. NetWorker will use either CCR when cloning to DD Boost devices or a normal clone when cloning to conventional storage media.

CCR format

When NetWorker clones data from a source DD Boost device to a target DD Boost device, usually at a geographically distant location, the operation uses CCR, also known as optimized clone or DD format clone. CCR is a fast process that uses low bandwidth, multiple parallel sessions, and low storage capacity. You can use CCR clones for data recovery or to create additional copies with minimal impact on the primary operations of production, backup, or recovery.

CCR operations use only IP connectivity between DD Boost devices on separate Data Domain systems, whether you have configured the participating devices for FC or IP.

For CCR operations on the same Data Domain system, EMC recommends that you replicate the data between two different SUs (MTrees), so you can apply different retention policies and manage the data independently. When you perform CCR operations to disks that reside within the same Data Domain system, CCR uses fast copy operation.

During the CCR process, the storage node reviews the incoming clone for data that NetWorker has already stored on the target DD Boost device. The storage node stores only unique data that does not exist on the device.

Normal clone format

When NetWorker clones data from a DD Boost device to conventional media, for example, or tape, the data reverts to the non-deduplicated format. This procedure creates a normal clone. The normal clone format is necessary for the data on conventional disk or tape storage to be fully recoverable, for example, for disaster recovery, without the need of a Data Domain system.

Native Data Domain replication considerations

EMC recommends that you do not use native Data Domain replication operations to clone data. Native replication is normally used to copy deduplicated data stored in CIFS, NFS, or VTL formats from one Data Domain system to another for disaster recovery purposes. Native replication clones data independently of NetWorker and DD Boost, and the NetWorker software cannot track or control native replication operations.

An exception would be to seed a new Data Domain system by collection replication to assist the migration of existing data. Migration versus native Data Domain replication on page 138 provides details.

Note

If you use Data Domain replication for non-DD Boost directories on the same system, ensure that the system and the network has enough capacity to enable NetWorker CCR operation with DD Boost devices.

Before you use native Data Domain replication with DD Boost devices review the following information:

- Directory replication (MTree replication) does not support DD Boost devices.
- Collection replication, which is the replication of the entire stored contents of a Data Domain system, renders DD Boost devices as read-only. This operation will replicate all DD Boost devices and the stored data onto a target Data Domain system. You cannot use the replicated DD Boost data for other replication operations, such as NetWorker CCR.

  Note

  When you perform a collection replication of a Data Domain system, the NetWorker software is not aware of any DD Boost devices on that system. Additional procedures, tests, and qualifications are required to configure NetWorker to detect the devices and enable data recovery of the replicated DD Boost data. Contact EMC Professional services for assistance.

Automated Multi-streaming (AMS)

The AMS feature improves cloning performance for large savesets when you use high bandwidth networks. Previously when you replicated savesets between two Data Domain devices on different machines, the replication process used to take longer in NetWorker. AMS significantly speeds up replication between DDRs by splitting up large files (files whose sizes are roughly greater than 3.5 GB) into multiple smaller 2 GB slices, replicating the slices individually, and finally re-creating the original large file on the destination DDR using those slices.

NetWorker 8.2.3 and NetWorker 9.0.1 and later features enhancements for clone-controlled replication (CCR), also known as DD to DD Managed File Replication. Also, enhancements to load balancing so that the load (save sets to clone) is spread evenly across the multi-threaded `nsrclone` process were implemented.

By default the AMS feature is disabled. You can turn on the feature by changing the command to ams_enabled=**`yes`** An example of how you can enable AMS is below:

```
racdd098:/nsr/debug # cat nsrcloneconfig
max_total_dd_streams=256
ams_enabled=yes
ams_slice_size_factor=31
ams_preferred_slice_count=0
ams_min_concurrent_slice_count=1
ams_max_concurrent_slice_count=20
max_threads_per_client=256
ams_force_multithreaded=yes
```

Note

Both Data Domains should be connected through a 10GB network.

Configuring the CCR environment

This section describes how to configure the network environment for CCR

Before you begin

- To use Data Domain encryption, global compression, or low-bandwidth optimization, enable these configurations on both the source and target Data Domain systems.

  Note

  If any of these configurations do not match on both the source and target Data Domain systems, clone operations will fail.

  The *EMC Data Domain Operating System Administration Guide* provides details on these settings.
- On the NetWorker server and both storage nodes, configure the Client resource **Aliases** field on the **Globals 1 of 2** tab with a list of all the names and aliases in use for the CCR. Include the fully qualified name, short name, aliases, and IP address.
- Select or create a target pool for the CCR, configured for **Backup Clone** type with the **Media Type Required** field set to **Data Domain**.
  If a DD Boost device that is targeted by the pool is not available during a CCR, and the media type required specifies Data Domain, then NMC displays a `Media Waiting` message.

Note

Do not use the Default Clone Pool. You cannot change the **Media type required** setting.

Creating pools to target DD Boost devices on page 109 provides details.

Complete the following steps to configure the network environment for CCR:

Procedure

1. Ensure that you have enabled valid licenses to the Data Domain systems that you will use for CCR operations, including a Replication license. Configuring the Data Domain system for DD Boost on page 96 provides details.
2. Ensure that the source and destination storage nodes are within the same datazone. A single NetWorker server must manage the clone operations and maintain the retention policies for all cloned copies. The server also monitors and reports on the storage operations.
3. Ensure Ethernet IP connectivity between the source and destination Data Domain systems. CCR occurs only over TCP/IP connectivity. If a DD Boost device participating in the CCR also has an FC connection, ensure IP access to the DD Boost device.
4. Ensure that you map the Data Domain FC server name to the IP address, if the Domain FC and IP hostnames differ.

   Note

   Do not use connections with ifgroup links for clone operations.

   Data Domain FC and IP hostnames are the same by default but they can be different. If they are different you must map the host Data Domain FC server name to its own IP address as follows:

   a. Open the Data Domain Enterprise Manager, and navigate to the **Data Management › DD Boost**. The Data Domain **Server Name** appears on the **Fibre Channel** tab.

      Alternatively, type the following command:

      ```
      ddboost fc dfc-server-name show
      ```

   b. Associate this server name to the IP address in the `/etc/hosts` file with the following command:

      ```
      net hosts add fc_server_name IP_address
      ```

      For example, if the Data Domain system has the IP address 10.99.99.99 and the IP hostname dd555-5.lss.mcm.com, and the DFC server name is dd-tenendo, then type the following command:

      ```
      net hosts add dd-tenendo 10.99.99.99
      ```

5. Mount the source DD Boost device on the source storage node.
6. Mount the target DD Boost device on the target storage node. The pool for the device must specify **Backup Clone** pool type.
7. Ensure that the target clone pool is properly specified for the clone method you use. Clone save sets will be written to this pool. You may need to use multiple or mixed approaches for control and flexibility.

   The following example use myccrpool as the name of a clone pool you created:

- For CLI clone operations, you would use the command `nsrclone -b myccrpool`. Cloning with nsrclone on page 93 provides details.

Strategies for cloning

Scheduled cloning occurs through configuration of data protection policies. You can configure cloning to occur concurrently or after a backup, as part of a single workflow. The decision of whether to clone data immediately after a backup or as a separate workflow depends on specific circumstances, such as the amount of resources that are required for the backup and recovery time objective.

You can use a clone action in one of the following ways:

- After a backup action in a backup workflow.
- In a separate workflow.
- As the head action in workflow that uses a Query or Save set protection group.

Note

The Backup Data Management chapter describes how you can clone save sets manually by using the `nsrclone` command.

The *EMC NetWorker Cloning Integration Guide* provides details on scheduling considerations for cloning.

Road map for configuring a new cloning data protection policy

This road map provides a high level overview of how to configure a new policy for clone operations.

Before you begin

Configure the backup policy to back up the data that is cloned.

Procedure

1. Create a group to define the data to clone.
2. Create a policy. When you create a policy, you specify the name and notification settings for the policy.
3. Within the policy, create a workflow. When you create a workflow, you specify the name of the workflow, the schedule for running the workflow, notification settings for the workflow, and the protection group to which the workflow applies.
4. Create one or more clone actions for the workflow.

Protection groups for a cloning workflow

You can use two types of protection groups to clone save sets in a workflow that are separate from backup workflows. The type of protection group that you use depends on the way that you plan to configure the workflow.

Use a save set group or a query group to specify a list of save sets if cloning occurs as the head action in a cloning workflow:

- Save set group—Use a save set group in clone-only workflows where you want to clone a specific list of save sets. Save set groups are similar to the manual clone operations in NetWorker 8.2.x and earlier.

- Query group—Use a query group in clone-only workflows where you want to clone save sets on an ongoing basis, based on the save set criteria that you define. Query groups are similar to the scheduled clone operations in NetWorker 8.2.x and earlier.

Note

To clone save sets in a backup workflow, use basic client group or a dynamic client group. Strategies for traditional backups on page 41 provides detailed information about how to create clone actions in a traditional backup workflow.

Create multiple protection groups to perform cloning in different ways as part of separate workflows, or to perform cloning for different save sets on different schedules. For example:

- Create a basic client group for a workflow that performs a traditional backup of the a client file system followed by cloning of the save sets that result from the backup.
- Create a query group that identifies full save sets in the last two days to clone.

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 8 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.
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.

Table 8 Save set criteria (continued)

Criteria	Description
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**.

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.

Workflows for scheduled cloning

A workflow can contain one or more clone actions.

Supported workflow path from a clone action

Another clone action is the only supported action after a clone action.

Figure 15 Workflow path from a clone action



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 9 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.

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

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

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

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

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

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

      Note

      The **Abort action** option 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 a clone workflow

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.

The following figure illustrates the visual representation of a clone workflow.

Figure 16 Visual representation of a clone 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.

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

- 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**.

Road map to add a clone workflow to an existing policy

This road map provides a high level overview of how to create a clone workflow and add the workflow to an existing backup policy.

Before you begin

Configure the backup policy to back up the data that is cloned.

Procedure

1. Create a query or save set group to define the data to clone.
2. Add the new group to an existing policy.
3. Create a workflow in the existing policy.
4. Create one or more clone actions for the workflow.

Example: Creating a policy that has a separate workflow for cloning

The following figure provides a high level overview of the configuration of a policy that contains two workflows, one for backups and one to clone a list of save sets.

Figure 17 Example of a policy with separate workflows for backup and cloning



Note

The amount of data and length of time that is required to complete the backup can impact the ability to clone data when the backup and clone workflows are in the same policy. For example, if the clone action starts before the backup action completes, there might not be any data yet to clone, or in other cases, only the save sets that completed at the start time of the workflow is taken into account. In both cases, NetWorker marks the Clone Workflow as successful, but there is no guarantee that all the data from the backup workflow was cloned.

Protection groups for a cloning workflow

You can use two types of protection groups to clone save sets in a workflow that are separate from backup workflows. The type of protection group that you use depends on the way that you plan to configure the workflow.

Use a save set group or a query group to specify a list of save sets if cloning occurs as the head action in a cloning workflow:

- Save set group—Use a save set group in clone-only workflows where you want to clone a specific list of save sets. Save set groups are similar to the manual clone operations in NetWorker 8.2.x and earlier.
- Query group—Use a query group in clone-only workflows where you want to clone save sets on an ongoing basis, based on the save set criteria that you define. Query groups are similar to the scheduled clone operations in NetWorker 8.2.x and earlier.

Note

To clone save sets in a backup workflow, use basic client group or a dynamic client group. Strategies for traditional backups on page 41 provides detailed information about how to create clone actions in a traditional backup workflow.

Create multiple protection groups to perform cloning in different ways as part of separate workflows, or to perform cloning for different save sets on different schedules. For example:

- Create a basic client group for a workflow that performs a traditional backup of the a client file system followed by cloning of the save sets that result from the backup.
- Create a query group that identifies full save sets in the last two days to clone.

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 10 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**

Table 10 Save set criteria (continued)

Criteria	Description
	• **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.
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**.

Editing an existing policy to create a workflow and clone action

Use the **Policies** window to create a workflow and create the clone action.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, expand **Policies**, and then select the existing policy.
3. In the right pane, right-click in the workflow section and select **New**, and select **Properties**.

   The **New Workflow** dialog box appears.
4. In the **Name** field, type the name of the workflow.

   The maximum number of characters for the name of the group is 64.
5. In the **Comment** box, type a description for the workflow. The maximum number of characters for the **Comment** field is 128.
6. From the **Send Notifications** list, select how to send notifications for the workflow:
   - To use the notification configuration that is defined in the policy resource to 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**.
7. 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.

8. 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.
9. In the **Groups** group box, specify the protection group to which the workflow applies.

   To use a group, select a protection group from the **Groups** list. To create a protection group, click the **+** button that is located to the right of the **Groups** list.
10. Click **Add**.

    The Policy Action Wizard appears.
11. In the **Name** field, type the name of the action.

    The maximum number of characters for the action name is 64.
12. In the **Comment** field, type a description for the action.
13. 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.

14. From the **Action type** list, select **Clone**.
15. 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.
16. 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**.
17. Click the icon on each day to specify whether to perform cloning.

    The following table provides details on the icons.

    Table 11 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**.

18. Click **Next**.

    The **Specify the Clone Options** page appears.
19. 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.

20. 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.

21. Click **Next**.

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

22. 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.

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

    Note

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

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

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

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

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

24. 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.

25. 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.
26. 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.
27. 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.
28. 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.
29. (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.

30. Click **Next**.

    The **Action Configuration Summary** page appears.

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

Clone reports

You can use the NMC Reports view to access reports of NetWorker clone operations on a Data Domain system. Generating reports on page 125 provides details.

Cloning with nsrclone

Use the `nsrclone` command to configure detailed CCR operations that you can launch by either running the `nsrclone` command manually or by scheduling a task on the operating system or an external scheduler that runs the `nsrclone` command. This method is best for larger environments where flexibility and control of conditions are necessary. Some examples are as follows:

- Start clone job B, which clones to tape storage, if and only if clone job A, which performs CCR, has successfully completed.
- Clone only specific save sets to specified storage nodes or specified devices.
- Perform a CCR from a host other than the NetWorker server or a NetWorker storage node. This command must specify the NetWorker server by its primary hostname. Use the hostname listed in the NMC Enterprise view. Otherwise, the CCR operation may produce normal clones instead.

Scripted solutions require additional knowledge and have external dependencies, such as operating systems and scripting tools, that are beyond the scope of this guide.

The *EMC NetWorker Administration Guide* and the *EMC NetWorker Command Reference Guide* provide more details. EMC Professional Services can also provide assistance.

CHAPTER 4

Software Configuration

This chapter includes the following topics:

- DD Boost configuration road map 96
- Configuring the Data Domain system for DD Boost 96
- Configuring NetWorker for DD Boost devices 99
- Configuring DD Boost devices with the wizard 100
- Configuring DD Boost devices manually 103
- Configuring clients to back up to DD Boost devices 111
- Deactivating and removing DD Boost devices 115

DD Boost configuration road map

Plan the DD Boost configuration with a high-level road map that outlines the sequence of basic configuration tasks that you must perform.

1. Configure the Data Domain system for use with NetWorker.
   Configuring the Data Domain system for DD Boost on page 96 provides details.
2. Configure NetWorker for use with the Data Domain system by using either the Client Configuration Wizard or manually create the client resource.
   Configuring NetWorker for DD Boost devices on page 99 provides details.
3. Configure NetWorker clients to back up to the Data Domain system.
   Configuring clients to back up to DD Boost devices on page 111 provides details.

Configuring the Data Domain system for DD Boost

Enable the Data Domain system for storage operations with DD Boost devices by using the Data Domain CLI to complete the following steps. You can also use the Data Domain System Manager as an alternative method for most of these steps.

The *EMC Data Domain Boost for OpenStorage Administration Guide* provides details.

Procedure

1. Ensure the network requirements. Network requirements on page 25 provides details.
2. Log in to the Data Domain system, log in as an administrative user, and then run the Data Domain CLI.
3. Use the `filesys` command to verify that the file system is running and the `nfs` command to verify that the NFS services are running, type the following commands:

   ```
   filesys status
   ```

   ```
   nfs status
   ```

   If NFS services are not running, type the following command:

   ```
   nfs enable
   ```

   Note

   For DD Boost functionality, you must enable NFS services on the Data Domain system, even if you do not configure users or shares are configured. You do not need to enable NFS on the NetWorker server, NetWorker storage nodes, or NetWorker clients.

4. (Optional) Use the `system show` command to verify the installed version of DD OS:

   ```
   system show version
   ```

   Note

   You cannot configure SMT on releases before DD OS 5.5, where only one DD Boost user could have access to the Data Domain system at a time.

5. Use the `user add` command to create one or more new user account. For SMT, you will assign the *none* role to the user:

   ```
   user add username password password [role role]
   ```

For example, to create three SMT user accounts named bob, joe and sue, type the following commands:

```
user add bob password mypwbob role none

user add joe password mypwjoe role none

user add sue password mypwsue role none
```

6. Use the `ddboost user` command to assign the new users as DD Boost users:

   ddboost user assign *username-list*

   where *username-list* is a comma separated list of usernames.

   For example:

   ```
   ddboost user assign bob, joe, sue
   ```

   Note

   To unassign one of more users from the DD Boost user list, type:

   ddboost user unassign *username-list*

7. (Optional) Use the `smt` command to enable SMT, create tenant units (TUs), and assign a default TU to each DD Boost user.

   Note

   You must log in to the Data Domain system as a user account that has the role of Global Storage Administrator to run the `smt` commands.

   a. To enable SMT, type:

      ```
      smt enable
      ```

   b. To create one or more TUs, type:

      smt tenant-unit create *tu-name*

      For example, to create a two TUs named tu1 and tu2, type:

      ```
      smt tenant-unit create tu1

      smt tenant-unit create tu2
      ```

   c. To assign a default TU to a DD Boost user, type:

      ddboost user option set *username* default-tenant-unit *tu-name*

      Note

      A DD Boost user can have only one default TU, but multiple DD Boost users can share the same default TU.

      For example, to assign TU tu1 to DD Boost users bob, type:

      ```
      ddboost user option set bob default-tenant-unit tu1
      ```

      To assign TU tu2 to DD Boost users joe and sue, type the following commands:

```
ddboost user option set joe default-tenant-unit tu2

ddboost user option set sue default-tenant-unit tu2
```

Because you assign a default TU to each DD Boost user, NetWorker automatically associates any storage unit (SU) created by a DD Boost user to their default TU. NetWorker does not expose the TUs.

- (Optional) To unassign a DD Boost user from its default TU, or to reassign the user to a different default TU, type the following command:

  ddboost user option reset *username* [default-tenant-unit]

  Note

  Avoid changing the owners of DD Boost SUs. A new owner cannot use the DD Boost devices from a previous owner. Create a device for the new owner instead.

- (Optional) To list the DD Boost users and their default TUs, or the DD Boost users within a specific default TU, type the following command:

  ddboost user show [default-tenant-unit *tenant-unit*]

  Note

  You can use the CLI to review tenant space usage and the performance data at both the TU and SU levels. As the global storage administrator, you can enable tenants to use the Data Domain CLI to review the space usage and the performance data of their TU and SUs. The Data Domain documentation provides details.

8. Use the `license add` command to add the OPENSTORAGE license key, and optionally, if you will use CCR, to add the Replication license key:

   license add *license_key*

9. Use the `ddboost` command to restart the DataDomain service, which applies the system modifications:

   ```
   ddboost disable

   ddboost enable
   ```

10. Use the `snmp` command to configure the system to receive and generate SNMP traps. SNMP traps enable users to monitor backup events captured by SNMP traps.

    snmp add ro-community *community_name*

    snmp enable

    snmp add trap-host *hostname*[:*port*]

    The typical value for *community_name* is “public”, which allows all users to monitor events.

11. Use the `ddboost option set` command to enable and verify Distributed Segment Processing (DSP). This is required for deduplication operations on the NetWorker storage nodes and NetWorker clients:

    ddboost option set distributed-segment-processing enabled

    To confirm that DSP is enabled, type **ddboost option show**.

12. If you plan to enable Fibre Channel (FC) connectivity on DD Boost devices, use the `ddboost fc` command to obtain the DD Boost over FC (DFC) server name.

    `ddboost fc dfc-server-name show`

    Note

    You will specify the DFC server name in the NetWorker device configuration procedure. FC-enabled clients can back up only to FC-enabled devices. IP-enabled clients can back up only to IP-enabled devices.

    For example, in the following output, the DFC server name is dd-tenendo:

    `ddboost fc dfc-server-name show`

    ```
    DDBoost dfc-server-name: dd-tenendo

    Configure clients to use "DFC-dd-tenendo" for DDBoost FC
    ```

    Note

    Do not use the "DFC-" prefix on the DFC server name, as suggested in the output of the `ddboost fc dfc-server-name show` command. This prefix is intended for use with other vendors only and will cause NetWorker communications to the DFC server to fail.

13. Use the NetWorker Device Configuration Wizard to create DD Boost devices and the Data Domain SU folders that contain the devices.

    Configuring DD Boost devices with the wizard on page 100 provides details.

Configuring NetWorker for DD Boost devices

After you configure a Data Domain system for the DD Boost environment, you can configure the NetWorker resources for devices, media pools, volume labels, clients, and groups that will use the DD Boost devices. Keep the following NetWorker considerations in mind:

- Each DD Boost device appears as a folder on the Data Domain system. A unique NetWorker volume label identifies each device and associates the device with a pool.
- NetWorker uses the pools to direct the backups or clones of backups to specific local or remote devices.
- NetWorker uses Data Protection policy resources to specify the backup and cloning schedules for member clients. EMC recommends that you create policies that are dedicated solely to DD Boost backups.
- VMware Backup Appliance does not support the SMT feature. The *EMC NetWorker VMware Integration Guide* provides details.

DD Boost device performance considerations

NetWorker does not limit the number of DD Boost devices that you can create. The number of required devices depends on device usage for backup and restore operations.

Increasing the number of DD Boost devices can impact Data Domain performance and maintenance. Typically, if you do not need multiple concurrent sessions or streams for recovery, then you can configure the device Target Sessions and Max Sessions settings

for multiple concurrent backup sessions. Avoid the removal of DD Boost devices. Deleting an AFTD or DD Boost device on page 115 provides details.

Methods to configure DD Boost devices

EMC recommends that you use the Device Configuration Wizard, which is part of the NetWorker Administration GUI, to create and modify DD Boost devices. The wizard can also create and modify volume labels and the storage pools for DD Boost devices.

After the wizard creates a DD Boost device, you can modify the device configuration by editing the Device resource that is created by the wizard.

Configuring DD Boost devices with the wizard

Before you begin

- Ensure that the network requirements are met. Network requirements on page 25 provides details.
- Review DD Boost storage characteristics. DD Boost storage structure and limits on page 22.

Complete the following steps to create or modify a DD Boost device with the Device Configuration Wizard.

Procedure

1. Use the NMC GUI to connect to the NetWorker server.
2. In the NetWorker **Administration** window, click the **Devices** tab.
3. In the left panel, right-click **Data Domain Systems** and select **New Device Wizard**.

   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 the **Data Domain** device type, and then click **Next**.
5. On the **Data Domain Preconfiguration Checklist** page, review the requirements, and then click **Next**.

   Configure the Data Domain system for NetWorker and define a DD Boost username. Configuring the Data Domain system for DD Boost on page 96 provides details.

6. On the **Specify the Data Domain Configuration Options** page, configure the following attributes:

   a. In the **Data Domain System** section, select one of the following options:

      - To use a Data Domain system on which you have previously created devices, select the **Use an existing Data Domain System** attribute.
      - To use a new Data Domain system, select the **Add a new Data Domain System** field, and then specify the FQDN or IP address of the Data Domain system.

Note

If you will use DFC connectivity, Do not use the "DFC-" prefix on the DFC server name, as suggested in the output of the `ddboost fc dfc-server-name show` command. This prefix is intended for use with other vendors only and will cause NetWorker communications to the DFC server to fail.

b. In the **DD Boost Credentials** section, provide the **DD Boost Username** and the **DD Boost Password** that you created on the Data Domain system.

c. In the **Secure Multi-Tenancy** area, to use only DD Boost devices in secure Storage Units (SUs), select **Configure Secure Multi-Tenancy (SMT)**, then perform one of the following tasks:

- To use an existing SU, select **Use an existing secure storage unit**, and then select the SU.
- To create a SU, select **Create a new secure storage unit**, and then specify the name of the SU.

Note

SMT restricts access of each SU to one owner according to the provided DD Boost credentials.

d. In the **Configuration Method** field, select **Browse and Select**, and then click **Next**.

Note

If you do not configure the SMT option, the wizard will create an SU for you on the Data Domain system, and name the SU after the shortname of the NetWorker server.

The next wizard page will enable you to browse the specified storage unit on the Data Domain system and select the folders that NetWorker will use as DD Boost devices.

7. On the **Select Folders to use as Devices** page, create a DD Boost device as follows:

a. Select the Data Domain system, and then click **New Folder**.

A new folder appears in the navigation tree. This folder is the new device.

Note

The navigation tree does not show the SU folder under the Data Domain system folder. However, the SU folder is verifiable in the final Review Configurations Settings wizard page. The wizard names the SU folder after the short hostname of the NetWorker server and places the devices that you create into this SU folder.

b. Type a name for the new device, and then select the checkbox next to the folder or device name.

The Device table displays the full NetWorker device name, the storage pathname, and details about the device.

Note

The device name refers to the subfolder created within the SU. The folder path must not contain other folders deeper than these device folders.

c. (Optional) To rename a DD Boost device as it appears in NMC, select the device in the table, and type a new name in the **NetWorker Device Name** field. Do not use special characters other than dot (.) and underscore (_). The **Storage Path** field remains unchanged.

Note

Implicit in the SU folder pathname on the Data Domain system are the hidden mount point folders: `/backup/ost` (DD OS 4.9) or `/data/col1` (DD OS 5.0 and later). Do not modify this folder structure, which all NetWorker server hosts use. The final wizard page, **Review Configurations Settings**, shows the complete location.
The `/backup` folder stores NFS service data and the clients that are configured for NFS access also have the ability to view, change, and delete the `/data/col1` directory that contains the DD Boost devices. If you use NFS devices, you can avoid the risk of potential interference by using alternative path names.

d. Click **Next**.

8. On the **Configure Pool Information** page, specify the following settings:
   a. Select the **Configure Media Pools for Devices** option.
   b. In the **Pool Type** section, select the type of data to send to the Data Domain device, either **Backup** for backups or **Backup Clone** for cloning or staging operation.
   c. In the **Pool** section, select **Create and use a new Pool** to create a pool to receive the data, or select **Use an existing Pool** to select a pool that exists on NetWorker server.
   d. Select the **Label and Mount device after creation** option.
   e. Click **Next**.
9. On the **Select Storage Nodes and Fibre Channel Options** page, specify the following settings:
   a. In the **Storage Node Options** section, specify the storage node that will manage the device. Select **Use an existing storage node** to select an existing storage node on the NetWorker server, or select **Create a new storage node** and specify the hostname of a storage node host. If the new Storage Node is also a dedicated storage node, select **Dedicated Storage Node**.
   b. Optionally, specify the following settings to enable FC data transport to this device:
      - Select **Enable Fibre Channel** for this device.
      - In the **Fibre Channel Host Name** field, type the hostname that the Data Domain system uses to identify itself for FC operations. By default, this hostname is the same name used for IP operations, but the hostnames can be different. The hostname must match the Server Name displayed on the Data Domain system in the **Data Management** › **DD Boost** › **Fibre Channel** tab of the **Data Domain Enterprise Manager**. The name is case-sensitive. Configuring the Data Domain system for DD Boost on page 96 provides details.

Note

All NetWorker clients that use an FC-enabled DD Boost device must be enabled for FC in the **Data Domain Interface** field.

c. Click **Next**.

10. On the **SNMP Monitoring Options** page perform the following steps:

    a. In the **Data Domain SNMP Community String** field, type the name of the SNMP community string.

    Note

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

    b. With the **Receive SNMP trap events** option selected, specify the SNMP Process port used by the Data Domain system and select the events in which to monitor. **Use the Reset** to defaults option to reset the events in which to monitor back to the default settings.

    Note

    The default SNMP process port is 162.

    c. Click **Next**.

    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 Configuration Settings** page, and then click **Configure**.

    Note

    The name listed as the SU is really the pathname for the device folder. The format is: *SU|device_name*, where *SU* is the short hostname of the NetWorker server.

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

12. On the **Device Configuration Results** page, review the information, and click **Finish** to exit the wizard.

Results

The **Data Domain Systems** window displays the new Data Domain device and the name of the volume.

Configuring DD Boost devices manually

EMC recommends that you use the Device Configuration Wizard to manually add a Data Domain system to the NetWorker datazone. and create DD Boost devices. You can modify

the Device resource that the wizard creates to modify the devices, and perform the tasks in the following sections.

Adding a host Data Domain system to NMC Enterprise view

Use the Add New Host Wizard to manually add a Data Domain system to the NetWorker datazone.

NetWorker lists the Data Domain systems as a host in the NMC Enterprise view. This view shows the Data Domain system status and the backup and recovery events that were performed by NetWorker managed by NMC. The **Enterprise** view also provides a live link to launch the Data Domain Enterprise Manager GUI. To manually add a Data Domain system to the NMC **Enterprise** view, perform the following steps:

Procedure

1. Ensure the network requirements. Network requirements on page 25 provides details.
2. From the **File** menu, select **New** › **Host** to run the **Add New Host** wizard.
3. Complete the wizard screens:
   - Type the Data Domain hostname.
   - Select **Data Domain**.
   - Select **Capture Events**.
   - Type the name of the SNMP community where NMC will retrieve Data Domain status information. By default, NMC uses the value configured on the Data Domain system with the `snmp add ro-community` command. Configuring the Data Domain system for DD Boost on page 96 provides details.
   - Type a value for the **SNMP Process Port**. By default, NMC uses the value that is configured on the Data Domain system with the `snmp add trap-hosthostname[:port]` command. Configuring the Data Domain system for DD Boost on page 96 provides details. This configuration must agree with the firewall configuration on the Data Domain system. Firewall requirements on page 27 provides details.
   - Select the **SNMP Traps** that you want to monitor.

   Configuring SNMP for an NMC managed Data Domain system on page 124 provides details on monitoring and alerts.

Configuring a DD Boost device manually

EMC recommends that you create a DD Boost device by using the Device Configuration Wizard. A DD Boost device appears as a folder on the Data Domain system. You associate each DD Boost device with a single NetWorker volume by labeling a device for a NetWorker pool.

Note

If you manually create a device with this procedure, NMC lists the device but this procedure does not create a corresponding device folder on the Data Domain system. If you try to label and mount such a device, an error appears.

To modify a DD Boost device complete the following steps:

Procedure

1. Use NMC to connect to the NetWorker server. In the **Administration** window, click the **Devices** view.

2. In the folder tree, expand **Data Domain Systems** and select the Data Domain system that stores the save sets.
3. In the right panel, right-click the name of the device that you want to modify, and then select **Properties**.
4. On the **General** tab, identify the DD Boost device by typing its name and access information:

Note

Multiple devices can share a single volume. Configuring volume sharing on multiple devices on page 107 provides details.

---

a. In the **Name** field, type a name for the Data Domain device.

For example:

dd_1

Configuring a DD Boost device manually on page 104 uses the following example values:

- NetWorker server short hostname = *dzone1*
- NetWorker remote storage node hostname = *dzone1_sn2*
- Data Domain hostname = *ddr1*
- DD Boost device name = *dd_1*

If you configure the device on a separate storage node host that is not the NetWorker server host as shown in Configuring a DD Boost device manually on page 104, it is a remote device. Specify the **Name** field in the following format:

rd=*remote_storagenode_hostname:device_name*

For example:

rd=*dzone1_sn2:dd_1*

b. In the **Device access information** field, type the Data Domain hostname followed by a colon and the path to the device folder.

If you are configuring a device with secure multi-tenancy (SMT) protection, the device folder must reside in a password-protected tenant unit on the Data Domain. Configuring the Data Domain system for DD Boost on page 96 provides details.

Use the following format:

DD_hostname:/*DD_storage_unit_name*/*device_name*

where, as a best practice, *DD_storage_unit_name* is the short hostname of the NetWorker server and *device_name* is a name for the device, which appears as a folder.

For example, the following figure uses the following name:

ddr1:/dzone1/dd_1

NetWorker does not limit the number device folders that you can create, but the Device access information field accepts one device folder only. Do not create any folders within a device folder.

Note

Implicit in this pathname are the hidden mount point folders, `/backup/ost` (DD OS 4.9) or `/data/col1` (DD OS 5.0 and later.) Do not modify this folder structure, which All NetWorker servers use.

Figure 18 Example of the device name and the access information for a DD Boost device



c. In the **Media type** field, select **Data Domain** from the list.

5. On the **Configuration** tab, in the **Save Sessions** area, set the number of concurrent save sessions (streams) and the number of `nsrmmd` (media storage) processes that the device can handle:

- In the **Target sessions** field, specify the number of save sessions that a `nsrmmd` process on the device handles before for another device on the Data Domain host takes the additional sessions. If another device is not available, then another `nsrmmd` process on the same device takes the additional sessions. Use this setting to balance the sessions load among `nsrmmd` processes.

  EMC recommends that you set this field to a low value. The default value is 6. The maximum value is 60.

- In the **Max sessions** field, specify the maximum number of save sessions that the device can handle. At the maximum limit, if no additional devices are available on the host, then another available Data Domain system takes the additional sessions. If no other Data Domain hosts are available, then the system retries the save sessions until a `nsrmmd` process become available.

  The default value is 60. The maximum value is 60.

  Note

  The **Max sessions** setting does not apply to concurrent recovery sessions.

- In the **Max nsrmmd count** field, specify the maximum number of `nsrmmd` processes that can run on the device. Use this setting to balance the `nsrmmd` load among devices.

  If you enabled Dynamic nsrmmds on the storage node, NetWorker automatically adjusts this value by using the formula max/target +4, with the default value being 14. Otherwise, the default value is 4.

  To modify this value, first adjust the two sessions fields, apply and monitor the effects, and then tweak the **Max nsrmmd count** value.

Note

NetWorker reserves at least one `nsrmmd` process for restore and clone operations.

6. In the **Remote user** and **Password** fields, type the DD Boost username and password, respectively.

   You can only define one DD Boost (OST) user. All NetWorker storage nodes and servers that access the Data Domain system must use the same username and password.

   Note

   Avoid changing the user of an existing device with a labeled volume. The new user will not have write permission to the files and directories that are created by the previous user and cannot re-label the volume. Create a device for the new user.

7. If you want the DD Boost device to use FC connectivity, complete the following steps:

   a. Select the **Enable fibre channel** field.

   b. In the **Fibre Channel Host Name** field, type the hostname that the Data Domain system uses to identify itself for FC operations. By default, this hostname is the same name used for IP operations, but the hostnames can be different. The hostname must match the Server Name displayed on the Data Domain system in the **Data Management** › **DD Boost** › **Fibre Channel** tab of the **Data Domain Enterprise Manager**. The name is case-sensitive.

      Configuring the Data Domain system for DD Boost on page 96 provides details.

      Note

      All NetWorker clients that use an FC-enabled DD Boost device must be enabled for FC in the **Data Domain Interface** field.

8. To save the device settings click **OK**.

   The NetWorker **Administration** window displays the Data Domain system and details of the device.

9. Ensure that the device is associated with a NetWorker storage volume before you try to use the device. Otherwise, an error appears. Labeling and mounting devices on the storage node on page 111 provides the procedure.

Configuring volume sharing on multiple devices

You can concurrently mount and share a single NetWorker storage volume with multiple DD Boost devices, to provide greater flexibility and performance gains.

- A volume that you simultaneously mounted on both an IP-enabled DD Boost device and an FC-enabled DD Boost device provides greater flexibility. Clients, including Client Direct clients, can back up and restore their data on the same volume over either IP or FC networks. Restoring by Client Direct over IP from an FC-enabled device on page 119 provides details on a volume sharing solution for restore operations.
- You can create multiple devices for shared volumes on the same storage node or on separate storage nodes.

- For clients that are not Client Direct clients, a shared volume can improve bandwidth for backup or restore operations to a DD Boost device because NetWorker can use the storage node that is closest to the requesting client.

  Note

  In some environments however, concurrent read or write operations to a volume from multiple storage nodes or Client Direct clients can result in disk thrashing that impairs performance.

You must create each device separately, with a different name, and you must correctly specify the path to the storage volume location.

For example, to create three devices, one on the NetWorker server host named dzone1 that uses local devices and two remote devices (rd) on storage nodes dzone1_sb2 and dzone1_sn3, specify the name of each device in **Name** field of each device as follows:

dd_1a

rd=dzone1_sn2:dd_1b

rd=dzone1_sn3:dd_1c

The **Device access information** field would specify the same single directory as a valid complete path for each alias.

For example, for a directory named dd_1 on the Data Domain storage host named ddr1, specify the correct pathname:

- If the storage node uses an automounter, you can specify the following pathname:

  `/net/ddr1/dzone1/dd_1`

- If the storage node uses an explicit system mount point, you can specify one of the following pathnames:

  `/mnt/ddr1/dzone1/dd_1`

  `/mnt/dzone1/dd_1`

Creating a volume label template for DD Boost devices

When you use the Device Configuration Wizard, the wizard automatically creates a label template for the volumes that a new device will use.

Note

The Device Configuration Wizard automatically creates a label template for the volumes, and this procedure does not apply if you use the wizard.

Each DD Boost device is associated with a single volume. The label template that is assigned to the pool determines the volume name. NetWorker mounts each volume in a DD Boost device. A label template provides a DD Boost device with a volume name and numbering to all storage volumes that belong to the same pool.

A label template defines the components of a volume label, which includes the volume name, a separator, and volume number. All the volumes in the same pool will have the same label name, for example, dd_myvol, but different volume numbers, for example, . 001.003.

For example, a Data Domain system may have three devices, each of which is mounted with a storage volume (Volume Name). If each device/volume is associated with the same pool, the volume names would be as follows:

- dd_myvol.001
- dd_myvol.002
- dd_myvol.003

To create a label template, perform the following steps:

Procedure

1. In the **NetWorker Administration** window, click **Media**.
2. In the browser tree, select **Label Templates**, and from the **File** menu, click **New**.

   The **Create Label Template** window appears.
3. In the **Name** and **Comment** fields, type a name and description for the label template. The label will associate a storage pool to a device.
4. In the **Fields** field, type the components of the label. Place each label component on a separate line. The template must include at least one volume number range component. NetWorker applies the label template to the volumes mounted on DD Boost devices in a Data Domain system.

   For example:

   dd_myvol

   001-999

   Host naming guidelines on page 29 provides guidelines for creating names.
5. Select a **Separator**, and click **OK**.
6. In the **Next** field, specify the next volume label in the sequence to be applied during the next label and mount operation, for example, dd_myvol.001.
7. Click **OK**.

Creating pools to target DD Boost devices

Typically, use the Device Configuration Wizard, which automatically creates a media pool. The following procedure describes the alternative manual method that uses the NMC property windows.

Each NetWorker client stores data to a media or target pool. This pool is used to direct the data from backup clients, or the data from storage volumes for clone operations, to the storage devices that are members of the pool.

Each DD Boost device is associated with a storage volume label when it is mounted. The Volume Name value of the storage volume implicitly associates the device with the specified pool.

Note

Dynamic Drive sharing (DDS) is not supported for DD Boost devices.

Complete the following steps to manually create a pool for Data Domain backups:

Procedure

1. Ensure that the devices that you assign to the pool were created in NetWorker.
2. Ensure that a label template has been created for the pool. Creating a volume label template for DD Boost devices on page 108 provides details.
3. From the **NetWorker Administration** window, click **Media**.
4. In the left navigation pane, select **Media Pools**, and from the **File** menu, select **New** to open the **Create Media Pool** window with the **Basic** tab selected.

5. In the **Name** field, type a name for each pool. Create names that clearly indicate whether the pool is for a Data Domain backup or a Data Domain clone operation.

   For example:

   DDsite1

   DDCLsite2

   A pool name that starts with DD would be a Data Domain pool, and a pool name that starts with DDCL would be a Data Domain clone pool. The pool name can also include the physical location where NetWorker stores the backup data. These conventions make the name easier to use for scripting and reporting.

   Host naming guidelines on page 29 provides guidelines for creating names.

6. (Optional) In the **Comment** field, type a description of the pool.
7. Select **Enabled**.
8. Select the **Pool type**:
   - To use the pool for backups, select **Backup**.
   - To use the pool for clone copies, select **Backup Clone**.

   Note

   You cannot modify the **Pool type** value after you create the device.

9. In the **Label Template** field, select a label template to associate with the pool.

   You can later apply the pool to DD Boost devices. Labeling and mounting devices on the Data Domain device on page 111 provides details.

10. On the **Selection Criteria** tab, under **Target Devices,** select all the DD Boost devices that this pool may use for storage. The pool may store data on any of these devices. Use the following practices:
    - Select only DD Boost devices for the pool. Do not mix DD Boost devices with other types of storage devices. If you modify a pool in this step, ensure that the pool excludes all devices that are not DD Boost devices.
    - Consider whether to include only FC-enabled or only IP-enabled DD Boost devices in the pool. When you include DD Boost devices with different connectivity, NetWorker might send data to an FC device one time, and to an IP device another time. If this is a concern, then do not mix FC and IP devices in the same pool.
    - Do not select devices that reside on more than one Data Domain system.

      Backups from a single NetWorker client can target any of these Data Domain systems. This behavior impairs the backup window and deduplication ratio.

11. Under **Media type required**, if you intend to use the pool for a Data Domain backup only, set this field to **Data Domain**. This setting ensures that only Data Domain devices use this pool.

    Note

    EMC recommends that you do not include different media types in a single pool. Backup fails for older NetWorker application modules on page 132 provides further details.

12. Click **OK**.

The *EMC NetWorker Administration Guide* provides details on media pools.

Labeling and mounting devices on the Data Domain device

EMC recommends that you use the Device Configuration Wizard to create a Data Domain device, which automatically labels and mounts the device. The following procedure describes the alternative manual method that uses the NMC property windows.

Before you can use a DD Boost device you must label and mount the device.

You can complete the following steps to label and mount a device.

Procedure

1. In the **NetWorker Administration** window, click **Devices**.
2. In the left navigation pane, select **Data Domain Systems**.
3. In the right panel, right-click the device you want to label and select **Label**.
4. In the **Label** window and **Pools** list box, select a pool to associate with the device.

   A label for the selected pool appears in the **Volume Label** field. This label will become the volume name for the device.

5. Select **Mount After Labeling** and click **OK**.

   The **Devices** window displays the device and the associated volume name.

Configuring clients to back up to DD Boost devices

You can create client resources to define backup data by using the Client Configuration wizard or manually. EMC recommends that you use the Client Configuration wizard to create client resources.

Configuring a backup client with the wizard

Use the NetWorker Client Configuration wizard to create and modify NetWorker backup clients.

Before you begin

If the client is to use a Client Direct backup, which is the default configuration, ensure that the client has access to the same network connectivity (IP or FC) that the target DD Boost devices use.

Note

Mac OS X clients only support the IP protocol.

Note

If you want to redirect existing client backups to new DD Boost devices, Redirecting backups from other devices to DD Boost on page 136 provides details.
The details for the settings referred to in this simplified wizard procedure are found in the next procedure Configuring a backup client with NMC property windows on page 112.

Procedure

1. Use NMC to connect to the NetWorker server, and then click **Protection**.

2. In the left navigation pane, right-click **Clients** and select **Client Backup Configuration › New**.
3. Complete the fields on the following wizard pages:
   - Show the **Client Name**.
   - Specify the **Backup Configuration Type**.
   - Specify the **Backup Options**:
     - In the **Deduplication** settings, select **Data Domain Backup**, if applicable. This setting ensures that the client backs up only to DD Boost devices if the pool used also contains other types of devices such as AFTDs. It is best not to have mixed devices in pools.
     - Set **Target Pool** to a pool associated with DD Boost devices. An alternative way to configure a client to use a pool is to specify the client or its group in the **Data Source** field of the Pool resource. Creating pools to target DD Boost devices on page 109 provides details.

       Note

       Current versions of NetWorker application modules support backup to DD Boost devices. Some earlier versions of modules do not support the client fields for **Data Domain backup** and **Pool**. In these cases, do not set these fields. Backup fails for older NetWorker application modules on page 132 provides details.
   - Select **Files to Back Up**.
   - Select the **Client Properties**.
   - Select the **Backup Group**.
   - Specify the **Storage Node Options**.
4. Complete the wizard.

Configuring a backup client with NMC property windows

EMC recommends that you use the Client Configuration wizard to create and modify NetWorker clients. The following procedure describes how to manually create a Client resource.

Before you begin

If the client is to use a Client Direct backup, which is the default configuration, ensure that the client has access to the same network connectivity (IP or FC) that the target DD Boost devices use.

Note

Mac OS X clients only support the IP protocol.

The *EMC NetWorker Administration Guide* provides details on NetWorker Client resource configurations.

You can complete the following steps to configure a NetWorker client for scheduled backups to a DD Boost device.

Procedure

1. Use NMC to connect to the NetWorker server and click **Protection**.
2. In the left navigation pane, select **Clients**:

- To create a Client resource, from the **File** menu, select **New**.
- To edit an existing Client resource, select the client name from the list in the right panel, and from the **File** menu, select **Properties**.

a. On the **General** tab, in the **Name** field, type the hostname for the client and, optionally, type a comment in the **Comment** field.

b. Select the **Block based backup** checkbox.

c. In the **Save Sets** field, click the **Browse** button to open the **Choose Directory** window. Browse to and select volumes or individual filesystem you want to back up. When finished selecting, click **OK**.

Type each item on a separate line. For example:

```
E:\
F:\
```

To back up all client data, type All in the **Save Sets** field.

Note

For Microsoft Windows systems, back up the SYSTEM or Volume Shadow Copy Service (VSS) SYSTEM on a periodic basis to prepare for disaster recovery of the client system.
The *EMC NetWorker Administration Guide* provides details for this step.

---

3. On the **General** tab, in the **Backup** area, complete the following steps:

a. To enable deduplicated backup data from this client to bypass the NetWorker storage node and be sent directly to the Data Domain system, select **Client Direct**. Review the following requirements:

- Ensure that you have not selected the **Checkpoint restart** field. If selected, backups revert to traditional storage node backups.
- Ensure that the client interface configuration, whether FC or IP, matches the DD Boost device interface configuration. If the interfaces do not match, then the storage node performs the backup and restore operations.
- Ensure that you have configured the Data Domain system to use the DD Boost devices. Configuring the Data Domain system for DD Boost on page 96 provides details.
- Ensure that you have configured the NetWorker Device resource for the Data Domain system with a **Remote User** field that specifies a DD Boost username. Configuring a DD Boost device manually on page 104 provides details.

Note

Client Direct access from a Linux host to a Data Domain system requires a glibc 2.3.4 or later library on the Linux system.

---

b. (Optional) In the Protection group list field, select the group in which to add the Client resource. If you have not created the protection group, you can create one later and add the client to the group.

4. On the **Apps & Modules** tab, perform the following tasks:

a. In the **Deduplication** area, select **Data Domain backup**. This ensures that NetWorker backs up the client data only to DD Boost devices, even if the selected

pool contains DD Boost and other types of devices. It is best not to include different device types in a single pool.

Note

Current versions of NetWorker application modules support backup to DD Boost devices. Some of the earlier module versions do not support the client fields for **Data Domain backup** and **Pool**. In this case, do not set these fields. Backup fails for older NetWorker application modules on page 132 provides details.

---

b. In the **Data Domain Interface** field, select the type of connectivity the client uses for DD Boost devices:

- Select **IP** for TCP/IP connectivity only.
  Do not select **IP** if the **Enable fibre channel** attribute is enabled on the **Configuration** tab of the DD Boost Device resource. This conflict in settings could cause backups to fail and restores to operate only through the storage node.
- Select **Fibre Channel** for FC connectivity only.
  You cannot select the FC-only setting when you create a Client resource for the NetWorker server resource because the NetWorker server requires IP connectivity to send control information to the hosts within the datazone.
- Select **Fibre Channel** for FC connectivity only.
  You cannot select the FC-only setting when you create a Client resource for the NetWorker server resource because the NetWorker server requires IP connectivity to send control information to the hosts within the datazone.
- To enable both FC and IP connectivity to the devices, select **Any**.
  If the NetWorker server contains multiple definitions of this Client resource, any changes to this field propagate to the other instances of the client.

c. To redirect an NDMP client from a tape backup to a DD Boost backup, change the **Backup** fields as follows:

a. Select the **NDMP** option.

b. In the **Backup Command** field, type the following command:

```
nsrndmp -T backup_type -M
```

where *backup_type* is dump, tar, or vbb.

The `-M` option specifies a backup with the Data Service Agent (DSA) option.

The *EMC NetWorker Administration Guide* provides details on the `nsrndmp` command.

5. On the **Globals (2 of 2)** tab, in the **Configuration** area, configure the following settings:

a. In the **Storage Nodes** field, type the hostnames of the remote storage nodes that receive the client backup data.

b. In the **Recover Storage Nodes** field, type the hostnames of the storage nodes that you use to restore the client data.

c. (Optional), In the **Backup target disks** field, specify an ordered list of AFTD and Data Domain disk devices that will receive data for this client. When you specify a value in this attribute, NetWorker ignores the values that you specify in the **Storage nodes** attribute. This attribute does not apply to the client resource of the NetWorker server, and applies to each instance of the client resource. You can specify devices that are local or remote to the NetWorker server.

6. When you have completed the client configuration, click **OK**.

   The NetWorker server window shows a check mark in the Scheduled backup column of clients that are enabled for scheduled backup.

Deactivating and removing DD Boost devices

To prevent NetWorker from using the DD Boost device, use one of the following procedures to deactivate a DD Boost device.

Converting a device to read-only

When you convert a device to read-only mode NetWorker cannot use of the device for backup operations. You can continue to use the device for read operations, for example, restore operations and as the read device for clone operations.

Procedure

1. Use NMC to connect to the NetWorker server, click the **Devices** view and select the **Data Domain Systems** folder in the navigation tree.
2. In the **Devices** table, right-click the device you want to convert to read-only and select **Unmount**.
3. Right-click this unmounted device, and select **Modify Device Properties**.
4. On the **General** tab, select **Read Only** and click **OK**.
5. Right-click the device, and select **Mount**.

Disabling a device

When you disable a device, NetWorker does not use the device for backup, recovery or clone operations. You can reenable the device to restore old data retained on the device.

Procedure

1. Use NMC to connect to the NetWorker server, click the **Devices** view and select the **Data Domain Systems** folder in the left navigation pane.
2. In the **Data Domain Systems** table, right-click the device you want to disable and select **Unmount**.
3. Right-click the device, and select **Enable/Disable**.
4. Confirm that the **Enabled** column of the table contains **No**, which indicates that you have disabled the device.

Deleting an AFTD or DD Boost device

The procedure to delete an AFTD or DD Boost device includes an option that will erase the data on the volume, denoted by the access path, that stores the device's data. You can erase the volume only if no other device in the system shares the volume.

Procedure

1. Use NMC to connect to the NetWorker server and select **Devices**. Click **Device** in the left navigation pane.
2. In the **Devices** table, right-click the device that you want to remove, and then select **Delete**.

   A confirmation window appears.

3. Specify the required setting in the confirmation window:
   - To delete the device from the NetWorker configuration without erasing the data on the device, click **Yes**.
   - To delete the device and erase the data on the device and volume access path, select **Permanently erase all data and remove media and index information for any selected AFTDs or Data Domain devices**, and then click **Yes**.

   Note

   If another device shares the volume that you want to erase, then an error message displays the name of the other device. Before you can erase the volume, you must delete all other devices that share the volume until the last one remaining.

4. If you did not unmount the device or did not remove the device from all the Pool resource configurations, then a confirmation window appears, which provides these details. To confirm the device unmount, the removal of the device from the pool, and the deletion of the device, click **Yes**.

CHAPTER 5

Restoring Data

This chapter includes the following topics:

- Restoring DD Boost deduplicated data 118
- Restoring by Client Direct over IP from an FC-enabled device 119
- Disaster recovery 119

Restoring DD Boost deduplicated data

You restore deduplicated data from DD Boost devices in the same way as you would restore non-deduplicated data. Each backup consists of two components that reside in different places:

- Deduplicated client backup data resides on the DD Boost devices on the Data Domain system.
- Backup metadata, which specifies how long you want to retain the data and allows you to browse the backups for recovery, resides in the media database and the client file indexes on the NetWorker server.

Restore requirements for deduplicated data

The requirements for the restore of deduplicated data from DD Boost devices are as follows:

- All the deduplicated data must be available on the Data Domain system. The retention periods for the backups must not have expired.
- Both the Data Domain system and the NetWorker storage node must be online during the restore of deduplicated data.

Supported NetWorker restore procedures

The data restore procedures are the same as for non-deduplicated NetWorker backups:

- You can use NetWorker to browse the client file index to select files or save sets to recover.
- You can perform directed restores for supported NetWorker clients and NetWorker storage nodes.
- You can try to restore expired backup data by using the NetWorker scanner program to reconstruct a media index from the surviving metadata.

The *EMC NetWorker Administration Guide* provides procedures for data recovery.

Note

For a Linux client which has 2 GB of memory (RAM), EMC recommends that you recover only up to a maximum of 3500 files at a time. If you attempt to recover more than this limit, an error message similar to the following appears:

```
readv from DD failed for read size 262144: Reading from a file failed
recover: Reading from a file failed [5001] ([31587] [140129876305664]
ddp_read() failed Offset 0, BytesToRead 262144, BytesRead 0 Err: 5001-Unable
to allocate file ddcl buffers rec_create: out of memory.
```

Restoring by Client Direct over IP from an FC-enabled device

You can use Client Direct over an IP network to restore data from a volume that you have mounted on an FC-enabled DD Boost device. Share the volume with an IP-enabled device as follows.

Procedure

1. Create an IP-enabled DD Boost device on which to mount the volume, but associate this device with a different storage node than the one that manages the FC-enabled DD Boost device.

   The storage node that you use for the IP restore must not have an FC-enabled DD Boost device available to the volume.

   Configuring DD Boost devices with the wizard on page 100 provides details.

2. Configure the devices to share the volume.

   Configuring volume sharing on multiple devices on page 107 provides details.

3. Modify the Client resource for the client that will receive the restored data.
   - Configure this client for Client Direct and IP connectivity.
   - On the **Globals (2 of 2)** tab, in the **Recovery storage nodes** field, specify the storage node that you associated to the IP-enabled DD Boost device.

     Note

     This option is now only available in the Diagnostic Mode view.

     Configuring a backup client with NMC property windows on page 112 provides details.

4. Mount the volume on the new IP-enabled DD Boost device and perform the restore by using the new IP-restore Client resource.

Disaster recovery

In this guide, a disaster is any loss of data in which the computing environment required to restore that data is not available. Disaster recovery is necessary when ordinary data recovery procedures are not sufficient to recover the computing environment and its data to normal day-to-day operations.

Causes of disaster

A disaster can result from any of the following situations:

- Debilitating hardware or software failures
- Computer viruses that corrupt the computing system
- Infrastructure interruptions, inconsistencies, or loss of services, such as problems with communications or the network connections that result in damage to the computing environment

Potential losses

Disaster recovery of the primary site must cover the potential loss of any or all the following systems at the primary site:

- The Data Domain server that stores the deduplicated client backups
- The NetWorker storage node that stores the deduplication metadata for the backups
- The NetWorker server that stores the metadata for the backups in the media database and client file indexes

Disaster recovery requirements

A complete disaster recovery environment provides a secondary site with systems that copy all the information involved in each completed backup at the primary site.

You can configure the two sites to provide disaster recovery for each other, with each serving as both a primary and secondary site with different datazones for different clients.

Disaster recovery requires the maintenance of the following systems:

- Data Domain system with deduplicated client data cloned from the primary Data Domain system
- Disaster recovery NetWorker storage node with deduplication metadata cloned from the primary NetWorker storage node
- Disaster recovery NetWorker server with metadata cloned from the primary NetWorker server

Disaster recovery environment on page 31 shows an example of a simple disaster recovery environment.

Disaster recovery scenarios

The procedures you use to recover from a disaster will vary depending on the circumstances, which could include the following factors:

- The deployment of the disaster recovery environment
- Which systems the disaster has affected
- The time required to successfully recover from the disaster

The *EMC NetWorker Disaster Recovery Guide* provides details.

CHAPTER 6

Monitoring, Reporting, and Troubleshooting

This chapter includes the following topics:

- Monitoring Data Domain events, statistics, and logs 122
- Generating reports 125
- Replacing a failed or old storage node 129
- Troubleshooting 131

Monitoring Data Domain events, statistics, and logs

NMC provides several ways to view for the backup statistics, the logs, and the alerts for connected Data Domain systems.

Viewing the statistics, logs, and alerts

When you use NMC to connect to the NetWorker server, the NetWorker Administration window provides a comprehensive view of the backup status, log, and alerts for connected Data Domain systems.

Note

In some logs and notifications, the Administration window lists Client Direct operations variously as direct file assess (DFA), direct file save, or DIRECT_FILE operations.

Procedure

1. Ensure that you have configured SNMP for the Data Domain system.

   Configuring SNMP for an NMC managed Data Domain system on page 124 provides details.

2. In the Administration window, click the **Devices** view.
3. In the folder tree, select **Data Domain Systems**, and perform one of the following operations:
   - Right-click a Data Domain system and select **Properties** to view system information, including the identity (name, hosts, model, OS version, serial number), configuration, and SNMP community string, access credentials, capacity status information, save stream status information, and system details.
   - To view backup information select a Data Domain system, as shown in the following figure:

     Figure 19 NetWorker Administration window displaying DD Boost devices



- The **Devices** area shows device and usage information, including the following:
  - Pre-Compression—Indicates the amount of space that the backup would have used if the data had not been deduplicated and compressed. NetWorker tracks this value as the size of backups.
  - Compression (Reduction)—Represents the data compression with the pre-compression and post-compression used values. Data compression is calculated with:
    [(1 - Post-comp Used) ÷ Pre-Compression] x 100%
  - /backup: post-comp—Indicates three values: the total capacity of the Data Domain system, the amount of disk space already in use, and the amount of space available.
  - /ddvar—Indicates the amount of log file space that is in use on the Data Domain file system.
- The **Status** area lists connectivity usage.
- The **Log** table shows a chronological list of events that occur during NetWorker server operations.
- The **Alerts** table lists messages for operational issues that can require administrative attention. Data Domain-specific alerts are available only if SNMP traps are configured.

Note

To delete individual messages from the **Alerts** table, open the NMC **Events** view, select the messages, right-click, and select **Dismiss**.

Viewing backup statistics in NMC

You can view the storage statistics for backups on a connected Data Domain system in NMC.

In the **NMC Enterprise** view, select a Data Domain host. A table shows the storage statistics for the selected system.

Viewing backup alerts (SNMP traps) in NMC

Alerts are messages for operational issues that can require administrative attention.

You can view backup alerts on a connected Data Domain system in NMC.

Procedure

1. Configure SNMP for the Data Domain system.

   Configuring SNMP for Data Domain on page 124 provides details.

2. In NMC, select the **Events** view.

   A table lists the backup alerts (SNMP traps) in chronological order.

   Note

   The same alert messages also appear in the NetWorker **Alerts** table.

Deleting individual messages from NMC Events and NetWorker Alerts

You can delete individual messages from the NetWorker Alerts and NMC Events tables by removing the messages from the NMC Events table. The two views show the same messages.

Procedure

1. In NMC, select the **Events** view.
2. Select the messages you want to remove from the Events table.
3. Right-click and select **Dismiss**.

   NSM deletes the selected messages.

Configuring SNMP for an NMC managed Data Domain system

You can configure NMC to monitor Data Domain alerts (SNMP traps) when you add a Data Domain system to the NMC Enterprise, or you can update an existing Data Domain system that is managed by the NMC server. Provided that you have viewing privileges, the NMC Enterprise view lists the Data Domain systems as network hosts.

Adding a host Data Domain system to NMC Enterprise view on page 104 describes how to add a Data Domain system to the NMC Enterprise view.

Procedure

1. Enable SNMP on the Data Domain system and configure the system to send traps to the NMC server. Configuring the Data Domain system for DD Boost on page 96 provides details.
2. In the NMC **Enterprise** view left panel, right-click the Data Domain system that you want to monitor and select **Properties**.
3. In the **Properties** window, on the **Manage Data Domain** tab, ensure that **Capture Events** is enabled.

   If you do not select the **Capture Events** checkbox, NMC monitors the status of the DD Boost devices but will not monitor Data Domain SNMP traps that are required to monitor events.

4. On the **Configure SNMP monitoring** tab, type a value for **SNMP Community String**. The typical setting is public, which allows all users to monitor events. The following figure shows an example.
5. Type a value for the port that the Data Domain system uses for SNMP traps in the **SNMP Process Port** attribute. Firewall requirements on page 27 provides details.
6. Select the SNMP traps that you want to monitor in the **SNMP Traps** section. Some traps are pre-selected. The following figure shows an example for Data Domain 5 alerts. Other versions may differ.

Figure 20 Data Domain alerts to monitor



7. Click **OK**.

Generating reports

Use the NMC Reports view to create statistical reports of NetWorker with Data Domain backup, recovery, and cloning activities.

Configuring a report

You can configure and display a Data Domain report for backup or clone in NMC.

Procedure

1. In the **NetWorker Management Console** window, click **Reports**.
2. Expand the **Reports** folder, expand the **Legacy Reports** folder, and then the **Data Domain Statistics** folder. Select the report that you want to view.

   Note

   There are three types of reports: summary, statement, and details.

   The Configure tab for the selected report type appears in the right panel.

3. In the **Configure** tab, customize the items that you want to include in the report. Select the item parameters and click the **Remove** (‹), **Add** (›), **Remove All** («), or **Add All** (») buttons as required.

   If you do not specify Save Time values, the report displays all the available data.

   The following table lists details of report configuration parameters. The specific parameters available depend on the type of report selected.

   The following figure shows an example report configuration.

4. To display the report, select the **View Report** tab.

Table 12 Data Domain report configuration parameters

Parameter	Description	Options
Server Name	Selects managed hosts within the enterprise.	Selected server names
Group Name	Selects one or more groups.	Selected group names
Client Name	Selects one or more clients. **Note** Monthly report does not include the Client Name parameter.	Selected client names
Save Set Name	Selects one or more save sets. Values are case-sensitive and you cannot use wild cards. **Note** Monthly report does not include the Save Set Name parameter.	Selected save set names
Save Time	Limits the report to a specified time range. The date/time format available depends on the language locale of the operating system.	Save time (within a range)

Figure 21 Report configuration



Types of backup reports

Backup reports are available in various formats. Most are basic reports. The Backup Summary and Monthly Client Statement are drill-down reports:

- Basic reports on page 127 describes details of basic reports.
- Drill-down reports on page 128 describes details of drill-down reports.
- Advanced reporting on page 129 describes advanced reporting functionality with the optional EMC Data Protection Advisor (DPA).
- For clone operations, there is no specific report. You can query and list the copies of save sets in the NetWorker Administration, Media view, under Save Sets.

Basic reports

A basic report collects statistics for a specific datazone component, a time span, or a field. You can modify the scope of a report by adjusting the parameters on the **Configure** tab in NMC.

The following table describes the basic reports available for Data Domain statistics.

Table 13 Data Domain basic reports

Report name	Purpose
Client Summary	For all or specified clients, displays the following statistics: • Amount of data—The amount of the data that NetWorker would have moved by using a conventional backup (protected data). • Target size—Size of the data after deduplication has taken place on the Data Domain system (stored data). • Deduplication ratio—Percentage of savings by using Data Domain deduplication.

Table 13 Data Domain basic reports (continued)

Report name	Purpose
	• Number of save sets—The number of save sets in the backup. • Number of files—The number of files in the backup.
Save Set Summary	For all or specified save sets, displays deduplication statistics on the following items: • Amount of data—The amount of the data that NetWorker would have moved by using a conventional backup. • Target size—Size of the data after deduplication has taken place on the Data Domain system. • Deduplication ratio—Percentage of disk space savings by using deduplication. • Number of save sets—The number of save sets in the backup. • Number of files—The number of files in the save set.
Save Set Details	Displays details about each save set, including backup duration and the following statistics: • Save Set ID • Save time • Backup level • Save Set size—Protected data size • Target size—Size of the data after deduplication has taken place on the Data Domain system (stored data size). • Deduplication ratio—Percentage of savings by using deduplication. • Number of files—The number of files in the save set.
Monthly Summary	Displays statistics on a month-to-month basis.
Daily Summary	Displays statistics on a day-to-day basis.

Drill-down reports

A drill-down report consists of multiple basic reports, connected as layers and all configured with the same parameters that the top layer uses.

You can run reports for groups, clients, or save sets. You can modify the scope of a report by adjusting the parameters on the **Configure** tab in NMC.

The following table lists the drill-down reports available for Data Domain statistics.

Table 14 Data Domain statistics drill-down report

Report name	Purpose	Sequence
Backup Summary	Reports backup statistics over a period, including a client summary.	1. Client Summary 2. Save Set Summary 3. Save Set Details
Monthly Client Statement	Reports backup statistics of individual clients on a month-to-month and day-to-day bases, down to individual save sets details.	1. Client Summary 2. Monthly Summary 3. Daily Summary 4. Save Set Details

Data Domain statistic reports

Use the Data Domain `gstclreport` command with a specified format to generate a specific Data Domain statistics report.

The Data Domain product documentation provides details.

Advanced reporting

NMC provides reports for only the recent backup history in a specific datazone. The optional EMC DPA software can provide extended reports of backups, trends, and analysis for one or multiple datazones, including reports of Data Domain systems. DPA is best for larger environments where you require additional analysis with forecasts and trends.

Replacing a failed or old storage node

Before you begin

Before you begin, ensure the following requirements are met:

- The replacement storage node has access to the original Data Domain system.
- The NetWorker server software is the same version as the original.
- The NetWorker server has all the same indexes and the same media database entries as before the disruption.

If a storage node fails or if you replace a storage node, you can recover the data stored on the associated DD Boost devices on the replacement storage node or on a different storage node. The success of the recovery depends on the state of the devices at the time of the loss:

- If the storage volumes were unmounted when the disruption occurred, the structure and integrity of the data remains intact and you can expect a complete recovery.
- If the volumes were mounted but not reading or writing data during the disruption, then complete recovery is still likely.
- If the devices were reading or writing at the time of the disruption, then data loss or data corruption is more likely to have occurred, and you cannot assure complete recovery.

If the volume structure of the devices is intact, then the NetWorker server can continue to perform its operations with the existing devices, with minimal impact.

If the replacement storage node has a different name or if you use the NetWorker server as the storage node, then you must re-create the devices in NetWorker as follows:

Procedure

1. Use NMC to connect to the NetWorker server. In the **NetWorker Administration** window, select the **Devices** view, and then select **Devices** in the left navigation pane.
2. For each affected original remote storage node-based DD Boost device, right-click the device, select **Properties**. Record the following information:
   - On the **General** tab:
     - Name
     - Device Access Information
   - On the **Operations** tab:
     - Volume Name
     - Volume Pool
3. Remove the original DD Boost devices from the NetWorker application. The device folders will continue to exist on the Data Domain system:
   a. In the **Devices** view, from the **Devices** tree, right-click and unmount each affected device. Mounted devices have a **Volume Name**.
   b. In the **Media** view, from the **Media Pool** tree, right-click each affected media pool (**Volume Pool**), select **Properties**, and on the **Selection Criteria** tab, remove each affected device from the **Target Devices** list.
   c. In the **Devices** view, **Devices** tree, right-click and delete each affected device.
4. Re-create the devices on the NetWorker application associated with a replacement storage node:
   a. In the **Devices** view, right-click the **Data Domain systems** tree and then run the **New Device Wizard**.
   b. Specify the Data Domain system and DD Boost (OST) credentials to gain access the system.
   c. On the **Select Folders to use as Devices** page, select the DD Boost devices (device folders) associated with the failed storage node.

      When you leave this page, a message notifies you that NetWorker had previously associated the devices with a different storage node. Confirm the selection.

   d. On the **Configure Pool Information** page, specify the media pool for the devices and cancel the **Label and Mount** selection. You must manually mount the devices on the new storage node later in this procedure.

      NOTICE

      If you enable **Label** and **Mount** at this point, NetWorker relabels the volume and you lose all the data. You cannot undo this action.

   e. On the **Select the Storage Nodes** page, select a storage node to handle the new devices by doing one of the following.
      - Select an existing storage node.

- Create a replacement storage node.
- Use the NetWorker server's storage node.
  The storage node must be running on the correct network and its hostname must be resolvable by DNS.

f. Complete the wizard.

5. Manually mount each new device:
   a. In the NMC window for the NetWorker server, click **Devices**.
   b. In the navigation tree, select the Data Domain system.
   c. In the right panel, right-click each device you want to mount, and select **Mount**.

      The device mounts on the storage node and will use the label associated with the pool you have specified.
6. Review the NMC log for any error messages.

Results

If this procedure does not report any errors, then the device and the volume are available for use. Backup and recovery operations may require further configuration depending on the original settings and the purpose of the device recovery.

Troubleshooting

The following sections will help you identify and resolve common issues of configuration and operation.

Name resolution issues

If connectivity issues are present, ensure that the network names are valid and consistent for the NetWorker server, the storage nodes, and the Data Domain systems. Use the same names that are consistent with the NetWorker software configuration.

Validate connections from the Data Domain system and the NetWorker server, and from the NetWorker server to the Data Domain system by using IP addresses and the network names. If you use aliases or short names in the configuration, then verify the aliases and short names. To validate connections, use one of the following methods:

- On the NetWorker server and storage nodes, use the `nslookup` command to verify that network names resolve to the correct IP address.
- On the Data Domain system, use the `net hosts` command.
  Host naming guidelines on page 29 provides suggestions for names.

Correct improper names by amending DNS entries or by populating the local hosts files.

Network connection issues

You can test the network connections for a Data Domain system by using the net lookup command through an SSH Telnet session, which requires administrator or system administrator permissions.

The Data Domain system can also show the current network configuration by using the net show and the other network related commands, available through the Data Domain interface. Log in and go to the specific Data Domain system. Then select the **Hardware** › **Network** tabs to access the commands.

EMC recommends that you diagram and verify all relevant network connections. A typical Data Domain network configuration provides a minimum of two network connections, one dedicated to administration and the other to backup data only. You can make effective use of 10 GbE connectivity or the use of multiple backup connections that you can aggregate or team together by using the `ifgroup` command on the Data Domain system.

Network requirements on page 25 provides suggestions for network connections.

Device access errors

The following error messages can occur when NMC cannot connect to a DD Boost device.

Volume unavailable error

This message appears when the Data Domain file system becomes inaccessible or disabled, and then you reenable the file system. For example, for service or testing, you could leave the devices in an unmounted state. Backup operations for the devices will not start and an error message similar to the following appears in the **Log** pane in the **Administration** window, and the `daemon.raw` file:

```
Waiting for 1 writeable volume(s) to backup pool
```

To resolve this issue, perform the following steps to mount and enable the device:

1. In the **NetWorker Administration** window, click the **Devices** view.
2. In the **Devices** table, right-click and select **Mount** for any unmounted DD Boost device.
3. To enable the device, in the **Enabled** column, right-click the device and select **Enable/Disable**.

NFS service errors

You must enable Data Domain NFS service for the NetWorker software to access DD Boost devices. Without NFS, an error message similar to the following appears, typically when NetWorker tries to label a device:

```
Failed to contact the Data Domain system. Host could be unreachable,
or username/password could be incorrect. Do you wish to configure
manually?
The user has insufficient privilege
```

Configuring the Data Domain system for DD Boost on page 96 describes how to enable NFS access.

Backup fails for older NetWorker application modules

Some older NetWorker application modules do not support the NetWorker Client resource fields for Data Domain Backup and Target Pool or Pool and you must not use these fields for DD Boost backups. The Data Domain Backup field ensures that backups use only DD Boost devices, even if the configured pool contains other device types, although pools with mixed devices is not a good practice.

The *EMC NetWorker Administration Guide* provides details on how to configure a pool to target DD Boost devices only.

The release notes for the specific NetWorker application modules provide details on supported Data Domain configurations.

Multiple recovery fails on AIX clients with less than 2 GB RAM

For NetWorker clients on AIX systems with less than 2 GB of RAM, a recovery that uses four or more parallel recovery save stream IDs might fail with an error message similar to the following:

```
93124:recover: readv from DD failed for read size 262144: Reading
from a file failed ([5001] memory no longer available)
```

To avoid this error, export the following environment variable on the client shell.

```
LDR_CNTRL=MAXDATA=0x70000000
```

Backing up streams from NetWorker to Apollo DD is rejected

Issue

When you perform a backup of 1024 streams from a NetWorker server to Apollo DD, the backup succeeds. However, if you run three NetWorker servers simultaneously with 3000+ save sets, the system is unable to reach the 1885 write streams limit on Apollo DD. Approximately 1600 write streams start on Apollo DD from the three NetWorker servers, and the system displays connection rejection messages in the `ddfs.info` log file on Apollo DD.

Workaround

To resolve this issue, ensure that you are in SE Mode, then change the attribute NFS_TOTAL_CONNS_PERCENT from the default 50 to 100.

```
reg set system.NFS_TOTAL_CONNS_PERCENT = 100
```

APPENDIX A

DD Boost Conversion and Upgrade

This appendix includes the following topics:

- Upgrading from legacy DD Boost configurations 136
- Converting DD Boost devices from IP to FC connectivity 136
- Redirecting backups from other devices to DD Boost 136
- Migrating data into Data Domain systems 137
- Migration scenarios 140

Upgrading from legacy DD Boost configurations

NetWorker 8.0 and later provides both write and read functionality on all upgraded and new DD Boost devices. Earlier releases had separate read-only DD Boost mirror devices for restore operations. The NetWorker 8.0 and later upgrade removes legacy read-only DD Boost mirror devices.

Note

After an upgrade to NetWorker 8.0 or later from an earlier release that used read-only devices, do not revert to the earlier release. If you revert, DD Boost devices created by NetWorker 8.0 or later will be unavailable and the legacy devices will require manual reconstruction.

Converting DD Boost devices from IP to FC connectivity

After you have met the FC support requirements, you can convert existing DD Boost devices that use Ethernet IP connections to use FC connections deployed as a SAN. No data is lost by the conversion and full DD Boost features are retained, including Client Direct operations for backup and restore.

Plan your device conversion with the following high-level road map that outlines the sequence of basic tasks that you must perform.

Procedure

1. Ensure that you meet all FC support requirements.

   FC support on page 26 provides details.

2. Configure the **Fibre Channel Options** of the DD Boost devices.

   Configuring DD Boost devices with NMC property windows on page 103 provides details.

3. Configure the **Data Domain Interface** field of the NetWorker clients for FC.

   Configuring a backup client with NMC property windows on page 112 provides details.

Redirecting backups from other devices to DD Boost

You can redirect the backups of existing NetWorker clients that do not use DD Boost devices to use new DD Boost devices.

To redirect the backups, configure the Pool resource to use DD Boost devices, as described in the following procedure.

After you redirect backups to DD Boost devices, ensure that you perform a full backup of the data. This practice avoids a dependency on the last full backup in the legacy storage environment and the potential need to restore from two different environments. To perform a full backup after you configure a DD Boost Pool resource, perform one of the following options:

- Configure the redirection of the backups to a DD Boost device at a time when the next scheduled backup for the client data will be a full backup.
- Configure the redirection, and then change the backup schedule to accommodate an initial full backup.

Complete the following steps to redirect data from existing scheduled client backups to use storage on DD Boost devices.

Procedure

1. Ensure that the required network connection, hostname resolutions, and licenses are available and ready to use. The following sections provide details:
   - Licensing in Data Domain systems on page 20
   - Host naming guidelines on page 29
   - Network requirements on page 25
2. Configure the Data Domain system for use with NetWorker. Configuring the Data Domain system for DD Boost on page 96 provides details.
3. If you plan to migrate existing save sets to the new DD Boost devices, migrate the save sets before scheduled redirected backups begin. Migration will “seed” the Data Domain system and help to reduce the bandwidth requirements for future backups. Considerations for migrating legacy save sets on page 137 provides details.
4. Use the NMC Device Configuration wizard to perform the following tasks:
   a. Select or create DD Boost devices on the Data Domain system.
   b. Select or create a Pool resource that is configured to send the save sets to DD Boost devices.

      Note

      The wizard creates and configures a pool for the Data Domain system that uses only DD Boost devices.

   c. Select or create a NetWorker storage node on which to label and mount the new devices.
   d. Complete the wizard pages.

   Configuring DD Boost devices with the wizard on page 100 provides details.
5. Test the backup environment to ensure that the new configuration operates correctly and that existing backups, that will not use DD Boost devices, continue to run as expected. For backups to new devices, test a restore from those devices.
6. Start the redirection with a full backup to the new devices. This practice avoids a dependency on the last full backup stored with the legacy storage environment and the potential need to restore from two different environments.
7. Monitor backup performance, and adjust the backup schedule to optimize the configuration for maximum throughput or additional clients. Monitoring Data Domain events, statistics, and logs on page 122 provides details.

Migrating data into Data Domain systems

When you successfully redirect client backups to the DD Boost devices, the existing save sets that are stored on the legacy devices or file systems become redundant. You can retain the legacy stored data until the data expires, or you can migrate the legacy data to the new devices.

The decision to retain or migrate the legacy data depends on the requirements that will differ between sites, clients, and backup types. For example, you may want to retain most of the legacy data and migrate only the backups of sensitive and high-priority clients or certain backup types.

Review the following information to help you decide to either retain or migrate the existing save sets as follows:

- Retain the existing save sets on the legacy storage system until they expire:
  - Provides the easiest way to make the transition; no migration is necessary.
  - Requires you to maintain the legacy storage for the life of the legacy data.
  - Requires that you maintain legacy storage environment, to perform recoveries of data on the legacy devices.
  - Provides features for storage, recovery, and clone operations that will differ between the legacy data and the new data.
- Migrate the existing save sets to the new DD Boost devices, which provides the following advantages:
  - Frees storage on the legacy storage system for removal or use by other clients.
  - Allows you to "seed" the new devices with the legacy client data. Seeding ensures that subsequent client backups will be deduplicated against the legacy data. This practice reduces the bandwidth and time required for the first backup window with the new devices.
  - Offers more flexible storage features for storage, recovery, and cloning, for example, multiple concurrent operations.
  - Maintains the NetWorker browse and retention policies and ensures that NetWorker manages all save sets.

Migration versus native Data Domain replication

EMC recommend that you do not use the native Data Domain replication feature to migrate data from one Data Domain system to another. NetWorker cannot track, manage, or recover legacy save sets that Data Domain replicates.

You can use the Data Domain replication feature to seed a new system to migrate the data. For example, you can perform native Data Domain replication over a local connection to quickly seed a new target Data Domain system, which you can then physically send to a distant location. Although NetWorker cannot immediately manage or restore the seeded data, this practice has advantages. The seeded data reduces the otherwise heavy bandwidth that is required for a data migration by using a NetWorker clone operation, or if you will not perform a migration, for the initial full backups to the target system. This practice can be especially effective if the remote location has limited network bandwidth.

Migration methods

Data migration is a one-time NetWorker clone operation which you can customize to different device types and time periods. You can include all the data in the migration or you can select a limited amount of data from a specific timeframe or a specific backup type, for example, weekly full backups.

The details of the migration procedure depend on the method that you use and the granularity of the data that you want to migrate:

- To perform a NetWorker scheduled clone operation, refer to Migrating legacy save sets to DD Boost devices on page 139.
- To run a NetWorker `nsrclone` script from a command line, refer to the *EMC NetWorker Administration Guide* for details.

- To perform a NetWorker staging (data movement) operation to move data from an AFTD, refer to the *EMC NetWorker Administration Guide* for details.

Migrating legacy save sets to DD Boost devices

After you choose a migration scenario, you can migrate the existing save sets to DD Boost devices. Part of this procedure requires that you create a special clone pool and configure a clone task.

Before you begin

Perform migrations before the scheduled NetWorker client backups begin using the new devices. Migration will seed the Data Domain system and help to reduce the bandwidth requirements for future backups.

Procedure

1. Decide which migration scenario you need. Migration scenarios on page 140 provides details.
2. Plan the migration schedule to ensure that sufficient DD Boost devices and bandwidth are available and that there will be minimal impact to the usual backup window.

   Note

   When you migrate existing deduplicated VTL or CIFS/NFS AFTD save sets, the deduplication software reverts the save sets to their native non-deduplicated format. The storage node then reads and stores the save sets in deduplicated format on the new DD Boost devices. This reversion process occurs for both Data Domain and non-Data Domain storage.

3. Create a clone pool for the DD Boost devices to be used for the migration:
   - In the **Data Source** field, select groups for the migration.

     Typically, you will migrate the same groups that you selected for the redirection of backups. Redirecting backups from other devices to DD Boost on page 136 provides details.
   - In the **Target Devices** field, select the DD Boost devices to store the migrated data.

     Creating pools to target DD Boost devices on page 109 provides details.
4. Configure a clone task with the **Write Clone Data to Pool** field selected for the clone pool.

   Road map for configuring a new cloning data protection policy on page 71 provides details on the scheduled clone option.
5. Run the clone action, either according to its schedule or by manual start.

   To amnually start the clone action, right-click on the workflow that contains the clone action, and select **Start**.
6. After the clone operation completes, verify that the cloned data appears on the target devices.

   DD Boost Conversion and Upgrade on page 135 provides details on the verification of NetWorker operations.
7. After you have verified the cloned save sets, remove the original save sets as required.

8. If you remove the original save sets, remove unused devices and pools as required. You cannot delete a pool until you delete or relabel in other pools all the volumes that belong to that pool.
9. Monitor the Data Domain system to ensure that adequate storage capacity is available. Monitor a complete backup cycle of all clients, including save set expirations.

   DD Boost Conversion and Upgrade on page 135 provides details.

Migration scenarios

This section offers some typical migration scenarios to help you migrate existing backup data from legacy devices or file systems to DD Boost devices. The best scenario for your situation will depend on the storage environment configuration and the available capacities and bandwidth.

Migration to DD Boost from conventional tape or AFTD

In these first two migration scenarios, you have added a Data Domain system to the existing NetWorker storage environment. You want to migrate and deduplicate the current legacy data, stored on tape or conventional disk, to DD Boost devices on the new system. The reason for this migration could be that you want to remove the old tape or disk system or free up space on the old system for other clients.

The number of client migrations that you will perform depends on whether you want to seed the devices for future backups or migrate all the legacy save sets. EMC recommends that you seed some of the data, because the new Data Domain system contains no data. If you migrate the data for one client to seed the DD Boost devices and some of the same data exists on other clients, then migrating the data for the additional clients has diminishing seed value.

There are two scenarios for this type of migration. In the first case, you create the DD Boost devices on a new storage node. In the second case, you create the devices on the existing storage node.

Migration to new devices on a different storage node

The figure in this section illustrates a scenario where the storage node named Cassie stored backups of the client named Olive on tape or conventional disk. You want to migrate these backups to a different storage node named Eric for storage on the Data Domain system.

In this scenario, you use the IP network to transfer the data from the original storage node Cassie to the new storage node Eric. The time that is required for the transfer depends on the capacity and bandwidth available on the IP network, regardless of the fact that the tape library is on a SAN. If restore operations must use the IP network during the transfer, then it takes additional bandwidth to ensure that the data transfer does not impact these operations.

Figure 22 Migration from conventional storage to DD Boost devices on a different storage node



Migration to new devices on the same storage node

You can eliminate data migration over the IP network between storage nodes by migrating data between devices on the same storage node. The figure in this section illustrates a scenario where you migrate data to DD Boost devices that were created on the original storage node named Cassie. During the migration, the storage node reads the data that is stored on tape or conventional disk and sends the deduplicated data to the Data Domain system for storage.

Although this scenario appears to be an ideal solution that avoids IP network restrictions, consider the following factors:

- The existing storage node is likely to be older and already at or near capacity. This situation limits the number of devices that you can add and the amount of data that you can transfer during backup operations.
- The existing storage node is not likely to have extra network connections available. If you need the legacy connections for backup and restore operations, this situation leaves limited bandwidth available for the additional DD Boost format.
- The network connection is less likely to have the recommended 10 GB capacity to maximize throughout from the storage node to the DD Boost devices.
- Although you use the same storage node for the same backup clients, change the device allocations and the pools. These changes can add confusion and result in configuration errors.

There are also advantages to this scenario. For smaller sites, to avoid network restrictions you can migrate the data to new devices on the same storage node. This scenario could also be an option for larger sites where you want to reuse multiple storage nodes or reconfigure the storage nodes to share one or more Data Domain systems. You can configure a storage node for data migration to seed the DD Boost devices as an interim step.

Figure 23 Migration from conventional storage to DD Boost devices on the same storage node



Migration to DD Boost from deduplicated VTL or AFTD storage

In the following two migration scenarios, you are already using an existing Data Domain system for VTL or CIFS/NFS AFTD deduplication storage. You want to migrate the stored data to new DD Boost devices on this same Data Domain system. Because the data is already present on the Data Domain system, you do not need to migrate the data to seed the DD Boost devices. The global deduplication format ensures that NetWorker does not resend data that exists on the Data Domain system.

These migration scenarios offer multiple concurrent operations for storage and recovery and more flexible storage options for cloning.

Although these migration scenarios use the same Data Domain system, change the pools and the device allocations to redirect the backups to the DD Boost devices. Copy or clone the save sets to migrate the data.

When you migrate existing deduplicated VTL or CIFS/NFS AFTD save sets, the process initially reverts the save sets to their native non-deduplicated format. The storage node then reads and concurrently stores the save sets in a deduplicated format on the new DD Boost devices. Data that exists in a deduplicated format on the Data Domain system is not deduplicated again. During the migration, only the metadata for the save sets are unique.

Migration to new devices on a different storage node

The figure in this section illustrates a legacy scenario where the storage node Cassie stored backup data from the client named Olive in VTL format over a SAN connection. You want to migrate this data to the new DD Boost devices on a different storage node named Dove.

This migration uses the SAN and the IP networks in two separate stages. First, the original storage node Cassie reads the non-deduplicated data provided by the Data Domain system over the SAN connection. Then the new storage node Dove reads this data and concurrently stores only unique data, in this case only the storage metadata, across the IP network to the Data Domain system. The limiting factor is the speed of the transfer across the IP network.

Figure 24 Migration from VTL to DD Boost devices on a different storage node



Migration to new devices on the same storage node

The figure in this section illustrates a scenario where you want to migrate legacy backup data from the client named OLIVE to new DD Boost devices on the original storage node named CASSIE. The existing storage node configuration is for VTL storage on a SAN. You have added the configuration for the new DD Boost devices that use the IP network.

Because this migration is between devices on the same storage node, this scenario fully uses the speed of the existing SAN connection. The storage node CASSIE reads non-deduplicated data over the SAN and concurrently stores only unique data, in this case only the storage metadata, across the IP network to the Data Domain system.

Figure 25 Migration from VTL to DD Boost devices on the same storage node



GLOSSARY

This glossary contains definitions for 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.

advanced file type device (AFTD) Disk storage device that uses a volume manager to enable multiple concurrent backup and recovery operations and dynamically extend available disk space.

attribute Name or value property of a resource.

authorization code Unique code that in combination with an associated enabler code unlocks the software for permanent use on a specific host computer. **See** license key.

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.

bootstrap Save set that is essential for disaster recovery procedures. The bootstrap consists of three components that reside on the NetWorker server: the media database, the resource database, and a server index.

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 Direct Feature that enables clients to deduplicate backup data and send it directly to AFTD or DD Boost storage devices, bypassing the NetWorker storage node. The storage node manages the backup devices but does not handle the backup data.

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 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.

clone-controlled replication (CCR) Creation of a replica of deduplicated data copied from one DD Boost device to another, which can be scheduled by the NMC clone feature and is indexed and tracked by the NetWorker server.

D

database
1. Collection of data arranged for ease and speed of update, search, and retrieval by computer software.
2. Instance of a database management system (DBMS), which in a simple case might be a single file containing many records, each of which contains the same set of fields.

datazone Group of clients, storage devices, and storage nodes that are administered by a NetWorker server.

DD Boost Optimized library and communication framework with a special Data Domain API that allows the backup software to define and interact with storage devices on the Data Domain system.

DD Boost device Logical storage device created on a Data Domain system that is used to store deduplicated NetWorker backups. Each device appears as a folder on the Data Domain system and is listed with a storage volume name in NMC.

DD OS Data Domain operating system.

deduplication Process used to compress redundant data.

deduplication backup Type of backup in which redundant data blocks are identified and only unique blocks of data are stored. When the deduplicated data is restored, the data is returned to its original native format.

deduplication ratio Reduction in storage space required to store data as a result of deduplication technology, usually combined with data compression, for example, a 20:1 space reduction.

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.

disaster recovery Restore and recovery of data and business operations in the event of hardware failure or software corruption.

distributed segment processing (DSP) Part of the DD Boost interface, which enables data deduplication to be performed on a host before the data is sent to the Data Domain system for storage.

E

enabler code Unique code that activates the software:

- Evaluation enablers or temporary enablers expire after a fixed period of time.
- Base enablers unlock the basic features for software.
- Add-on enablers unlock additional features or products, for example, library support.

See license key.

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.

I

ifgroup A private network configured on the Data Domain system consisting of multiple network interfaces logically designated as a single group IP address. The ifgroup provides dynamic load balancing, fault tolerance within the group, and better network bandwidth usage than traditional network aggregation.

L

label Electronic header on a volume used for identification by a backup application.

license key Combination of an enabler code and authorization code for a specific product release to permanently enable its use. Also called an activation key.

M

managed application Program that can be monitored or administered, or both from the Console server.

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

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

metadata Hash information that identifies stored sub-file information for deduplication, and is required to revert deduplicated client backup data to the regular nondeduplicated format.

MTree Shortened from "managed tree," also referred to as storage units, logical partition of the namespace in a Data Domain file system that can be used to group a set of files for management purposes. MTrees are normally associated with a single NetWorker datazone.

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.

notification Message sent to the NetWorker administrator about important NetWorker events.

O

online indexes Databases located on the NetWorker server that contain all the information pertaining to the client backups (client file index) and backup volumes (media index).

optimized clone **See** clone-controlled replication (CCR)

P

pathname Set of instructions to the operating system for accessing a file:

- An absolute pathname indicates how to find a file by starting from the root directory and working down the directory tree.
- A relative pathname indicates how to find a file by starting from the current location.

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.

R

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

remote device

1. Storage device that is attached to a storage node that is separate from the NetWorker server.
2. Storage device at an offsite location that stores a copy of data from a primary storage device for disaster recovery.

replication Process of creating an exact copy of an object or data. This is different than NetWorker cloning. **See** clone

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.

resource database NetWorker database of information about each configured resource.

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.

retrieve To locate and recover archived files and directories.

S

save NetWorker command that backs up client files to backup media volumes and makes data entries in the online index.

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.

save stream Data and save set information that is written to a storage volume during a backup. A save stream originates from a single 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.

storage device See device.

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.

storage unit (SU) Logical unit of disk storage on a Data Domain system that is associated with a NetWorker datazone.

T

trap Setting in an SNMP event management system to report errors or status messages.

V

virtual tape library (VTL) Software emulation of a physical tape library storage system.

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 name Name that you assign to a backup volume when it is labeled.

INDEX

A
Add New Host Wizard 104
AFTD disk configurations 23, 139
authorization code 20

B
backup
    redirect 136
backup failure
    older NetWorker application modules 132
bootstrap 112

C
cascaded replication configuraiton 33
CCR.See clone-controlled replication 17
CIFS formats 23, 139
cleaning storage space 24
client
    configuration wizard 111
    non-wizard configuration 112
Client Direct
    configuration 112
    description 16
clone
    feature 67
    native non-deduplicated 68
    normal 68
    nsrclone 93
    operations 17
    requirements 69
    to native format 17
clone-controlled replication
    cascade configuration 33
    feature 17
    for disaster recovery 31
    process 67
    shared datazones 34
compression not supported 18
configuring
    Data Domain for NetWorker 96
    NetWorker for Data Domain 99

D
daily summary report 127
Data Domain
    backup attribute 111, 112, 136
    replication 68, 138
    Retention Lock not supported 18
    server 18
Data Domain backup storage 22
Data Domain Extended Retention 24
DD Boost
    enabling 96
    feature 16
    structure 16
DD Boost credentials 96
DD Boost devices
    creating 100
    creation by property windows 103
deduplication
    feature 16
    ratio 28
device
    creating 100
    deactivate 115
    delete 115
    disable 115
    erase data 115
    identity 104
    mounting 111
    non-wizard configuration 104
    removal 24
Device Configuration Wizard 100
direct file access (DFA) 16
disaster recovery configuation 31
distributed segment processing 16, 96
DPA advanced reports 129
dynamic drive sharing 109

E
Ethernet 18
evaluation period 20

F
FC conversion from IP 136
Fibre Channel 18
Fibre Channel support 26
filesys clean command 24
filesys show space command 24
firewall requirements 27
full device 24

I
ifgroup
    multiple storage nodes 31, 35
    NIC usage 25

L
label
    device error 132
    template 108
legacy save sets migration 139
license requirements 20

M

max nsrmmd count 104
max sessions 104
memory requirement, storage node 23
migration
    legacy save sets 139
    scenarios 140
    tape to local devices 141
    tape to remote devices 140
    VTL to local devices 143
    VTL to remote devices 142
monitoring
    backup operations 122
    backup statistics 123
    devices 124
mounting a device 111
MTree 22
multiple datazones 34

N

naming guidelines 29
native non-deduplicated format 139
NDMP deduplication backup 112
net hosts add command 29
network connectivity 25
NetWorker
    application modules 19
    backup reports 129
    client 19
    storage node 19
NetWorker Management Console
    description 18
NetWorker Management Console See Console 19
NFS formats 23, 139
NFS service errors 132
NFS services 96, 100
NIC connections 25
NMC See Console 19
normal clone 68
nsrclone
    command 69
    script 93

O

openstorage license 96
optimized clone 17, 67

P

parallel cloning 17
pool
    associated with device 108
    configuration 112
    removing 24

R

recover
    deduplicated data 118
    disaster 119
redirect backups 136
remote device 104
replication license 96
replication, Data Domain 68
reports
    advanced 129
    basic 127
    configuring 125
    drill-down 128
    reviewing 127

S

Save Set Details
    NetWorker Backup Statistics Report 127
Save Set Summary
    NetWorker Backup Statistics Report 127
save streams 23
secure multi-tenancy 16, 22, 96, 100
seed
    as a data migration step 141
    before redirected backups begin 136, 139
    by native Data Domain replication 138
    diminishing value of 140
    not required on existing system 142
sessions 23
shared datazones 34
short names, preferred over long names 29
snmp
    add ro-community command 104
    add trap-host command 104
    traps 96, 104
statistics, monthly backup report 127
storage node
    dedicated 35
    replace 129
storage space, cleaning 24
storage unit
    maximum on Data Domain 22
    structure 22
synthetic full backup 17

T

target
    devices 109
    sessions 104
tenant unit 96
throughput, maximize 25
troubleshooting
    connectivity issues 131
    device access errors 132
    IP addresses 30
    network connections 131
    network issues 131
    no writable volume 132

V

version requirements 18
virtual tape library 23, 139
volume name 108
Volume Shadow Copy Service 112
volume sharing 23

volume,erase 115

W
wizard
Add New Host 104
Client Configuration 111
device configuration 100