

EMC[®] NetWorker[®] and EMC[®] Data Domain[®] DD Boost Deduplication Devices

Release 8.0

Integration Guide

P/N 300-013-562

REV 06

Copyright © 2010 - 2012 EMC Corporation. All rights reserved. Published in the USA.

Published November 2012

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 the technical documentation and advisories section on the EMC online support website.

CONTENTS

Preface

Revision History

Chapter 1 Introduction

Features.....	16
DD Boost integration	16
Client Direct with DD Boost backup	16
DD Boost storage structure.....	17
Cloning and remote storage of backups	17
NMC management of DD Boost operations	18
Features not supported by the integration	18
Network environment.....	18
Data Domain storage system.....	18
Data Domain Archiver system.....	19
NetWorker Management Console server.....	19
NMC user interface.....	19
NetWorker server.....	19
NetWorker client	19
NetWorker storage node.....	19
NetWorker application modules	20
Licensing	20
NetWorker licenses	20
Data Domain licenses	21

Chapter 2 Planning and Practices

About DD Boost devices.....	24
DD Boost devices on Data Domain systems.....	24
DD Boost devices on Data Domain Archiver systems	25
Memory and connection requirements	26
Devices in mixed environments.....	27
Devices and media pools	27
Reclaiming expired storage space	27
Removing a DD Boost device	27
Deduplication ratio and efficiency.....	28
Data types.....	28
Change rate	28
Retention period	28
Frequency of full backups.....	29
Host naming guidelines	29
IP addresses	30
Examples	30
Network connectivity and firewall.....	30
Network connectivity and capacity	31
Firewall requirements.....	31
Example backup environments	32
Client Direct deduplication environment.....	32

	Disaster recovery environment	33
	Cascaded replication environment	34
	Shared datazones environment.....	36
	Dedicated storage node environment.....	37
Chapter 3	Data Redirection and Migration	
	Redirecting backups to the new devices	40
	Redirect backups to DD Boost devices.....	40
	Migrating legacy save sets to the new devices	42
	Migration versus native Data Domain replication.....	42
	Migration methods.....	43
	Migrating legacy save sets to DD Boost devices.....	43
	Migration scenarios	44
	Migrating to a new DD system from traditional tape or disk storage	44
	Migrating to the existing DD system from non-integrated storage	46
Chapter 4	Software Configuration	
	Configuring Data Domain for NetWorker	50
	Configure the Data Domain system for NetWorker	50
	Configuring NetWorker for Data Domain	51
	DD Boost device performance considerations	51
	Methods	52
	Creating a DD Boost device with the wizard	52
	Configuring a DD Boost device with the property windows	54
	Add a host Data Domain system to NMC Enterprise view	54
	Configure a DD Boost device	55
	Create a volume label template for DD Boost devices	58
	Create pools to target DD Boost devices	59
	Label and mount devices on the storage node.....	61
	Configuring clients to back up to Data Domain	61
	Configure a backup client with the client wizard	62
	Configure a backup client with the property windows	63
	Deactivating and removing a device	66
	Convert a device to read-only	66
	Disable a device.....	66
	Delete a device	67
Chapter 5	Cloning and Replicating Data	
	About cloning and replicating with Data Domain	70
	Datazone requirement.....	70
	Clone pool	70
	Data Domain native replication support	70
	Clone reports	71
	Clone formats	71
	Clone-controlled replication format	71
	Regular clone format	71
	Clone methods.....	72
	Immediate clone	72
	Scheduled clone	72
	nsrclone command	72
	Clone requirements.....	73
	Scheduling a clone operation.....	74

	Cloning by pools	76
	Starting a scheduled clone operation on demand.....	77
	Manually start a clone operation from the Configuration window.....	77
	Manually start a clone operation from the Monitoring window	77
Chapter 6	Recovering Data	
	Recovering deduplicated data	80
	Recovery requirements.....	80
	Supported NetWorker recovery procedures.....	80
	Disaster recovery	80
	Causes of disaster.....	80
	Potential losses	81
	Disaster recovery requirements	81
	Disaster recovery scenarios.....	81
Chapter 7	Monitoring, Reporting, and Troubleshooting	
	Monitoring Data Domain events, statistics, and logs	84
	View backup statistics, log, and alerts in NetWorker.....	84
	View backup statistics in NMC.....	85
	View backup alerts (SNMP traps) in NMC.....	85
	Configure SNMP for Data Domain	86
	Backup, recovery, and clone reports.....	88
	Configure a report	88
	View a report.....	89
	Data Domain statistic reports	91
	Advanced Reporting	91
	Troubleshooting.....	91
	Name resolution issues.....	91
	Network connection issues	92
	Device access errors.....	92
	Backup fails for older NetWorker application modules	93
	Replace a failed or old storage node.....	93
Appendix A	Upgrading from NetWorker 7.6 SP1	
	Data Domain storage folder structure and session changes	98
	Planning the upgrade.....	98
	Glossary	
	Index	

FIGURES

	Title	Page
1	Client Direct (DFA) backup versus traditional storage node backup.....	33
2	Clone-controlled replication for disaster recovery	34
3	Clone-controlled replication cascaded to multiple Data Domain systems	36
4	Data Domain system shared across two NetWorker datazones	37
5	Single datazone with dedicated storage nodes and one high bandwidth link	38
6	Migration from traditional storage to NW-DD devices on a different storage node	45
7	Migration from traditional storage to NW-DD devices on the same storage node.....	46
8	Migration from VTL to NW-DD devices on a different storage node	47
9	Migration from VTL to NW-DD devices on the same storage node.....	48
10	DD Boost device name and access information.	56
11	Clone properties, general settings	75
12	NetWorker Administration showing DD Boost devices	85
13	Data Domain alerts to monitor.....	87
14	Report configuration	89

PREFACE

As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions 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 representative if a product does not function properly or does not function as described in this document.

Note: This document was accurate at publication time. New versions of this document might be released on the EMC online support web site. Check the EMC online support web site 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 NetWorker Data Domain devices within an EMC NetWorker backup and storage management environment.

Audience

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

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

Related documentation

This document is part of the NetWorker documentation set. The following EMC publications provide additional information:

- ◆ *EMC Information Protection Software Compatibility Guide* provides a list of client, server, and storage node operating systems supported by the EMC information protection software versions.
- ◆ *EMC NetWorker Installation Guide* provides instructions for installing or updating the NetWorker software for clients, console, and server on all supported platforms.
- ◆ *EMC NetWorker Cluster Installation Guide* contains information related to installation of the NetWorker software on cluster servers and clients.
- ◆ *EMC NetWorker Administration Guide* describes how to configure and maintain the NetWorker software.
- ◆ *EMC NetWorker Release Notes* contain information on new features and changes, fixed problems, known limitations, environment, and system requirements for the latest NetWorker software release.

- ◆ *EMC NetWorker Licensing Guide* provides information about licensing NetWorker products and features.
- ◆ *EMC NetWorker License Manager 9th Edition Installation and Administration Guide* provides installation, setup, and configuration information for the NetWorker License Manager product.
- ◆ *EMC NetWorker Error Message Guide* provides information on common NetWorker error messages.
- ◆ *EMC NetWorker Command Reference Guide* provides reference information for NetWorker commands and options.
- ◆ *EMC NetWorker Avamar Integration Guide* provides planning, practices, and configuration information for the use of Avamar deduplication nodes within an integrated NetWorker environment.
- ◆ *EMC VMware Integration Guide* provides planning, practices, and configuration information for the use of VMware and the vStorage API for Data Protection (VADP) within an integrated NetWorker environment.
- ◆ *EMC NetWorker Management Console Online Help* describes how to perform the day-to-day administration tasks in the NetWorker Management Console and the NetWorker Administration window.
- ◆ *EMC NetWorker User Online Help* describes how to use the NetWorker User program, which is the Microsoft Windows client interface for the NetWorker server, to back up, recover, archive, and retrieve files over a network.

Conventions used in this document

EMC uses the following conventions for special notices:

NOTICE

NOTICE presents information related to hazards.

Note: A note presents information that is important, but not hazard-related.

Typographical conventions

EMC uses the following type style conventions in this document:

Normal	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications
Bold	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and man pages Used in procedures for: <ul style="list-style-type: none"> Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus What the user specifically selects, clicks, presses, or types
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> Full titles of publications referenced in text Emphasis, for example, a new term Variables
Courier	Used for: <ul style="list-style-type: none"> System output, such as an error message or script URLs, complete paths, filenames, prompts, and syntax when shown outside of running text
Courier bold	Used for specific user input, such as commands
<i>Courier italic</i>	Used in procedures for: <ul style="list-style-type: none"> Variables on the command line User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user
[]	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 nonessential information 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, licensing, and service, go to the EMC online support (registration required) at:

[EMC online support](#)

Technical support — For technical support, go to EMC online support and select Support. On the Support page, you will see several options, including one to create a service request. 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.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to:

BSGdocumentation@emc.com

REVISION HISTORY

Email your clarifications or suggestions for this document to:

BSGdocumentation@emc.com

The following table lists the revision history of this document.

Revision	Date	Description of added or changed sections
06	November 27, 2012	Updated the following sections: <ul style="list-style-type: none">• Added “Retention tier storage restriction” on page 17 and removed “Data Domain Archive system” formerly on page 18.• Clarified “NetWorker 8.0 device changes and downgrade restrictions” on page 16.• Clarified “Data movement between tiers” on page 25.• Clarified “Clone-controlled replication format” on page 71.
05	November 16, 2012	Updated the following: <ul style="list-style-type: none">• Clarified intro of “Configure a backup client with the client wizard” on page 62.• Removed unnecessary note in “About cloning and replicating with Data Domain” on page 70• Clarified “Data Domain native replication support” on page 70.• Clarified “Data Domain storage folder structure and session changes” on page 98.
04	October 10, 2012	Updated the following: <ul style="list-style-type: none">• Updated “Configure a backup client with the property windows” on page 63 - Details on NDMP deduplication clients.• Changed the book title to: EMC NetWorker and EMC Data Domain Deduplication Devices Integration Guide.
03	August 8, 2012	Updated “Network connectivity” on page 26 - DD Boost supports LAN, WAN, MAN.
A02	July 16, 2011	Updated the following: <ul style="list-style-type: none">• “Delete a device” on page 67 - Clarified procedure.• “Clone requirements” on page 73 - Data Domain encryption, if used, must be consistent on both the source and target systems.
A01	June 21, 2012	First release of this document for Restricted Availability (RA).

CHAPTER 1

Introduction

This chapter includes the following sections:

- ◆ Features 16
- ◆ Network environment 18
- ◆ Licensing 20

Features

Review the following features of the EMC® NetWorker® release 8.0 integration with the EMC Data Domain® software to help plan your deployment.

Note: An upgrade from the NetWorker 7.6 SP1 integrated solution with Data Domain requires special procedures. [Appendix A, “Upgrading from NetWorker 7.6 SP1,”](#) provides details.

DD Boost integration

NetWorker software uses DD Boost components to integrate with logical storage devices (DD Boost devices) on Data Domain systems and perform data deduplication backups. Data deduplication dramatically reduces the amount of data that is sent to the storage devices and the bandwidth required for the data transmission.

Each NetWorker client and storage node installed with DD Boost can perform up to 60 concurrent backup and recovery sessions (save streams) per DD Boost device. This high use of sessions reduces the number of devices needed and thereby reduces the impact on the Data Domain system performance and maintenance. The resulting performance gain can compare favorably against conventional advanced file type device (AFTD) or virtual tape library (VTL) interfaces that do not allow the same high rate of sessions.

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

The DD Boost software consists of the following two components:

- ◆ The distributed segment processing (DSP) component, which reviews the data that is already stored on the Data Domain system and sends only unique data for storage.
- ◆ The DD Boost library API, which enables the NetWorker software to communicate with the Data Domain system.

NetWorker 8.0 device changes and downgrade restrictions

NetWorker 8.0 provides both write and read functionality on existing and new DD Boost devices. This is unlike earlier releases that used separate read-only DD Boost mirror devices for restore operations. The NetWorker 8.0 installation removes legacy read-only DD Boost mirror devices.

NOTICE

After an upgrade to NetWorker 8.0, a downgrade to an earlier NetWorker release is *not* supported. After the downgrade, DD Boost devices created with NetWorker 8.0 will be unavailable and legacy devices and data will require manual reconstruction.

Client Direct with DD Boost backup

Client Direct, also known as direct file access (DFA), is a NetWorker feature that enables clients with IP network access to the Data Domain system to send backup data directly to AFTD or DD Boost storage devices, bypassing the NetWorker storage node. The storage node manages the devices used by the backup clients but does not handle the backup data. The Client Direct feature is used by default when it is available.

The Client Direct feature leverages the DD Boost distributed segment processing (DSP) software component that is installed as part of the NetWorker 8.0 client software. During backup, the DSP software on the client deduplicates the backup data before the Client Direct component sends the deduplicated data directly to the DD Boost devices. By working together, the DD Boost and Client Direct features can provide highly efficient data deduplication, transmission, and storage for multiple concurrent client backup operations. Backup bottlenecks are removed from the storage node, and network bandwidth is better utilized.

Traditional storage node backups can be recovered by using Client Direct functionality, if it is available. If not, you can recover by using traditional storage node functionality.

DD Boost storage structure

A Data Domain system stores deduplicated NetWorker backups on special DD Boost (NetWorker Data Domain) storage devices. The NetWorker server, storage nodes, and Client Direct clients can access DD Boost devices, which are structured as follows:

- ◆ Storage on the Data Domain system is based on storage unit (SU) parent folders that contain the DD Boost devices. These parent folders are also called managed trees (MTrees). By default, each SU handles a single NetWorker datazone and is named after the short hostname of the NetWorker server.
- ◆ Data Domain Operating Systems 5.0 and 5.1 should use no more than 14 active SUs.
- ◆ Each DD Boost storage device is identified by an SU child sub folder location and is associated with a single NetWorker storage volume. You can configure multiple DD Boost device hosts to share a single volume, which can result in performance gains in some environments.
- ◆ There is no recommended limit to the number of DD Boost devices that you can create, but for best performance, use fewer devices and more backup sessions on each device.

[“Configure the Data Domain system for NetWorker” on page 50](#) provides an optional advanced step to manually create device folders if the wizard is not used.

Retention tier storage restriction

The Data Domain Extended Retention software option, available with DD OS 5.2 and later, extends the Data Domain storage structure with internal tiers that enable long-term retention of deduplicated backup data. Data Domain operations are used to migrate the data from the active tier to the retention tier.

On a Data Domain system with Extended Retention software enabled, NetWorker software works with the active tier only and is not aware of any migration activity between the internal storage tiers.

Cloning and remote storage of backups

Backups that are stored on a Data Domain system can be copied to remote storage for added protection and efficient disaster recovery. These copies are configured and launched by the NetWorker Management Console (NMC) clone feature and are indexed and retained according to NetWorker storage policies.

There are two types of clone operations, both of which require the DD Boost software on the NetWorker storage nodes:

- ◆ The clone-controlled replication (CCR) or “optimized clone” operation is typically used to replicate data from a DD Boost device to another DD Boost device at a different location. This operation preserves the deduplicated data format and minimizes bandwidth usage between the Data Domain systems.
- ◆ The clone to native format operation is used to clone data from Data Domain storage to traditional disk or tape storage. This operation reverts the data to its native non-deduplicated format, which is required for the data to be recoverable from a traditional disk or tape device.

NMC management of DD Boost operations

NMC software manages all the configuration, monitoring, and reporting operations for DD Boost devices. The NMC server must have network access to each Data Domain system that it manages.

The NMC Device Configuration Wizard simplifies the configuration of storage devices, backup clients, storage pools, volume labeling, and save set cloning.

Features not supported by the integration

The integration does not support the following Data Domain or NetWorker features:

- ◆ Data Domain Retention Lock.
- ◆ Data Domain MTree quotas (added in DD OS 5.2).

MTree quotas manage the usable capacity of storage units and impact the associated DD Boost devices. When a storage unit reaches its MTree quota setting, a running DD Boost backup operation terminates. The Data Domain documentation provides details.

Network environment

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

Data Domain storage system

A Data Domain system that runs DD OS 5.0 or later software can store deduplicated backup or cloned data on NetWorker DD Boost storage devices that are accessed by the a NetWorker client or storage node. For full compatibility with NetWorker software, install the latest upgrade of DD OS software.

A Data Domain system can support mixed environments that may include DD Boost devices, VTLs, and CIFS or NFS AFTD disk configurations. Additional licenses might be required to enable the functionality of the Data Domain system.

Data Domain Archiver system

Data Domain Archiver is a specific model of Data Domain storage system that has been designed for the management of long-term data retention. NetWorker DD Boost devices may be created on a Data Domain Archiver that runs DD OS 5.0 or later software. The archive tier does not require additional capacity licensing.

A Data Domain Archiver system can support mixed environments that may include DD Boost devices, VTL, and CIFS or NFS AFTD disk configurations.

NetWorker Management Console server

The NMC server provides a user interface for NetWorker configurations and monitoring. The NMC server must run NetWorker 8.0 or later software and have network access to NetWorker storage nodes and Data Domain systems.

NMC user interface

The NMC user interface can be accessed and operated from a computer with an Internet web browser and Java Runtime Environment (JRE) software installed. Multiple users can access the NMC server concurrently from different web browsers.

The *EMC NetWorker Installation Guide* provides information on supported web browsers and versions of the JRE.

NetWorker server

NetWorker 8.0 or later software must be installed on the NetWorker server that stores, clones, or recovers data for its clients on Data Domain systems.

The *EMC NetWorker Release Notes* provides details on supported Data Domain features.

NetWorker client

A NetWorker client may be a supported host whose data requires backup and restore services. The NMC server, NetWorker server, and NetWorker storage nodes are also NetWorker clients. The NetWorker client software must be installed and running before DD Boost devices are configured.

- ◆ NetWorker clients that use Client Direct deduplication must have network access to the Data Domain system where their data will be stored and must have NetWorker 8.0 or later client software installed. This software includes the DD Boost 2.4.1 or later library. Client Direct is enabled by default.
- ◆ Other NetWorker clients that do not use Client Direct deduplication may use a supported earlier NetWorker release, although release 7.6 and later is recommended.

The *EMC Information Protection Software Compatibility Guide* provides information on supported releases.

NetWorker storage node

Backup data from NetWorker clients that do not use Client Direct deduplication may be deduplicated on a NetWorker storage node before it is sent to the Data Domain system for storage.

NetWorker 8.0 or later software must be installed on each storage node that stores or clones data on Data Domain systems. All storage nodes within the same NetWorker datazone should be of the same release and service pack level and must have access to the Data Domain systems where the data will be stored.

NetWorker application modules

Client Direct deduplication backup and recovery is supported for clients with NetWorker application modules, such as NMDA, NMM and NMSAP, provided the client has access to the Data Domain system and the appropriate version of NetWorker client and application module is installed. The application module release notes provides details.

Licensing

A fresh installation of a NetWorker release (not an upgrade) enables evaluation of all features, including Data Domain features, which require no enabler (license key) for 30 days. You can extend this evaluation period by 15 additional days by typing the word `grace` as an authorization code before the end of the 30-day period. After the evaluation period, you cannot perform a backup unless you apply permanent enablers.

NetWorker licenses

Two alternative types of licensing can enable the NetWorker server to interact with a Data Domain system, as described in the following sections. Upgrade enablers might also be required for earlier NetWorker releases.

Note: Unlike earlier releases, NetWorker 8.0 Data Domain storage software enablers support multiple interfaces and network identities. In earlier releases, it was recommended to use an `ifgroup` software configuration of NICs on the Data Domain system to allow multiple storage nodes to use a single network identity and a single Data Domain Storage System Enabler. Although `ifgroup` is no longer needed to share licensing, bandwidth aggregation is still a valid and recommended use case of `ifgroup` for DD Boost devices.

The *EMC NetWorker Licensing Guide* and your EMC sales representative can provide licensing details.

Traditional licensing

Traditional licensing requires the NetWorker server to be installed with an individual enabler for each client, module, storage device, or other component or feature, including Data Domain storage. For Data Domain storage, this licensing requires an enabler for Data Domain usage and an enabler that specifies the capacity of this usage as follows:

- ◆ The Data Domain Storage System Enabler (DD Boost Enabler) enables an unlimited number of Data Domain systems and DD Boost devices within the datazone.
- ◆ The Data Domain Capacity Entitlement Enabler enables use of the Data Domain systems within the datazone up to the purchased capacity.

NetWorker Capacity Entitlement licensing

NetWorker Capacity Entitlement licensing enables an unlimited number of clients, modules, storage devices, and other NetWorker associated components or features to be used within a datazone, including Data Domain storage. This type of licensing requires the following enablers:

- ◆ The NetWorker Source Capacity Datazone Enabler enables all of the NetWorker components or features in the datazone, including Data Domain storage.
- ◆ The Tiered Capacity Entitlement License Enabler enables protection up to the purchased capacity for all of the supported components or features in the datazone, including Data Domain storage.

The *EMC NetWorker Licensing Guide* provides requirements for environments that use a NetWorker Capacity Entitlement license. Contact your EMC sales representative for details.

Data Domain licenses

You must enable the Data Domain or Data Domain Archiver system with the following licenses for NetWorker DD Boost operations:

- ◆ DDBOOST license
- ◆ REPLICATION license (if CCR will be used)

To verify if the licenses are enabled on the Data Domain system, type the **license show** command. The command output includes OPENSTORAGE if a DD Boost license is installed, and additionally REPLICATION if a replication license is installed. [“Configure the Data Domain system for NetWorker” on page 50](#) provides details.

For upgrade and other inquiries, contact the Data Domain service and support portal at <https://my.datadomain.com>.

CHAPTER 2

Planning and Practices

This chapter includes the following sections:

- ◆ About DD Boost devices 24
- ◆ Deduplication ratio and efficiency 28
- ◆ Host naming guidelines 29
- ◆ Network connectivity and firewall 30
- ◆ Example backup environments 32

About DD Boost devices

The NetWorker 7.6 SP2 or later software enables the integration of NetWorker features with Data Domain systems by storing backup or clone data on NetWorker Data Domain (DD Boost) devices.

DD Boost devices on Data Domain systems

Data Domain systems that run DD OS 5.0 or later software support DD Boost devices that run NetWorker 7.6 SP2 or later software.

Note: The NetWorker 8.0 release, unlike earlier releases, does not create read-only mirror devices. [“NetWorker 8.0 device changes and downgrade restrictions” on page 16](#) provides details.

DD Boost storage structure and restrictions

The NetWorker integration environment creates storage unit (SU) folders on the Data Domain system that are the parent folders for the DD Boost devices, which appear as sub folders. On Data Domain systems with DD OS 5.0 and later, each SU appears as a separate logical partition, also called an MTree (managed tree).

The Data Domain folder structure has the following features and constraints:

- ◆ A single SU folder is configured per NetWorker datazone. By default, the Device Configuration Wizard names the SU folder with the short hostname of the NetWorker server.
- ◆ You should not use more than 14 active SU folders on a single Data Domain system. You can define up to 99 SUs.

Note: Data Domain systems with DD OS 5.0 or later that have more than 14 active SU folders for DD Boost devices can suffer impaired performance and become suspended.

- ◆ There is no recommended limit to the number of DD Boost devices that you can create. However, increasing the number of devices can impact Data Domain performance and maintenance.

Device access configuration

The configuration of a DD Boost device identity is designed to handle multiple accesses:

- ◆ Each device is identified with a single NetWorker storage volume.
- ◆ Each device is defined by a single access path, although the access path can be represented in different ways for different client hosts.
- ◆ Each device can be created multiple times, with each copy of the device being given a different name as an alias.

[“Configure a DD Boost device” on page 55](#) provides details on device access information.

Device performance

DD Boost devices use multiple concurrent **nsrmmd** (data mover) processes per device and multiple current save sessions (streams or threads) per **nsrmmd** process. Optimal device configuration for backup or clone operations reduces the number of active devices required and thereby reduces the impact on Data Domain system performance and maintenance.

You should balance session load among the available DD Boost devices so that new sessions attach to devices with the least load. To enable optimum performance, you can adjust the device Target sessions, Max sessions, and Max nsrmmd count attributes.

[“Configure a DD Boost device” on page 55](#) provides details on session settings.

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

DD Boost devices on Data Domain Archiver systems

You can treat SU folders or MTrees and DD Boost devices on Data Domain Archiver systems the same way as they are on standard Data Domain systems, with the considerations described in this section.

The Data Domain documentation provides details on Data Domain Archiver features.

Data Domain Archiver tier structure

On DD Archiver systems, data can be stored in two tiers of a storage units, the active tier and the archive tier:

- ◆ All incoming data is first placed in the active tier file system, similar to standard Data Domain systems. This tier is used for short-term data storage and recovery.

The active tier can be used for backup provided appropriate data movement and retention policies are used for the devices. The creation of separate SUs for backup data is recommended.

- ◆ Data on the active tier can subsequently be moved to the archive tier by the DD Archiver, based on data movement policies applied at the SU level.

DD OS 5.0 and 5.1 support up to 14 active and 99 defined storage units on a DD Archiver system. There is no recommended limit on the number of DD Boost devices that you can create on the system.

Data movement between tiers

Each storage unit has a single data movement policy that applies to all of the devices it serves within its corresponding NetWorker datazone, whether they are NetWorker AFTD or DD Boost devices.

A Data Domain SU data movement policy determines when data is moved from devices in the active tier to devices in the archive tier. Policies should typically not move data to the archive tier for short periods, less than 30 days retention. The policy and movements are internal to the DD Archiver system and the NetWorker software has no awareness of them.

You can assign alternative data movement policies to the client data by using the Device Configuration Wizard to create additional SUs specific to this purpose. For example, data can be written to different devices in separate SUs with different archive policies. Additionally, data can be moved within the same DD Archiver system by using clone-controlled replication (CCR).

For CCR on the same DD Archiver system, you should replicate between two different SUs so that different retention policies can be applied and the data can be managed efficiently.

Memory and connection requirements

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

Memory

Ensure the following memory requirements:

- ◆ Allowing for other types of devices and services on a typical storage node, a storage node should have a minimum of 8 GB of RAM if hosting DD Boost devices.
- ◆ DD Boost clients require a minimum of 4 GB of RAM at the time of backup to ensure optimum performance for Client Direct backups.
- ◆ Each DD Boost device requires an initial 24 MB of RAM on the storage node and Client Direct client. Each DD Boost save session requires an additional 24 MB. To run 10 sessions requires 24 + 240 MB. The default max sessions of 60 sessions per DD Boost device requires 24 + 1440 MB.

Network connectivity

Ensure the following network access:

- ◆ The storage nodes and all Client Direct clients must have network access to the Data Domain system where the data will be stored and recovered from.
- ◆ If the NetWorker server and the NMC server are located on different systems, then both require access to the Data Domain system for administration and monitoring purposes.
- ◆ NetWorker support for DD Boost devices does not distinguish network types (LAN, WAN, or MAN) and can successfully operate where packet loss is strictly 0% and latency is less than 20 milliseconds.

Note: Currently, the DD Boost devices support only an IP network and do *not* support SAN (fiber channel) data transport.

[“Network connectivity and firewall” on page 30](#) provides further details on network bandwidth and connections.

Devices in mixed environments

A Data Domain system can support mixed NetWorker environments, which can consist of DD Boost devices, VTLs, and CIFS or NFS AFTD disk configurations. Each different storage environment should use a different interface connection, either NIC for IP or FC port for SAN data transport. Data stored on these other environments can be migrated to DD Boost devices through a NetWorker clone process.

Devices and media pools

Media pools are used to direct client backups or clone copies of backups to specific storage devices. You should configure pools for DD Boost devices with the following important consideration.

A pool that contains the DD Boost devices should *not* contain any other type of storage media and should *not* contain devices on more than one Data Domain system. This ensures that the backup window and deduplication ratio are not impaired. It also ensures that each backup is stored on a single Data Domain system, which simplifies recovery and facilitates Data Domain file replication, if desired.

Periodically, review and remove unused pools if they are no longer relevant to the storage environment.

Reclaiming expired storage space

When a backup on a Data Domain system reaches its retention expiry date, all its references to the data become invalid. However, unlike traditional storage systems, space is not immediately freed up on the device because other backups or clients might 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 result in 1 GB of space being made immediately available on the Data Domain system.

The DD OS **filesys show space** or **df** command shows the data that is eligible to be deleted from the Data Domain system under the Cleanable GiB column.

A review of deduplication references and storage space occurs during a periodic cleanup operation, which is scheduled by default to run once per week.

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

The DD OS **filesys clean** command describes all the available options for managing how the storage space for redundant or expired data, such as expired NetWorker save sets, can be reclaimed and recycled.

Removing a DD Boost device

You must use special procedures to remove DD Boost devices.

[“Delete a device” on page 67](#) provides details.

Deduplication ratio and efficiency

Deduplication efficiency is typically measured by the deduplication ratio. This ratio represents the reduction in storage space that results from the data deduplication/compression technology. Ratios of 20:1 are considered to be broadly achievable and reductions of even 5:1 are extremely valuable.

A number of factors can contribute to the deduplication ratio, including retention periods, the type of data being backed up, change rate, the frequency of full backups, and the use of encryption and compression.

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

Data types

Some types of user data such as text documents, slide presentations, spreadsheets, most database types, source code, and e-mail are known to contain redundant data and are good deduplication candidates.

Some other types of data such as 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 will typically produce high deduplication ratios if the data has a low change rate.

Change rate

Data that does not change much between backups (low change rate) produces high deduplication ratios and is a good candidate for deduplication. Data that has already been stored will not be stored again.

When a completely new save set is deduplicated, the number of unique blocks within it 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 percent) data reduction.

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

Retention period

The deduplication ratio increases with longer data retention periods. The longer the stored save sets are retained, the greater the chances that identical data already exists in storage that can be used to deduplicate each subsequent backup, and the greater the efficiency of the storage usage.

Frequency of full backups

Frequent full backups result in high deduplication ratios, but also increased data processing on the NetWorker storage node. For example, compare deduplication backups performed on a daily full basis with those performed on a weekly full with daily incremental basis. 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 because the storage node sends 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 should 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.

The following guidelines are offered to help you create consistent, easily identifiable hostnames that facilitate configuration, reporting, and troubleshooting in the NetWorker Data Domain environment:

- ◆ Create names that are unique across all NetWorker datazones. They should identify the network role, such as administration, backup, cloning, or production, with possibly a location or server name.
- ◆ Use a single hostname associated with each NIC or IP interface within the same NetWorker datazone.
- ◆ Include source and target abbreviations in names to quickly identify whether 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. Ensure that these names are included in the Data Domain `/etc/hosts` file.
- ◆ Specify all aliases, such as long and short names, IP address, and so on for the NetWorker server and storage nodes in their respective Client resources (Globals 1 of 2 tab).
- ◆ Ensure that all hostnames in the network can be consistently resolved from multiple locations in both directions. For example: shortname to IP, longname to IP, IP to shortname, IP to longname.
- ◆ 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 might not fit into some views. Examples of a long and short name:


```
NWDD365-1.burloak.lab.mycorp.com:/NWDZ_Dr1
NWDD365-1:/NWDZ_Dr1
```
- ◆ Use standard alphanumeric characters, including dot “.” and underscore “_” with no spaces or special characters, except that pool and hostnames may *not* use underscores “_”.

- ◆ Use fixed formats (text field length and case) and include leading 0s in numbers, with a maximum of 50 characters.
- ◆ Avoid the use of dates in names where the dates might change or become meaningless in the future.

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 IP addresses to be used to interface with an ifgroup for Advanced Load Balancing and Link Failover features.
- ◆ For clone-controlled replications, the hosts file on the source Data Domain system should list the IP address of the target Data Domain system. Otherwise, the clone-controlled replication will use the same network that is used to access the source and target systems for backup and restore operations.

The Data Domain documentation provides details.

Examples

The following are some example name formats.

DD Boost devices

Format: *DD-DDsystem name-Device00-99*

For example: DD-Tenendo-Device01

Folders on Data Domain system

DD Boost device names should refer to the NetWorker storage node and indicate whether they are for backup (BACK) or clone (CLON) operation.

Format: *storage node_BACK or CLON_DD00-99*

For example: DZBurl_BACK_DD01

Volume labels for Data Domain

Format: *DD-MEDIA or 5-letter descriptor-000-999*

For example: DD-MEDIA-001

Network connectivity and firewall

Because DD Boost devices are network-based, supported on Ethernet IP networks, consideration should be given to hostname resolution, connectivity and capacity, NFS access, and firewall requirements.

The network hostname resolution method depends on your environment and manufacturer recommendations. [“Host naming guidelines” on page 29](#) provides details.

Network connectivity and capacity

The recommended minimum network connectivity to the Data Domain system is to use two 1 GbE network links: one of which is dedicated to administration and the other to data backup. Network connections and types can improve on this depending on the Data Domain system model.

You can maximize throughput of the Data Domain system by using multiple connections or high bandwidth connections. For example, you can use multiple 1 GbE connections for dedicated storage nodes and storage devices. 10 GbE connections are also available and these can be used instead of, or in addition to, 1 GbE interfaces.

In environments where 10 GbE connectivity is not available or cost prohibitive, two alternatives are available:

- ◆ Use a dedicated 1 GbE connection from a storage node directly to the Data Domain system. This provides a private, high-bandwidth data connection and avoids the latency and complexity of a shared Ethernet connection. However, a separate traditional Ethernet connection is also required for administration and NMC access.
- ◆ Use two or more NICs on the Data Domain system with 1 GbE connections aggregated together by using the Data Domain **ifgroup** command. This will provide increased capacity and can offer some resiliency. The Data Domain system provides automatic Advanced Load Balancing and Link Failover for NIC connections.

Firewall requirements

Regardless of the network connections used, communication through a firewall requires the use of specific ports and protocols for backup, monitoring, and replication across sites. Ensure that the following firewall ports are open between the Data Domain, NetWorker, and NMC servers:

- ◆ TCP 111 (NFS portmapper)
- ◆ TCP 161 (for NMC server to query for alerts and statistics)
- ◆ TCP 162 (SNMPTRAP for NMC server to monitor status and events)
- ◆ TCP 2049 (NFS)
- ◆ TCP 2051 (Replication, if clone-controlled replication is used, Data Domain to Data Domain systems)
- ◆ TCP *xxxx* (select a random port for NFS mountd, 2052 is the default)
On the Data Domain system, type the following command from SE mode:

```
# nfs set mountd-port xxxx
```

Restart the Data Domain system.

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

Example backup environments

This section provides some examples of how the NetWorker Data Domain integration can be deployed in backup environments. The use of two interfaces, 1 GbE for administration and 10 GbE for data, is recommended, as shown in the examples:

- ◆ [“Client Direct deduplication environment” on page 32](#)
- ◆ [“Disaster recovery environment” on page 33](#)
- ◆ [“Cascaded replication environment” on page 34](#)
- ◆ [“Shared datazones environment” on page 36](#)
- ◆ [“Dedicated storage node environment” on page 37](#)

Client Direct deduplication environment

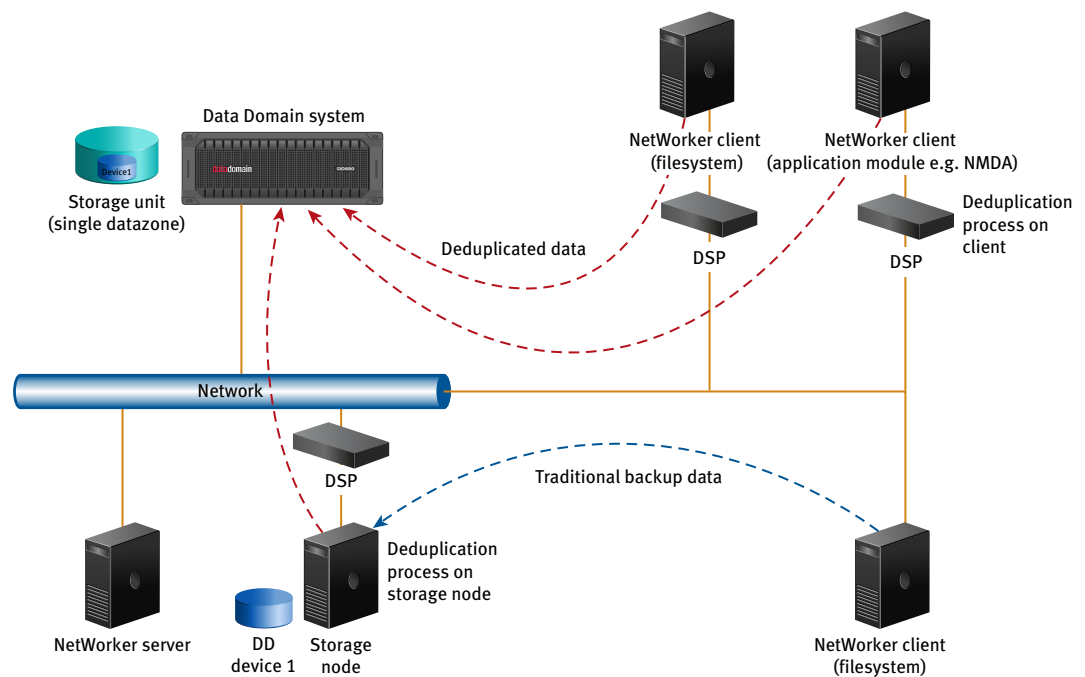
[“Client Direct with DD Boost backup” on page 16](#) describes the Client Direct feature, which leverages client DSP software to send deduplicated data directly from a client to a Data Domain system and bypass the storage node.

Client Direct functionality is enabled by default in the Client resource settings for backup and recovery operations. If the client does not have a direct connection to the Data Domain system, then the traditional storage node workflow is automatically used.

[Figure 1 on page 33](#) shows an example Client Direct environment.

Client Direct deduplication provides the following advantages:

- ◆ Deduplication on the client host dramatically reduces the bandwidth for the backup data transfer.
- ◆ A single storage volume can be shared among DD Boost devices and among multiple backup hosts and storage nodes. Configuring multiple hosts and multiple sessions per device, instead of creating multiple devices, can improve performance and maintainability.
- ◆ Client Direct deduplication offers an alternative to an environment that uses a NetWorker client installed with a dedicated NetWorker storage node, as described in [“Dedicated storage node environment” on page 37](#). In the dedicated storage node environment, additional licensing and configuration is required, but might not be supported or suitable for the clients and applications that are to be protected.



GEN-001725

Figure 1 Client Direct (DFA) backup versus traditional storage node backup

Disaster recovery environment

An example of a disaster recovery (DR) environment would be to use clone-controlled replication to copy individual save sets or complete volumes from a primary Data Domain storage system to a secondary Data Domain system at a geographically distant location. Each cloned replication (optimized clone) is a complete and independent copy of the source deduplicated data that is managed by NetWorker policies. For added protection, some or all of the stored data may be further cloned from the secondary system to tape storage.

Both the primary and secondary Data Domain systems must be configured, enabled, and managed within a single datazone. There may be either a single or separate storage nodes for local and remote operations within the datazone. Target devices must be properly configured on the secondary Data Domain system.

Note: Clone-controlled replication across datazones or to Data Domain devices that are not managed by NetWorker are *not* supported.

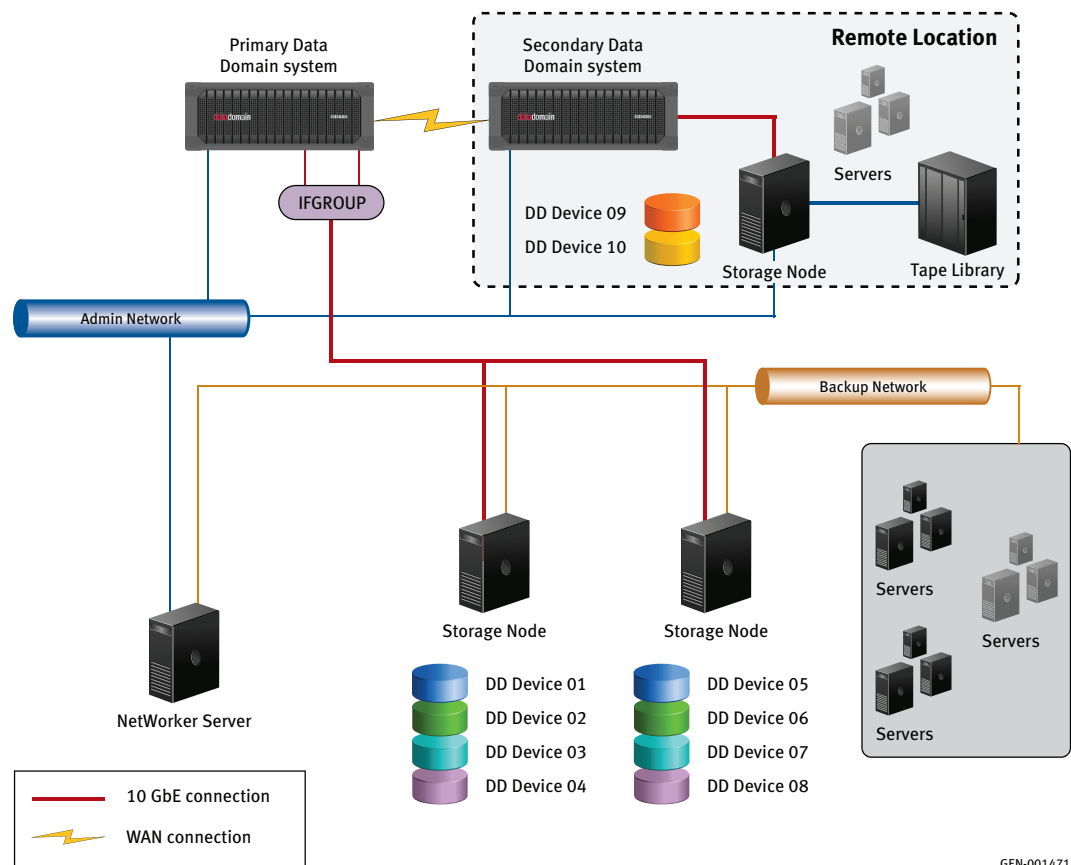
Figure 2 on page 34 illustrates an example of a disaster recovery environment. Two Data Domain Storage System Enablers are required on the NetWorker server, one for the primary and one for the remote secondary Data Domain systems:

1. The NetWorker server initiates the backup of client groups within its datazone.
2. Two storage nodes in the datazone write the backup data to media pools, which target specific NetWorker Data Domain storage devices on the primary system.

- The storage nodes communicate with the primary Data Domain system and store deduplicated backup data on the system's devices over dedicated 10 GbE network connections.

Note: An **ifgroup** software configuration of 1 GbE NICs on the Data Domain system will allow multiple storage nodes to use the same network identity and this bandwidth aggregation can improve performance for DD Boost devices.

- Clone-controlled replication is used 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 the purpose of disaster recovery.
- As a further option, a NetWorker storage node, attached to the secondary Data Domain system, is then able to create an additional NetWorker clone copy of the data for one of the backup groups and stores it on traditional disk or tape. The data in this copy is reverted to its native non-deduplicated format, which is necessary for storage on traditional disk or tape.



GEN-001471

Figure 2 Clone-controlled replication for disaster recovery

Cascaded replication environment

A variation of the disaster recovery environment is the cascade replication environment. Once a deduplicated backup has successfully completed, it may be used to create multiple clone copies in other locations, either simultaneously from the original deduplicated backup or in sequence from clone-controlled replications of the original.

Each clone replication is a complete and independent copy of the source backup, complete with NetWorker policies. There is no limit to the number of clone copies that can be made, provided that the source save set used for each clone has successfully completed.

As with the previous example, each Data Domain system in the environment must be configured, enabled, and managed within the same single NetWorker datazone. Target devices must be properly configured on the Data Domain systems that receive the clone copies.

[Figure 3 on page 36](#) illustrates an example of a cascaded replication environment with three separate Data Domain systems at three different sites. Three Data Domain Storage System Enablers are required on the NetWorker server, one for each Data Domain system:

- ◆ 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 or country.
- ◆ The third site is a disaster recovery (DR) site located at a geographically distant location where the communication links would be limited or subject to significant latency and bandwidth restrictions. This might be in a different country or 250 kilometers (150 miles) or more distant from either of the other two sites.

This example environment operates as follows:

1. The NetWorker server initiates the backup of production site client groups within its datazone.
2. The production site storage node assigns the backup data to media pools, which target specific NetWorker Data Domain storage devices on the primary Data Domain system.
3. The storage node communicates with the primary Data Domain system and stores deduplicated backup data on the system's devices over dedicated 10 GbE network connections.
4. Once the backup successfully completes, clone-controlled replication is used to store optimized clone copies of the backups on the primary Data Domain system over a network to Data Domain systems at other locations. These clone copies may be created by either a sequential or concurrent method:
 - In the *sequential* method, only one clone operation is performed 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, the original backup on the primary Data Domain system is used to create an optimized clone copy on a local secondary Data Domain system. Once this has been successfully completed, this copy is used to create an additional optimized clone copy on the geographically distant Data Domain system.

[Figure 3 on page 36](#) illustrates data paths 1a and 1b for this method.

- In the *concurrent* method, it may be possible to perform the clone operations simultaneously. This method impacts the production system and requires more replication bandwidth.

Note: The performance required for the concurrent method depends on many factors and would need to be validated by testing at the individual sites.

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

Figure 3 on page 36 illustrates data paths 2a and 2b for this method.

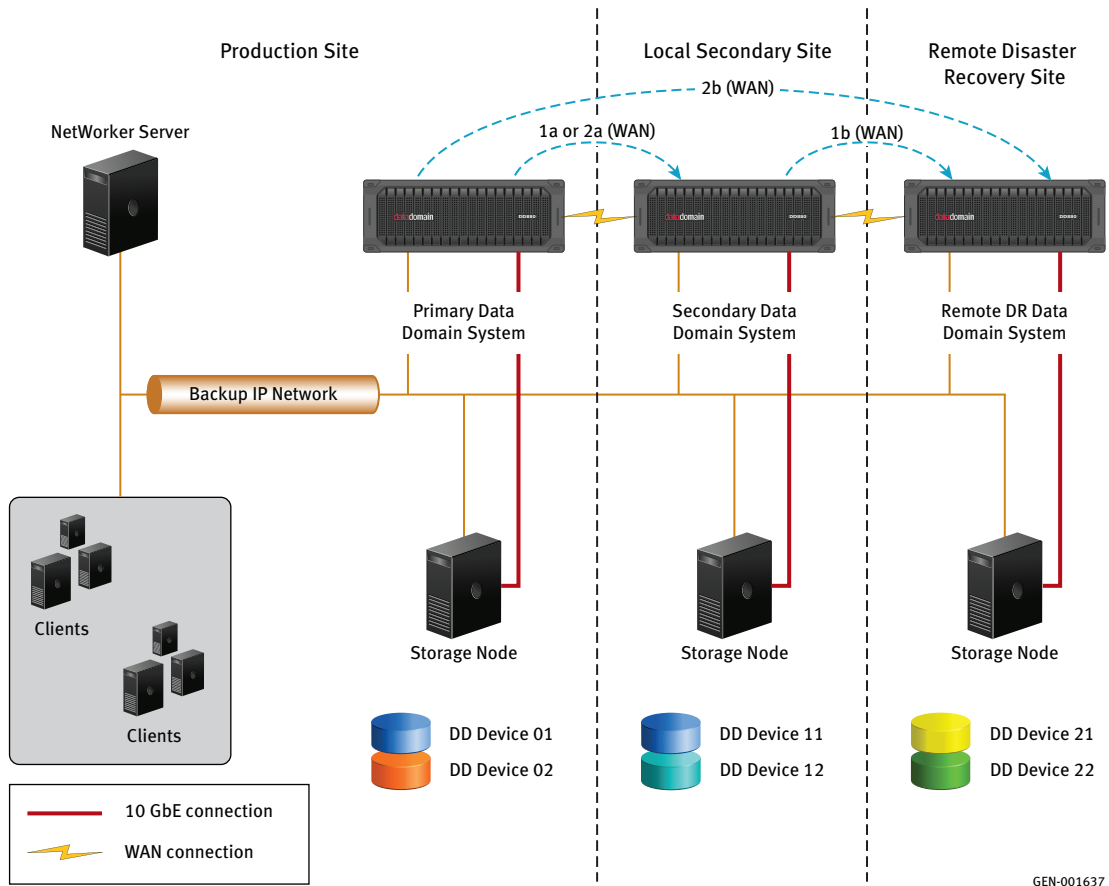


Figure 3 Clone-controlled replication cascaded to multiple Data Domain systems

Shared datazones environment

Backups from two datazones may be stored on a single Data Domain system. In this configuration, consideration needs to be given to dividing the stream counts and memory resources to manage the two datazones as separate entities and not allow one datazone to impact the performance of the other datazone. The total number of streams and devices cannot exceed the total capacity of the Data Domain system.

Figure 4 on page 37 illustrates an example of a dedicated 10 GbE network connection that is shared by three storage nodes in two 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 initiate backups of client groups within their respective datazones.

- The three storage nodes write the backup data to specific pools, which are targeted to storage devices on the Data Domain system.

Note: A DD Boost device *cannot* be shared across datazones.

“DD Boost devices on Data Domain systems” on page 24 provides details on save stream usage and memory requirements. “Traditional licensing” on page 20 provides details on Data Domain system enablers:

- The storage nodes store the deduplicated backup data on the DD Boost devices over a shared 10 GbE connection.
- Additional backup to tape storage may be made either directly from a storage node or by a NetWorker clone operation from the Data Domain system.

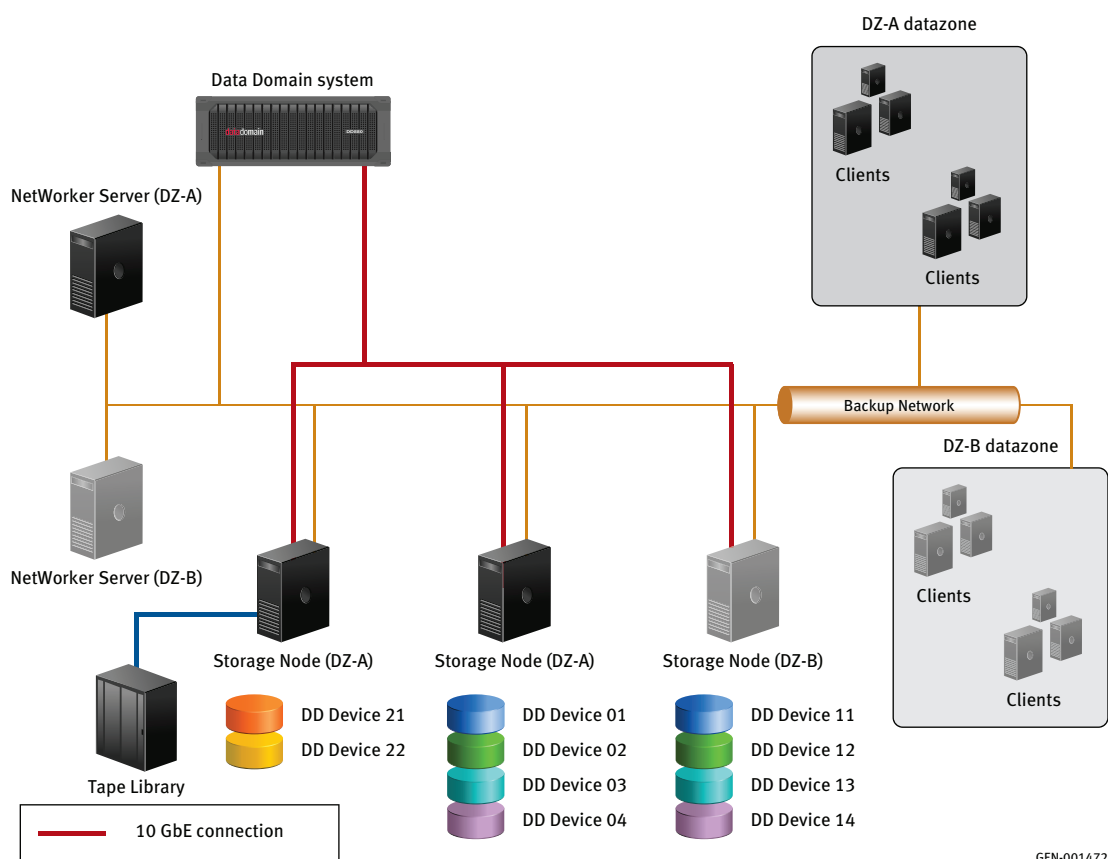


Figure 4 Data Domain system shared across two NetWorker datazones

GEN-001472

Dedicated storage node environment

Deduplication backups are supported for high-volume NetWorker clients that each have a dedicated storage node. For example, a dedicated storage node may be deployed on a client that runs NetWorker Module for Database Applications (NMDA). Other NetWorker application modules are also supported in this configuration with the exception of NMM 2.2 SP1 or earlier, which do not support dedicated storage node operation.

This environment can coexist with client groups that use shared NetWorker storage nodes in the same datazone. However, because this is a private network, the connection (and the devices it uses) would not be available to other NetWorker clients.

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

Figure 5 on page 38 illustrates a mixed environment of shared and dedicated storage nodes:

1. The NetWorker server initiates backups of the various regular client groups as well as the module clients on the dedicated storage nodes.
2. The storage nodes write the backup data to specific pools, which are targeted to storage devices on the Data Domain system.

“DD Boost devices on Data Domain systems” on page 24 provides details on save stream usage and memory requirements.

3. The storage nodes store the deduplicated backup data on the DD Boost devices.

Note: An **ifgroup** software configuration of 1 GbE NICs on the Data Domain system will allow multiple storage nodes to use the same network identity and this bandwidth aggregation can improve performance for DD Boost devices.

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

The EMC Information Protection Software Compatibility Guide provides information on NetWorker Application modules compatible with Data Domain systems.

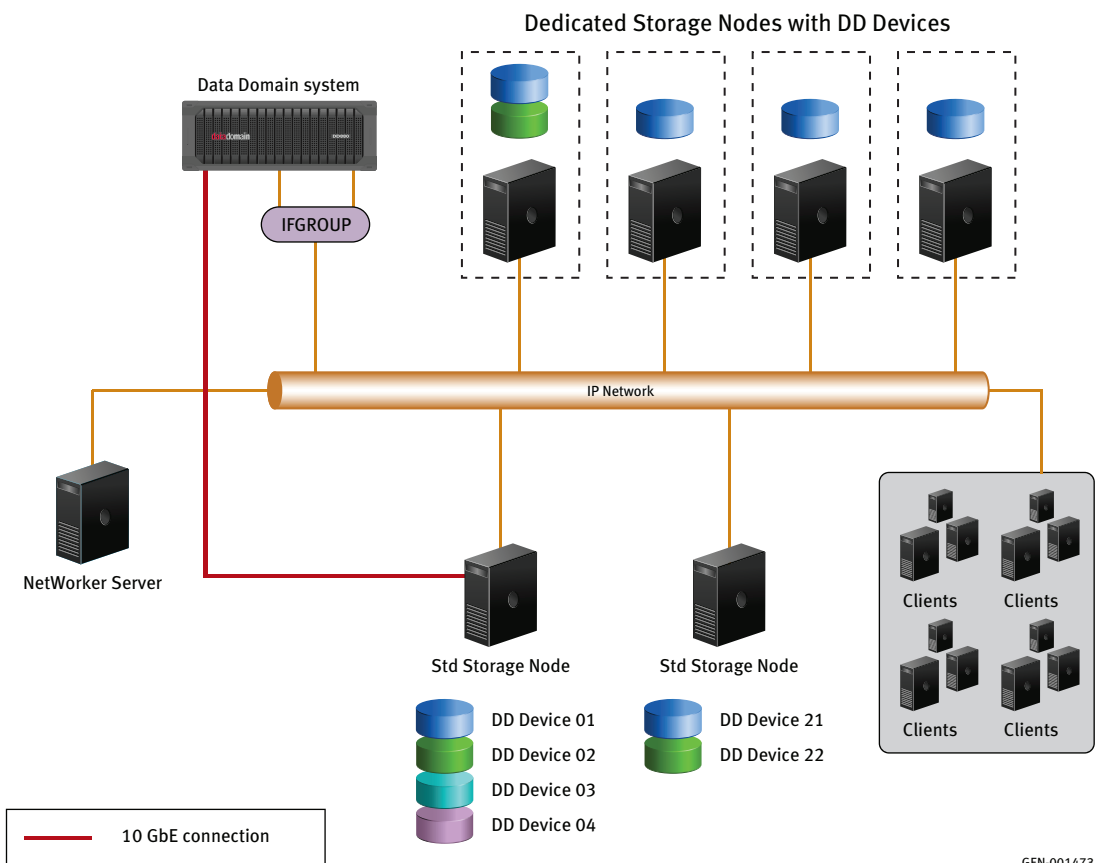


Figure 5 Single datazone with dedicated storage nodes and one high bandwidth link

CHAPTER 3

Data Redirection and Migration

This chapter includes the following sections:

- ◆ [Redirecting backups to the new devices](#) 40
- ◆ [Migrating legacy save sets to the new devices](#) 42
- ◆ [Migrating legacy save sets to DD Boost devices](#) 43
- ◆ [Migration scenarios](#) 44

Redirecting backups to the new devices

Following the upgrade to NetWorker 7.6 SP2 or later, you can redirect the backups of existing NetWorker clients to begin using the new DD Boost devices.

Clients are members of backup groups, which use pools to associate the backups with storage devices. If you redirect the backups at the pool level, as described in this procedure, then minimal effort is required to configure the redirection.

Begin your new redirected backups with a full backup. This avoids a dependency on the last full backup stored with the legacy storage environment and the potential need to restore from two different environments. Begin with a full backup by either:

- ◆ Configure the redirection at a time when the next scheduled backup will be a full backup.
- ◆ Configure the redirection and change the backup schedule to accommodate an initial full backup.

Redirect backups to DD Boost devices

Plan and perform your redirection operation carefully. Read through the following overview procedure completely before you begin.

Note: [“Upgrading from NetWorker 7.6 SP1” on page 97](#) provides information on redirecting backups from existing DD Boost devices.

To redirect existing scheduled backups to begin using storage on DD Boost devices:

1. Ensure that the required network connection, hostname resolutions, and licenses are available and ready to use. The following sections provide details:
 - [“Licensing” on page 20](#)
 - [“Host naming guidelines” on page 29](#)
 - [“Network connectivity and firewall” on page 30](#)
2. Ensure that the Data Domain system is properly configured. [“Configuring Data Domain for NetWorker” on page 50](#) provides details.
3. If you are planning to migrate your existing save sets to the new DD Boost devices, it is best to migrate *before* the redirected backups are scheduled to begin. Migration will “seed” the Data Domain system and help to reduce the bandwidth requirements for future backups. [“Migrating legacy save sets to the new devices” on page 42](#) provides further details.
4. If you are upgrading from a NetWorker 7.6 or earlier server on a Microsoft Windows system, which used an existing CIFS AFTD on a Data Domain system, remove the local administrator account (created to match the Data Domain system user account) from the NetWorker Remote Exec Service:
 - a. Open **Control Panel** > **Administrative Tools** > **Services**.
 - b. Right-click **NetWorker Remote Exec Service** and select **Properties**.
 - c. On the **Log On** tab, select **This Account**, and click **Browse**.

- d. Type the local administrator username that was created to match the Data Domain system user account and click **OK**.
- e. Type the password for the account, and click **OK**.
- f. Right-click the **NetWorker Remote Exec Service** and click **Stop**.

Note: This also stops the NetWorker Backup and Recovery Service and the EMC GST service.

- g. Right-click the **NetWorker Backup and Recovery Service** and click **Start**. This also starts the NetWorker Remote Exec Service.
 - h. If the NetWorker server is also the NMC server, right-click the **EMC GST Service** and click **Start**.
5. Run the NMC Device Configuration Wizard. The following tasks are specific to DD Boost devices:
 - a. Select or create DD Boost devices on the Data Domain system where the backup save sets will be stored.
 - b. Select or create a Pool resource that will target the save sets to the devices.

Note: The wizard enforces that a NetWorker Media Pool created for the Data Domain system uses *only* DD Boost devices. This can be verified in the NMC property windows (**Media Pool, Selection Criteria** tab) for **Target Device > Devices**. The wizard also sets **Target Media > Media type required** to Data Domain.

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

[“Creating a DD Boost device with the wizard” on page 52](#) provides details.

6. In the NMC **NetWorker Administration** window:
 - a. In the **Media** view, open the pool created or modified in [step 5](#).
 - b. In the **Data Source** attribute, ensure that the groups selected for this pool are intended for DD Boost devices.
 - c. Ensure that these backup groups are removed from their legacy pools.

[“Create pools to target DD Boost devices” on page 59](#) provides details.

Note: Existing individual NetWorker Client resources should *not* require any changes to their settings, including the **Data Domain backup** attribute. The redirection is more efficiently configured in the Pool resource.

7. Test the backup environment to ensure that the new configuration operates correctly and existing backups that should be unaffected continue to run as expected. For backups to new devices, the test should also include a restore from those devices.

8. Start your redirection with a full backup to the new devices. This avoids a dependency on the last full backup stored with the legacy storage environment and the potential need to restore from two different environments.
9. 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 84](#) provides details.

Migrating legacy save sets to the new devices

When you redirect client backups to the DD Boost devices, you will be left with existing save sets that are stored on the legacy devices or filesystems. The legacy stored data may either be retained until expiry or migrated to the new devices.

Whether to retain or migrate your legacy data is a judgment that depends on requirements that will differ between sites, clients, and backup types. For example, in some cases, you may choose to retain most of your legacy data and migrate only the backups of sensitive or high-priority clients or backup types. The following points provide some considerations for each of these options:

- ◆ Retain existing save sets on the legacy storage system until they expire:
 - Easiest way to make the transition; no migration is necessary.
 - The legacy storage must be maintained for the life of the legacy data.
 - If data needs to be restored, it would need be restored from either the legacy environment or the new devices, depending on when it was stored.
 - Cloning, storage, and recovery features would be different between the legacy and new data.
- ◆ Migrate existing save sets to the new DD Boost devices:
 - Frees up the legacy storage system for removal or use by other clients.
 - Allows the new devices to be “seeded” with legacy client data. This ensures that subsequent client backups will be deduplicated against the legacy data, which can significantly reduce the bandwidth and time required for the first regular backup window.
 - Offers more flexible storage features with regard to cloning and multiple concurrent storage and recovery operations.
 - Maintains NetWorker browse and retention policies and ensures that all save sets are managed by the NetWorker software.

Migration versus native Data Domain replication

Do not use the native Data Domain replication feature to migrate data from one Data Domain system to another. The NetWorker software will not be able to track, manage, or recover legacy save sets that are replicated in this manner.

An exception is where “seeding” a new system may be used. For example, native Data Domain replication can be performed over a local connection to quickly and easily seed a new target Data Domain system, which would then be physically shipped to a distant location. Although NetWorker cannot immediately recover the seeded data, it will reduce

the otherwise heavy bandwidth required either for a data migration by a NetWorker clone operation, or if no migration, for the initial full backups to the target system. This method can be especially effective if the remote location has limited network bandwidth.

Migration methods

Data migration is essentially a one-time NetWorker clone operation that is tailored to the different device types and time scales. The migration can include all the data or be limited to data selected from a specific time frame or backup type, such as weekly full backups.

The details of the migration procedure depend on the method to be used and the granularity of the data to be migrated:

- ◆ A NetWorker scheduled clone operation. [“Migration scenarios” on page 44](#) provides details.
- ◆ A NetWorker nsrclone script to be run from a command line. The *EMC NetWorker Administration Guide* provides details.
- ◆ A NetWorker staging (data moving) operation. The *EMC NetWorker Administration Guide* provides details.

Migrating legacy save sets to DD Boost devices

After you have chosen the migration scenario, migrate your existing save sets to DD Boost devices. Part of this procedure requires the creation of a special clone pool and a NetWorker scheduled clone operation:

1. Choose your migration scenario. [“Migration scenarios” on page 44](#) provides details.
2. Plan your migration schedule to ensure that sufficient DD Boost devices and bandwidth are available and the impact to the regular backup window is minimized.

Migration should be performed *before* NetWorker client backups are scheduled to begin by using the new devices. Migration will seed the Data Domain system and help to reduce the bandwidth requirements for future backups.

Note: When existing deduplicated VTL or CIFS/NFS AFTD save sets are migrated, the save sets are first reverted to their native non-deduplicated format, which the storage node reads and stores 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 to target DD Boost devices for the migration:
 - In the **Data Source** attribute, select groups for the migration.
Typically, you will migrate the same groups that you selected for the redirection of backups. [“Redirect backups to DD Boost devices” on page 40](#) provides details.
 - In the **Target Devices** attribute, select the DD Boost devices where the migrated data will be stored.

[“Create pools to target DD Boost devices” on page 59](#) provides details.

4. Configure a NetWorker scheduled clone operation with the **Write Clone Data to Pool** attribute selected for your clone pool.

[“Scheduling a clone operation” on page 74](#) provides details on the scheduled clone option.

5. Allow the clone operation to proceed, either according to its schedule or by manual start.

[“Starting a scheduled clone operation on demand” on page 77](#) provides details.

6. After the clone operation completes, verify the cloned data on the target devices.

[Chapter 7, “Monitoring, Reporting, and Troubleshooting,”](#) provides details on the verification of NetWorker operations.

7. Once the cloned save sets have been verified, if desired, you may remove the original save sets.

8. If the original save sets are removed, you may remove old unused devices and pools to ensure that they are not mistakenly used or cause confusion. A pool may *not* be deleted until all volumes that belong to that pool are also deleted or relabelled into other pools.

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.

[Chapter 7, “Monitoring, Reporting, and Troubleshooting,”](#) provides details.

Migration scenarios

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

Some example migration scenarios are as follows:

- ◆ [“Migrating to a new DD system from traditional tape or disk storage” on page 44:](#)
 - [“Migrate to new devices on a different storage node” on page 45](#)
 - [“Migrate to new devices on the same storage node” on page 45](#)
- ◆ [“Migrating to the existing DD system from non-integrated storage” on page 46:](#)
 - [“Migrate to a new devices on a different storage node” on page 47](#)
 - [“Migrate to new devices on the same storage node” on page 48](#)

Migrating to a new DD system from traditional tape or disk storage

In these first two migration scenarios, you have added a new Data Domain system to your existing NetWorker storage environment. You want to migrate and deduplicate your legacy data, which is currently stored on traditional tape or disk, to the new system. The reason for this migration might be that the old storage system is due to be removed or you wish to free up space on the old system for other clients.

How many clients you migrate depends on how much data you want to move and the purpose of the migration, whether for client recovery or simply for seeding for future backups. Because the new Data Domain system contains no data, at least some minimal

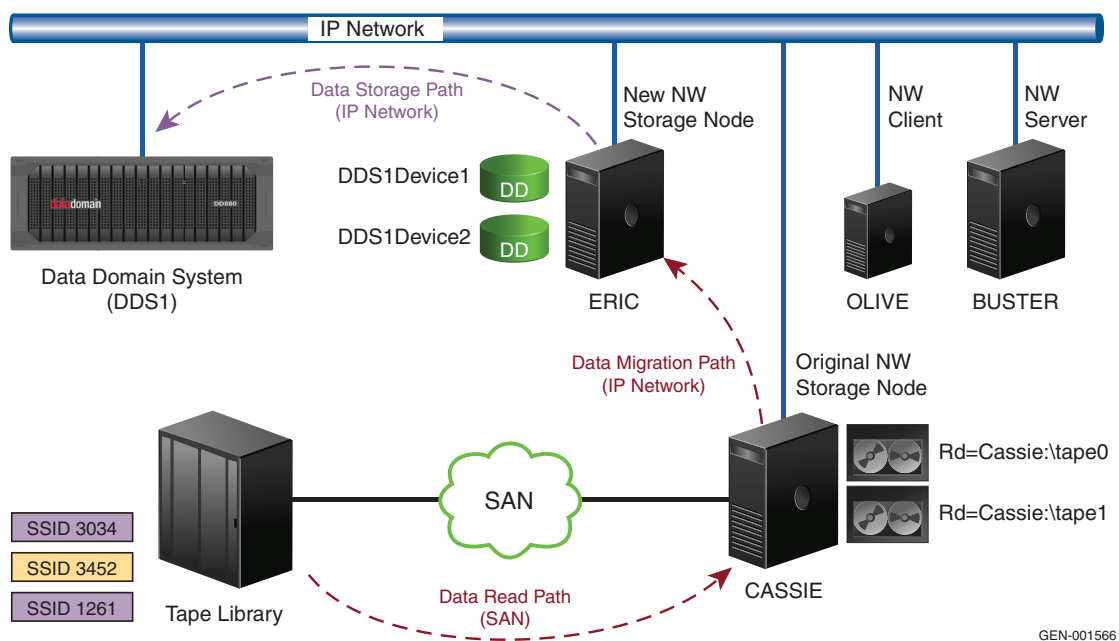
seeding is desirable. If the data for one client is migrated to seed the Data Domain system and the same data is likely to be found on other clients, then migrating the data for the additional clients might have diminishing seed value.

Two migration scenarios of this type are presented. In the first case, the DD Boost devices are created on a new storage node. In the second case, the devices are created on the existing storage node.

Migrate to new devices on a different storage node

Figure 6 on page 45 illustrates a scenario where legacy backups of the client OLIVE, which were stored by the storage node CASSIE on traditional tape or disk, will be migrated to a different storage node ERIC for storage on the Data Domain system.

In this scenario, the IP network is used to transfer the data from the original storage node CASSIE to the new storage node ERIC. The time required for the transfer will depend on the capacity and bandwidth available on the IP network, irrespective of the fact that the tape library is on a SAN. If this same IP network is used for backup and restore operations, then additional effort and time will be required to ensure that the data transfer does not impact the regular backup operations.



GEN-001566

Figure 6 Migration from traditional storage to NW-DD devices on a different storage node

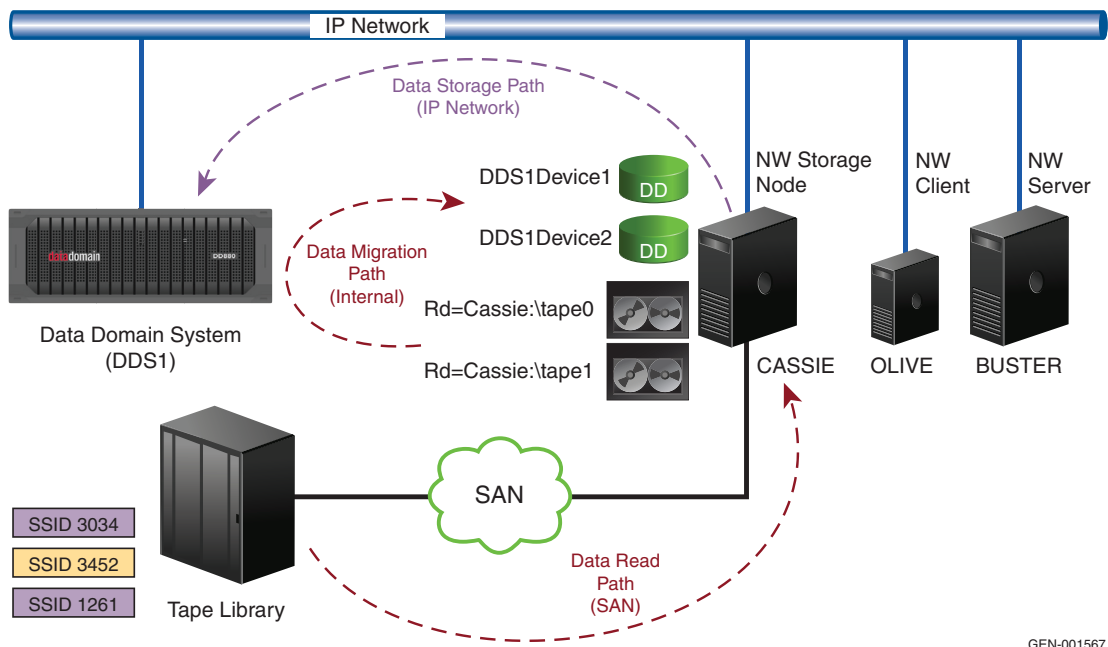
Migrate to new devices on the same storage node

You can eliminate data migration over the IP network between storage nodes by using the same storage node for both the old and new devices. Figure 7 on page 46 illustrates a scenario where DD Boost devices are created on the original storage node CASSIE. During this migration, these new devices read the data stored on traditional tape or disk and send deduplicated data to the Data Domain system for storage.

Although this scenario appears to be an ideal solution because it avoids IP network restrictions, the following factors should be considered:

- ◆ The existing storage node is likely to be older and already at or near capacity. This will limit the number of devices you can add and the amount of data that can be transferred during regular backup operations.
- ◆ The existing storage node is unlikely to have extra network connections available. If the connections are currently used for backup and restore operations, then the bandwidth available for the additional Data Domain storage format will be limited.
- ◆ The network connection is less likely to have 10 GB capacity, which is recommended to maximize throughput from the storage node to Data Domain systems.
- ◆ Although the same storage node will be used for the same backup clients, changes to pools and device allocation are still required. This could add confusion and be more prone to configuration errors.

For smaller sites, it may be a desirable to avoid network restrictions by migrating to new devices on the same storage node. This could also be an option for larger sites where multiple storage nodes are to be reused or reconfigured to share one or more Data Domain system. A storage node could be configured for data migration to seed the Data Domain system as an interim step.



GEN-001567

Figure 7 Migration from traditional storage to NW-DD devices on the same storage node

Migrating to the existing DD system from non-integrated 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, migration for the purpose of seeding is not required. The global deduplication format ensures that the same data is not resent.

The reason for this migration would be that it offers more flexible storage options with regard to cloning and multiple concurrent storage and recovery operations.

Although these migration scenarios use the same Data Domain system, changes to pools and device allocation are still required to redirect backups to the DD Boost devices. A copy or clone operation is required to migrate the data.

When existing deduplicated VTL or CIFS/NFS AFTD save sets are migrated, the save sets are first reverted to their native non-deduplicated format, which the storage node reads and concurrently stores in deduplicated format on the new DD Boost devices. Data that already exists in deduplicated format on the Data Domain system will not be deduplicated again but only the metadata will be unique.

Migrate to a new devices on a different storage node

Figure 8 on page 47 illustrates a scenario where legacy backup data from the client OLIVE was being stored by the storage node CASSIE in VTL format over a SAN connection. This data will be migrated to the new DD Boost devices on the different storage node DOVE for storage on the Data Domain system.

During this migration, the SAN and the IP network are used in two separate stages. First, the original storage node CASSIE reads the non-deduplicated data from 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 the storage metadata, across the IP network to the Data Domain system. The speed of the transfer of non-deduplicated data across the IP network would be the limiting factor.

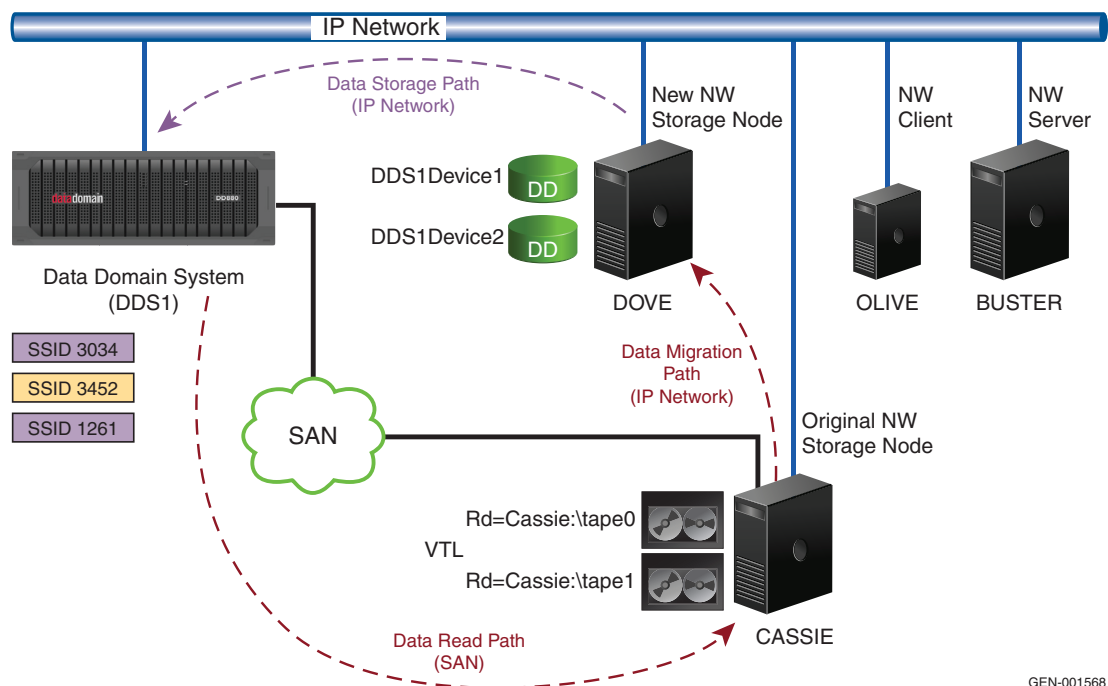


Figure 8 Migration from VTL to NW-DD devices on a different storage node

Migrate to new devices on the same storage node

Figure 9 on page 48 illustrates a scenario where legacy backup data from the client OLIVE will be migrated to new DD Boost devices on the original storage node CASSIE. This storage node is configured for VTL storage on a SAN and also for the new DD Boost devices, which will use the IP network.

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

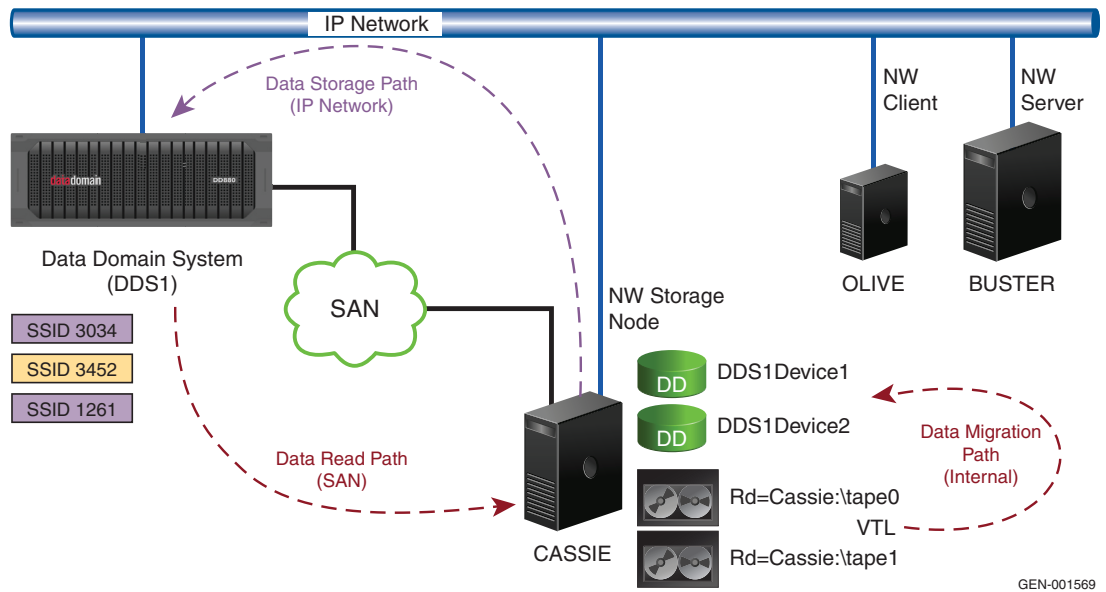


Figure 9 Migration from VTL to NW-DD devices on the same storage node

CHAPTER 4

Software Configuration

This chapter includes the following sections:

- ◆ [Configuring Data Domain for NetWorker](#) 50
- ◆ [Configuring NetWorker for Data Domain](#) 51
- ◆ [Creating a DD Boost device with the wizard](#) 52
- ◆ [Configuring a DD Boost device with the property windows](#) 54
- ◆ [Configuring clients to back up to Data Domain](#) 61
- ◆ [Deactivating and removing a device](#) 66

Configuring Data Domain for NetWorker

Configure the Data Domain system for the NetWorker backup environment as described in the following sections.

Configure the Data Domain system for NetWorker

The following steps use the Data Domain CLI. The Data Domain documentation provides information on using the Data Domain Enterprise Manager GUI instead.

To enable a Data Domain system for storage operations with a NetWorker server:

1. Ensure the required versions and licenses are available and ready to use. [“Network environment” on page 18](#) provides details.
2. On the Data Domain system, log in as an administrative user and open the Data Domain CLI.
3. Verify that the filesystem and NFS services are running:

```
# fileysys status
# nfs status
```

If NFS is *not* enabled, type the following command:

```
# nfs enable
```

Note: The DD Boost functionality requires NFS service to be enabled on the Data Domain system, even if no users or shares are configured. The NetWorker server, storage nodes, or clients are *not* required to be enabled for NFS.

4. If required, confirm the version number of the currently installed Data Domain operating system:

```
# system show version
```

5. Create the backup user and set the DD Boost user name:

```
# user add username
# ddbboost set user-name username
```

Note: Use the same *username* in both of these commands. Only one DD Boost (OST) user may have access to Data Domain at a time.

6. Add the OPENSTORAGE license key supplied by Data Domain:

```
# license add license_key
```

7. (Optional) If clone-controlled replication (“optimized cloning”) will be used, add the REPLICATION license key supplied by Data Domain:

```
# license add license_key
```

8. Apply the modifications by disabling and re-enabling the Data Domain services:

```
# ddbboost disable
# ddbboost enable
```

9. Specify and enable users to monitor backup events captured by SNMP traps and add the host to receive 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.

10. Verify that the **boost**, distributed segment processing (DSP), option is enabled. This is required for deduplication operations on the NetWorker clients or storage nodes:

```
# ddbboost option show
```

If boost is disabled:

```
# ddbboost option set distributed-segment-processing enabled
```

Note: The **boost** option must be enabled. The disabled setting is *not* supported for production use.

11. Use the Device Configuration Wizard to create DD Boost devices and the Data Domain storage unit (SU) folders that contain them.

“[Creating a DD Boost device with the wizard](#)” on page 52 provides details.

Configuring NetWorker for Data Domain

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

- ◆ Each DD Boost device appears as a folder on the Data Domain system. Each device is identified by a unique NetWorker volume label that associates it with a pool.
- ◆ Pools are used to direct backups or clones of backups to specific local or remote devices.
- ◆ NetWorker groups define backup and clone schedules for member clients and are managed by pools. Although not mandatory, it is advised to create groups that are dedicated solely to DD Boost backups.

DD Boost device performance considerations

There is no recommended limit on the number of DD Boost devices that can be created. However, the number of devices you use depends on your backup and recovery requirements.

- ◆ Increasing the number of DD Boost devices can impact Data Domain performance and maintenance. Typically, where multiple concurrent recovery sessions (streams) is not a priority, configure device Target and Max Sessions for multiple concurrent save sessions.
- ◆ Avoid the removal of devices. “[Delete a device](#)” on page 67 provides details.

Methods

The Device Configuration Wizard, which is part of the NMC software, is designed to reduce user errors and should be used to create and configure the DD Boost devices. It also creates and configures the required volume labels and storage pools that are necessary to store client backups on DD Boost devices.

After a device has been created with the wizard, you may modify its configuration, if required, by using the individual NMC property windows:

- ◆ [“Creating a DD Boost device with the wizard” on page 52](#)
- ◆ [“Configuring a DD Boost device with the property windows” on page 54.](#)

Creating a DD Boost device with the wizard

The Device Configuration Wizard is designed to reduce user errors when creating DD Boost devices and is the preferred method.

To create a device with the wizard:

1. Ensure the required versions and licenses are available and ready to use. [“Network environment” on page 18](#) and [“Licensing” on page 20](#) provide details.
2. In the NMC **Enterprise** view, select the NetWorker server name, and double-click the **NetWorker** application to launch it.
3. In the NetWorker Administration application, click **Devices**.
4. In the left panel, right-click **Data Domain Systems** and select **New Device Wizard**.

On each wizard page that appears, specify the options and values required for the backup configuration and click **Next**.

Note: To navigate to *visited* wizard pages to modify them, click the links in the steps panel. The number of steps may vary according to the type of configuration chosen.

5. In **Select the Device Type**, specify **Data Domain** device type.
6. In **Data Domain Preconfiguration Checklist**, review the requirements.
7. In **Specify the Data Domain Configuration Options**:
 - a. In the **Data Domain System Name**, choose the appropriate option:
 - For **Use an existing Data Domain System**, select an existing system.
 - For **Create a New Data Domain System**, specify the FQDN, or IP address, of the Data Domain system.
 - b. In **Data Domain OST Username** and **Data Domain OST Password**, specify the DD Boost username and password, respectively.

Note: Only one DD Boost (OST) username and password may be defined. All NetWorker storage nodes and servers that access the Data Domain system must use the same username and password. [“Configure the Data Domain system for NetWorker” on page 50](#) provides details on setting the DD Boost username.

- c. In **Configuration Type**, select **Browse and Select**.
8. In **Select Folders to use as Devices**, create a new DD Boost device:

Note: The wizard does *not* show existing devices that were previously created by NetWorker 7.6 SP1. [Appendix A, “Upgrading from NetWorker 7.6 SP1,”](#) provides details.

- a. Select a top-level folder (the Data Domain system) and click **New Folder**.

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

Note: For NetWorker 7.6 SP2 and later, the folder 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 this SU folder after the short hostname of the NetWorker server and places the devices you create within this SU folder.

- b. Type a name for the new device.

To list the details of any device, select its checkbox. The table displays the full NetWorker device name and the storage path name.

Note: For NetWorker 7.6 SP2 and later, the device name refers to the sub folder that is created within the SU. There should be *no* folders deeper in the folder path than these device folders. Do *not* create device folders within device folders.

- c. (Optional) To rename a DD Boost device as it will appear in NMC, select the device in the table and type a new name in the **NetWorker Device Name** attribute. No special characters are permitted except dot (.) and underscore (_). The **Storage Path** attribute remains unchanged.

Note: Implicit in the path to the SU folder 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). This folder structure is used by all NetWorker server hosts and should *not* be modified. The **Review Configurations Settings** wizard page shows this location.

Note: Although NFS service data is stored in the /backup directory, clients that are configured for NFS access also have the ability to view, change, and delete the /data/col1 directory that contain the DD Boost devices. If NFS devices are used, alternative configurations may be advised to avoid the risk of potential interferences.

9. In **Configure Pool Information**:

- Ensure **Configure Media Pools for Devices** is selected.
- Specify the pool that will target clients to the devices, either **Backup** for backups or **Backup Clone** for cloning or staging operations, or create a new pool.

Note: If you create a new pool, do *not* select an existing pool.

- Ensure **Label and Mount device after creation** is selected.

10. In **Select Storage Nodes**, select or create the storage node that handles the devices.
11. In **SNMP Monitoring Options**, type the Data Domain SNMP community string and specify the events to be monitored.

If the community is not known, unselect **Gather Usage Information**.

SNMP monitoring allows NMC to display the Data Domain system status and monitor backup and recovery events. It also provides a launcher link for the Data Domain interface.

12. In **Review Configuration Settings**, review the information and click **Configure**.

Note: For NetWorker 7.6 SP2 and later, what appears to be listed as the Storage Unit name is actually the path for the device folder in the format: *SU/device_name*, where *SU* is named after the short hostname of the NetWorker server.

The DD Boost device will be configured, mounted, and labeled for the specified pool.

13. The final **Device Configuration Results** page is informational only. Click **Finish** to exit the wizard.
14. In the NMC Devices view, verify that the device is labelled and mounted, ready for use. This view should list the device with its appropriate Volume Name.

Configuring a DD Boost device with the property windows

The Device Configuration Wizard is the recommended method to create DD Boost devices. You may modify the devices and perform other tasks by following the manual procedures in the following sections:

- ◆ [“Add a host Data Domain system to NMC Enterprise view” on page 54](#)
- ◆ [“Configure a DD Boost device” on page 55](#)
- ◆ [“Create a volume label template for DD Boost devices” on page 58](#)
- ◆ [“Create pools to target DD Boost devices” on page 59](#)
- ◆ [“Label and mount devices on the storage node” on page 61](#)

Add a host Data Domain system to NMC Enterprise view

Normally, a Data Domain host is added to the NMC Enterprise view by the Device Configuration Wizard. As an alternative, a host may be added by the Add New Host Wizard.

Data Domain systems that are used for storage with NetWorker software should be listed as hosts in the NMC Enterprise view. This allows NMC to display the Data Domain system status and monitor events during backup and recovery operations. It also provides you with a live link that launches the Data Domain GUI.

To add a host Data Domain system to the NMC Enterprise view:

1. From the **File** menu, select **New > Host** to run the **Add New Host** wizard.
2. Follow the wizard screens:
 - Type the Data Domain Host Name
 - Select **Data Domain**

- Select **Capture Events**
- Type the SNMP community string, where NMC will retrieve Data Domain status information. The default is the value set on the Data Domain system with the **snmp add ro-community *community_name*** command. [“Configure the Data Domain system for NetWorker” on page 50](#) provides details.
- Type a value for the SNMP Process Port. The default is the value set on the Data Domain system with the **snmp add trap-host *hostname[:port]*** command. [“Configure the Data Domain system for NetWorker” on page 50](#) provides details. This setting should agree with the firewall setting on the Data Domain system. [“Firewall requirements” on page 31](#) provides details.
- Select the **SNMP Traps** that you want to monitor.

[“Configure SNMP for Data Domain” on page 86](#) provides details on monitoring and alerts.

Configure a DD Boost device

Normally, DD Boost devices are created by the Device Configuration Wizard. You may use this manual procedure to *modify* an existing device. A device appears as folder on the Data Domain system and each device is associated with a single NetWorker volume by the label and mount procedure.

Note: If you manually create a device with this procedure, NMC will list the device but no corresponding device folder will be created on the Data Domain system. If you try to label and mount such a device, an error appears. [“Configure the Data Domain system for NetWorker” on page 50](#) provides an optional step to manually create a storage unit on the Data Domain system.

To modify a DD Boost device:

1. Ensure that the required versions and licenses are installed.
[“Network environment” on page 18](#) and [“Licensing” on page 20](#) provide details.
2. In the NMC window for your NetWorker server, click the **Devices** view.
3. In the folder tree, expand **Data Domain Systems** and select the Data Domain system where your save sets are stored.
4. In the right panel, right-click the name of the device to modify and select **Properties**.
5. On the **General** tab, in the **Identity** area, type the DD Boost device name and access information:

Note: Multiple devices can share a single volume. [“Configure multiple devices for a single volume” on page 58](#) provides details.

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

For example:

dd_1

The following values are used in [Figure 10 on page 56](#):

- 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 [Figure 10 on page 56](#), it is a “remote device” and this **Name** attribute must be specified in the following format:

rd=remote_storagenode_hostname:device_name

For example:

rd=dzone1_sn2:dd_1

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

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, as in [Figure 10 on page 56](#):

ddr1:/dzone1/dd_1

There is no recommended limit to how many device folders you can create, but this attribute accepts one device folder only. Do *not* create folders within a device folder.

Note: Implicit in this path are the hidden mount point folders, /backup/ost (DD OS 4.9) or /data/col1 (DD OS 5.0 and later). This folder structure is used by all NetWorker server hosts and should *not* be modified.

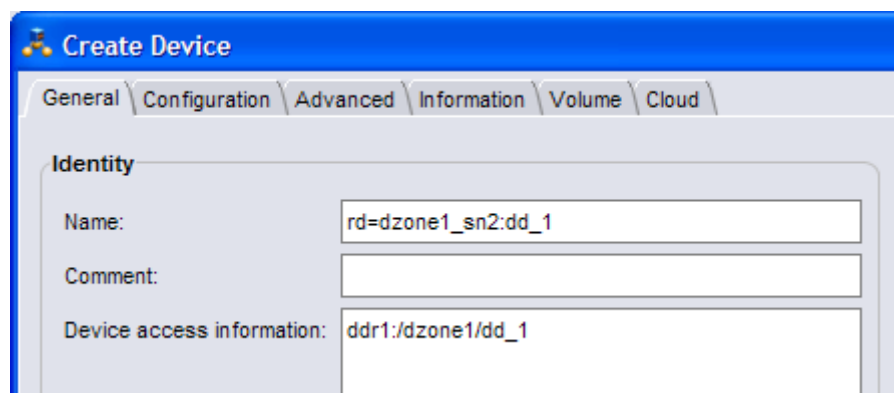


Figure 10 DD Boost device name and access information.

- c. In the **Media type** attribute, select **Data Domain** from the list.
6. Ensure that **Auto media management** is *not* selected. It does not apply to DD Boost devices.

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

- In the **Target sessions** field, specify the number of sessions that a **nsrmmd** process on the device will handle before for another device on the host will take the additional sessions. Use this setting to balance the sessions among **nsrmmd** processes.

If another device is not available, then another **nsrmmd** process on the same device will take the additional sessions.

Typically, set this attribute to a low value. The default value is 1. The maximum value is 60.

[“Configure multiple devices for a single volume” on page 58](#) provides details on shared volumes.

- In the **Max sessions** field, specify the maximum number of sessions that the device can handle. If no additional devices are available on the host, then another available Data Domain system takes the additional sessions, or retries are attempted until sessions become available.

The default value is 6, which is recommended. The maximum value is 60.

Note: The Max sessions setting does not apply to concurrent recover 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 **Dynamic nsrmmms** is enabled on the storage node, this value is automatically adjusted by the formula $\text{max}/\text{target} + 4$, with the default value being 14. Otherwise, the default value is 4.

To modify this value, first adjust the sessions attributes, apply and monitor the effects, and then update the **Max nsrmmd count** value.

Note: At least one **nsrmmd** process will always be reserved for restore or clone operations.

8. In the **Remote user** and **Password** attributes, type the DD Boost username and password, respectively.

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

9. Click **OK** to save the device settings.

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

10. Ensure that the device is associated with a NetWorker storage volume before attempting to use it. Otherwise, an error appears. [“Label and mount devices on the storage node” on page 61](#) provides the procedure.

Configure multiple devices for a single volume

In some environments, multiple devices sharing a single NetWorker storage volume can result in performance gains. For example, a read or write request can be sent to the storage node that is closest to the requestor. Consideration should be given to use cases because in some environments concurrent read or write operations to a volume from many storage nodes could result in disk thrashing that impairs performance.

Multiple devices can be created on separate storage nodes or on the same storage node. Each device must be created separately, must have a different name, and must correctly specify the path to the storage volume location.

For example, for NetWorker 7.6, if you create three devices, one on the NetWorker server host named dzone1 that uses the server's local storage node and two remote devices (rd) on remote storage nodes, the Name attributes for the three devices, each created separately, might be specified by different aliases as follows:

```
dd_1a
rd=dzone1_sn2:dd_1b
rd=dzone1_sn3:dd_1c
```

The Device access information attribute for each of these aliases would specify a single directory as a valid complete path. For example, if a directory is named dd_1 on the storage host named ddr1:

- ◆ If the storage node uses an automounter, you might specify the following path:
/net/ddr1/dzone1/dd_1
- ◆ If the storage node uses an explicit system mountpoint, you might specify one of the following paths:
/mnt/ddr1/dzone1/dd_1
/mnt/dzone1/dd_1

Create a volume label template for DD Boost devices

Normally, the Device Configuration Wizard automatically creates a label template for the volumes that a new device will use. This procedure describes the alternative manual method that uses the NMC property windows.

Note: The Device Wizard automatically creates a label template for the volumes and this procedure does not apply if the wizard is used.

Each DD Boost device must be mounted with a single NetWorker storage volume, which associates it with a specific **Volume Name** and membership within a pool. A label template must be created to provide 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. 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 volume label:

1. From the **NetWorker Administration** window, click **Media**.
2. In the browser tree, select **Label Templates**, and from the **File** menu, click **New** to open the **Create Label Template** window.
3. In the **Name** and **Comment** attributes, type a name and description for the label template. The label will associate a storage pool to a device.
4. In the **Fields** attribute, type components of the label. Place each label component on a separate line. The template must include at least one volume number range component. The label template will be applied 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** attribute, 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**.

Create pools to target DD Boost devices

Normally, the media pool is automatically created by the Device Configuration Wizard. 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 in clone operations, to the storage devices that are members of the pool.

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

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

To manually create a pool for Data Domain backups by using the property windows:

1. Ensure that the devices that will target the pool have already been created. If you do *not* use the Device Configuration Wizard to create the devices, [“Configure the Data Domain system for NetWorker” on page 50](#) provides details.

2. Ensure a label template has been created for the pool. [“Create a volume label template for DD Boost devices” on page 58.](#)
3. From the **NetWorker Administration** window, click **Media**.
4. In the browser tree, 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** attribute, type a name for each pool. Create names that clearly indicate whether the pool is for Data Domain backup or clone operations.

For example:

DDsite1
DDCLsite2

For example, 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 the backup will be stored. These conventions make the name easier to use for scripting and reporting.

[“Host naming guidelines” on page 29](#) provides guidelines for creating names.

6. In the optional **Comment** attribute, 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**. [Chapter 5, “Cloning and Replicating Data,”](#) provides details on clones.

Note: **Pool type** cannot be modified after it is created.

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

The pool will later be applied to DD Boost devices. [“Label and mount devices on the storage node” on page 61](#) provides details.

10. In the **Data Source** attribute, select the groups (and the implicit backup schedules) that the clients in this pool may belong to.

An alternative way to configure a client to use a pool is to specify the **Target Pool** attribute of the client resource. [“Configure a backup client with the client wizard” on page 62](#) provides details.

Note: A client may belong to multiple groups through multiple client definitions.

11. Click the **Selection Criteria** tab.
12. Under **Target Devices**, select all of the DD Boost devices that this pool may use for storage. The pool may store data on any of these devices:
 - Select *only* DD Boost devices for the pool. Do *not* mix with other types of storage devices. The pool **Selection Criteria** may also include client attributes that need to be removed or updated.

Note: If you modify an existing pool in this step, ensure that the pool excludes all devices that are *not* DD Boost devices.

- Do *not* select devices on more than one Data Domain system.

Otherwise, backups from a single NetWorker client may be targeted to any of these Data Domain system. This will impair the backup window and deduplication ratio.

13. Under **Media type required**, if the pool is intended for Data Domain backup, set this attribute to **Data Domain**. This ensures that only Data Domain devices can be associated with this pool.

Note: As a best practice, do *not* include different media types in a single pool. [“Backup fails for older NetWorker application modules” on page 93](#) provides further details.

14. Click **OK**.

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

Label and mount devices on the storage node

Normally, devices are labeled and mounted by the Device Configuration Wizard. The following procedure describes the alternative manual method that uses the NMC property windows.

A DD Boost device must be identified with a volume label (Volume Name) and its associated pool, and must be mounted for use by a NetWorker storage node.

To label and mount a device:

1. From the **NetWorker Administration** window, click **Devices**.
2. In the browser tree, select the Data Domain system.
3. In the right panel, right-click the device to be labeled and select **Label**.
4. In the **Label** window, **Pools** list box, select a pool to associate with the device.

A label for the selected pool appears in the **Volume Label** attribute. This will become the **Volume name** that is applied to the device.

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

The NMC list shows the device with its associated **Volume name**.

The *EMC NetWorker Administration Guide* provides details.

Configuring clients to back up to Data Domain

The NMC software provides different methods to create NetWorker clients and their schedules. The wizard method is recommended, but you may use the individual property windows method as described in the following sections.

Configure a backup client with the client wizard

The NetWorker Client Configuration Wizard is designed to reduce user errors and is the preferred method for creating and configuring NetWorker backup clients and their backup schedules.

Note: If you want to redirect existing client backups to new DD Boost devices, [“Redirecting backups to the new devices” on page 40](#) provides details.

To create or modify a client and its scheduled backup with the wizard:

1. If the client is to use Client Direct backup, which is the default configuration, ensure the client has access to the same network interface that is used and referenced by the target DD Boost device.
2. In the NMC **Enterprise** view, select the NetWorker server name, and double-click the **NetWorker** managed application to launch it.
3. In the **Configuration** view, under the NetWorker server name, right-click **Clients** and select **Client Backup Configuration > New**.
4. Follow the instructions on the wizard pages:
 - Specify 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 will back up to Data Domain devices if the pool used also contains other types of devices, although pools with mixed devices is *not* recommended.
 - Set **Target Pool** to a pool that is associated with Data Domain devices.

An alternative way to configure a client to use a pool is to specify the client or its group in the **Data Source** attribute of the Pool resource. [“Create pools to target DD Boost devices” on page 59](#) provides details.

Note: Current versions of NetWorker application modules support backup to DD Boost devices. Some of the earlier versions of modules do not support the client attributes for **Data Domain backup** and **Pool** and these attributes should *not* be set. [“Backup fails for older NetWorker application modules” on page 93](#) provides details.

- Select Files to Back Up.
- Select the Client Properties.
- Select the Backup Group.
- Specify the Storage Node Options.

[“Configure a backup client with the property windows” on page 63](#) provides details on settings for backups to DD Boost devices.

Configure a backup client with the property windows

As an alternative to the client wizard method, new or existing NetWorker backup clients may be configured for the DD Boost backup environment by using the individual NMC client configuration windows.

To configure a NetWorker client for scheduled backups to a DD Boost device:

1. If the client is to use Client Direct backup, which is the default configuration, ensure the client has access to the same network interface that is used and referenced by the target DD Boost device.
2. In the NMC **Enterprise** view, select the NetWorker server name and double-click the NetWorker application to launch it.
3. In the window for your NetWorker server, click **Configuration**.
4. In the browser tree, select **Clients**:
 - To create a new Client resource, select the **Clients** icon and from the **File** menu, click **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**.
5. To create a new Client resource, perform the following sub-steps. Otherwise skip this step:

Note: The following attributes do *not* normally need to be modified for clients whose backups have been redirected to DD Boost devices. [“Redirecting backups to the new devices” on page 40](#) provides details:

- a. On the **General** tab, in the **Name** attribute, type the hostname for the client and, optionally, type a comment in the **Comment** attribute.
- b. Select values for **Browse Policy** and **Retention Policy**:
 - The browse policy determines how long the details of individual backed-up files are maintained in a browsable index for quick recovery through the GUI or command line.
 - The retention policy determines how long backed-up data is protected and available for recovery even though the browse policy has lapsed. Recovery might require rebuilding an index.
- c. Select the **Scheduled Backups** checkbox.
- d. In the **save sets** attribute, click the **Browse** button to open the **Choose Directory** window. Browse to and select all the directories or individual files to be backed up. When finished selecting, click **OK**.

Each item should appear on a separate line. For example:

```
D:\accounting
sales
```

To back up *all* client data, type **All** in the **save sets** attribute.

Note: For Microsoft Windows systems, the SYSTEM or Volume Shadow Copy Service (VSS) SYSTEM save sets should be backed up on a regular basis for disaster recovery of the client system.

The *EMC NetWorker Administration Guide* provides details for this step.

6. On the **General** tab, in the **Backup** area, configure the following:
 - a. Select the **Client Direct** attribute to enable backup data from this client to bypass the NetWorker storage node and send deduplicated backup data directly to the Data Domain system. The following requirements must be met:
 - The client must have network access to the Data Domain system.
 - The **Checkpoint restart** attribute must *not* be selected. If selected, backups will revert to traditional storage node backups.
 - The Data Domain system and the client host must be enabled with DD Boost interface. [“Configuring Data Domain for NetWorker” on page 50](#) provides details.
 - The NetWorker Device resource for the Data Domain system must be configured with a **Remote User** attribute that specifies a DD Boost (OST) username. [“Configure a DD Boost device” on page 55](#) provides details.
 - **Data Domain backup** must be selected on the Client’s **Apps and Modules** tab, see the appropriate step later in this procedure.

If this attribute is *not* selected, or if the Data Domain system is unable to receive deduplication backups directly from the client, then the Data Domain backup will be deduplicated on the storage node before it is sent to the Data Domain system for storage.

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. In the **Group** attribute, select a backup group. A backup group consists of clients that are backed up according to the schedule for the group.

This attribute might show groups pre-selected based on the pool selected. It is recommended to *not* separate groups from their associated pools at the client level.

Note: Groups that are specifically devoted to deduplication backups help you monitor deduplication operations and configure backup schedules for either deduplication or traditional storage.

- c. In the **Pool** attribute, select a pool that targets the DD Boost devices you want to use. The pool selected in this attribute overrides any other pool that might be configured for the client or its save sets.

An alternative way to configure a client to use a pool is to specify the client or its group in the **Data Source** attribute of the pool resource. [“Create pools to target DD Boost devices” on page 59](#) provides details.

- d. (Optional) A Data Domain backup applies only to the save sets (the data portion of backups), and does *not* back up the NetWorker server's bootstrap and index files to a DD Boost device. Normally, the bootstrap and index files are backed up to a traditional disk or tape device that is locally attached to the NetWorker server.

If you want to back up the bootstrap and index files to a DD Boost device, create a Pool resource specifically for them. In the pool's **save sets** attribute, specify the following:

```
bootstrap*
index*
```

- e. In the **Schedule** attribute, select a backup schedule.

The schedule selected in this attribute overrides any other schedules that might be configured for the client or its save sets.

7. Select the **Apps & Modules** tab.
8. In the **Deduplication** area, select **Data Domain backup**. This ensures that the client data will be backed up *only* to DD Boost devices, even if the pool selected for the backups contains a mix of other device types. Pools with a mix of device types are *not* recommended.

Note: Current versions of NetWorker application modules support backup to DD Boost devices. Some of the earlier versions of modules do not support the client attributes for **Data Domain backup** and **Pool** and these attributes should *not* be set. [“Backup fails for older NetWorker application modules” on page 93](#) provides details.

9. For NDMP clients that back up to DD Boost devices, attributes in the **Backup** area are required.

Typically, the NDMP protocol is used to back up to tape devices. The following settings specify deduplication backups to DD Boost devices instead:

- a. Select the **NDMP** option.
- b. In the **Backup Command** attribute, type the following:

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

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

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

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

- c. On the **General** tab, ensure that the client is a member of a pool that targets the DD Boost device you want to use.
 - d. On the **Apps & Modules** tab, ensure that **Data Domain backup** is selected.
10. Click the **Globals (2 of 2)** tab.
 11. In the **Configuration** area, configure the following:
 - a. In the **Storage nodes** attribute, specify storage nodes that the client will back up to, especially if this client backs up to a remote storage node.

- b. In the **Clone storage nodes** attribute, if this client is a source storage node for clone operations, specify the target storage node where the clones will be stored.
12. When you have completed the client configuration, click **OK**.

The window for your NetWorker server should list the client with a checkmark in the **Scheduled backup** column to indicate that scheduled backup is enabled.

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

Deactivating and removing a device

To deactivate a DD Boost device so it does not interfere with normal backup operations, use one of the following options:

- ◆ [“Convert a device to read-only” on page 66](#)
- ◆ [“Disable a device” on page 66](#)
- ◆ [“Delete a device” on page 67](#)

Convert a device to read-only

Conversion of a device to read-only prevents the use of the device for backup operations. The device can still be used for read operations, such as restore and clone.

To convert a device to read-only:

1. In the NMC window for your NetWorker server, click the **Devices** view and select the Devices folder in the navigation tree.
2. In the **Devices** table, right-click the device to be converted to read-only and select **Unmount**.
3. Right-click this unmounted device and select **Properties**.
4. In the Device Properties window, select **Read only** and click **OK**.
5. Right-click the device and select **Mount**.

Disable a device

Disabling a device prevents further operation of the device. The device may be re-enabled to restore old data, which is retained but not active.

To disable a device:

1. In the NMC window for your NetWorker server, click the **Devices** view and select the Devices folder in the navigation tree.
2. In the **Devices** table, right-click the device to be disabled and select **Unmount**.
3. Right-click this unmounted device and select **Enable/Disable** to disable.
4. Inspect the **Enabled** column of the table to verify that the device is disabled.

Delete a device

The procedure for deleting a device includes an option for also erasing the volume (access path) that stores the device's data. The volume can be erased only if no other device in the system shares the volume.

To delete an AFTD:

1. In the NetWorker server **Device** view, click **Devices** in the navigation tree.
2. In the **Devices** table, right-click the device to be removed and select **Delete**.

A confirmation window appears.

3. In the confirmation window:
 - To delete the device from the NetWorker configuration only, without erasing the device's data, click **Yes**.
 - To delete the device and erase the device's data and volume access path, select the **Permanently erase all data and remove media and index information for any selected AFTDs or Data Domain devices** option, and click **Yes**.

Note: If the volume that you want to erase is shared by another device, then an error message displays the name of the other device. You must delete all other devices that share the volume until the last one remaining before you can erase the volume.

4. If the device is mounted or the device is a member of a pool, then a second confirmation window displays the details of the device and pool. To confirm the device unmount, the removal of the device from the pool, and the deletion of the device, click **Yes**.

The Data Domain documentation provides further details.

CHAPTER 5

Cloning and Replicating Data

This chapter includes the following sections:

- ◆ About cloning and replicating with Data Domain 70
- ◆ Clone formats 71
- ◆ Clone methods..... 72
- ◆ Clone requirements..... 73
- ◆ Scheduling a clone operation..... 74
- ◆ Cloning by pools 76
- ◆ Starting a scheduled clone operation on demand..... 77

About cloning and replicating with Data Domain

For added data protection, save sets that have been successfully written to a DD Boost device can be copied to a different location with the NetWorker clone feature. A clone is a complete and independent copy of the data, which can be used for data recovery or to create further clones. Single save sets or the entire volume of a DD Boost device may be cloned. A NetWorker clone retains the original NetWorker browse and retention policies by default, but these can be changed to allow the clone copy to have different policies.

Datzone requirement

A clone of stored data may *not* be created in a different NetWorker datzone. For the NetWorker server to manage and monitor clone operations, the storage nodes at both the source and target locations must be clients of the same NetWorker server. The NetWorker server maintains browse and retention policies for all cloned copies and can monitor and report on their storage operations.

Clone pool

In order for the NetWorker software to copy save sets from one device to another, a special pool must be created. This pool, known as a “clone pool,” is assigned to devices on the target system. It must be properly configured and available for use. [“Create pools to target DD Boost devices” on page 59](#) provides details. The *EMC NetWorker Administration Guide* provides additional details.

Data Domain native replication support

As a recommendation, replication of data on DD Boost devices should use NetWorker clone-controlled replication. This allows the data to be tracked and maintained by the NetWorker software. [“Clone-controlled replication format” on page 71](#) provides details.

The replication operations that are native to Data Domain systems are typically used to copy deduplicated data stored in CIFS, NFS, or VTL formats from one Data Domain system to another for disaster recovery purposes. These native operations are independent of NMC and NetWorker. They provide copies of data without NetWorker knowledge or control.

Note: If Data Domain replication is used for other (non-DD Boost directories) on the same system, ensure that there is adequate network load and bandwidth capacity such that NetWorker clone-controlled replication operations will not be impaired.

Native Data Domain replication support for DD Boost devices is as follows:

- ◆ Native Data Domain directory replication is *not* supported for DD Boost devices.
- ◆ Native Data Domain MTree replication is *not* supported for DD Boost devices.
- ◆ Data Domain collection replication (the replication of the entire stored contents of a Data Domain system) is supported and will replicate the DD Boost devices and stored data onto a target Data Domain system. However, the target Data Domain system cannot be the same system that is used for other replication operations such as NetWorker clone-controlled replication.

NOTICE

If collection replication is performed, the Data Domain devices on the target system become read-only and NMC and NetWorker will not be aware of any DD Boost devices on the target system. NetWorker awareness and recovery from the target Data Domain system would require further procedures, testing, and qualifications to be performed. EMC Professional services are strongly recommended.

Clone reports

NetWorker Data Domain clone operation reports can be viewed in the NMC Reports view.

[“Backup, recovery, and clone reports” on page 88](#) provides details.

Clone formats

Copies of data stored on a DD Boost device may be created by the NetWorker clone process. The type of clone produced depends on the type of media on which the clone copy will be stored, whether a clone-controlled replication or a regular clone.

Clone-controlled replication format

Data that NetWorker clones from a source DD Boost device to a target DD Boost device, typically at a geographically distant location, retains its deduplication format and NetWorker policies. This is called clone-controlled replication (CCR), or optimized clone.

Note: For CCR operations on the *same* Data Domain system, it is recommended to replicate between two different SUs (MTrees) so that different retention policies may be applied and the data can be managed efficiently.

Clones created by CCR may be used for data recovery or to create additional copies. For example, additional clones may be created on multiple Data Domain systems with minimal impact on production or primary backup and recovery operations.

CCR is a fast process that uses low bandwidth and low storage capacity. During the cloning process, the storage node for the target DD Boost device reviews the incoming clone for data that is already stored on the target device and transfers only unique data and reference pointers to be stored on the device.

NetWorker control ensures that a CCR will *not* begin until after the related backup group at the source has finished its backup.

Regular clone format

When data on the DD Boost device is cloned to a traditional disk or tape, it is reverted to its native non-deduplicated format, known as “regular clone” format. Regular clone format is necessary for the data on traditional disk or tape storage to be completely recoverable, for example for disaster recovery, without the need of a Data Domain system.

Clone methods

NetWorker clone operations may be configured by several methods, which are suitable to different environment and storage needs. In some cases, it may be necessary to use multiple or mixed approaches to achieve the desired control and flexibility.

As choices, you can configure clone operations to be run by automatic start (auto-clone), by a schedule, or by a customized script.

Immediate clone

Also known as “auto-clone.” Save sets can be automatically cloned when the backup group that contains them is completed. This clone method is suitable for smaller environments, or a small number of clients, where the clone operations need to be completed quickly and immediately within the backup window:

1. In the NetWorker **Administration** window, select **Configuration**.
2. Create a **Group** resource and specify the **Clones** option and **Clone pool**.

“[Create pools to target DD Boost devices](#)” on page 59 provides details on creating a clone pool.

Scheduled clone

NetWorker scheduled clone operations can be configured and run in NMC according to a schedule for predetermined clients, pools, save sets, and devices. This method is suitable for environments where copies of save sets need to be regularly provided, typically as part of a well-defined maintenance cloning window, which runs independently of the main backup operation.

“[Scheduling a clone operation](#)” on page 74 provides details.

nsrclone command

A NetWorker **nsrclone** script can be created and used to run clone operations and be launched either manually or as a scheduled task run from the operating system or an external scheduler. This method is typically used in larger environments where flexibility and conditional controls are required.

One example would be to start Clone Job B, which clones to tape storage, if and only if Clone Job A successfully completes. Another example would be to select multiple specific save sets to be cloned to specific devices or storage nodes.

Scripted solutions require additional knowledge and have external dependencies, such as operating system, scripting tool, and so forth.

The *EMC NetWorker Administration Guide* and the *EMC NetWorker Command Reference Guide* provides details. EMC Professional Services are also available.

Clone requirements

To clone data from one Data Domain device to another by NetWorker clone-controlled replication (optimized cloning), ensure that the following requirements are met. These requirements assume the previous creation of a clone target pool, for example, named *newclonepool*:

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.

[“Configure the Data Domain system for NetWorker” on page 50](#) provides details.

3. Ensure that the Client resource for the NetWorker server and both storage nodes specify, in their **Aliases** attribute (Globals 1 of 2 tab), all of their names in use (fully qualified name, short name, aliases, and IP address).

Note: If a **nsrclone** command or script is used to perform an optimized clone from a host that is *not* the NetWorker server, then this command must specify the NetWorker server by its primary hostname (as listed 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.

Note: The Default Clone pool does *not* allow any modification and **Media type required** cannot be set in that pool. The use of the default clone pool is *not* recommended.

[“Create pools to target DD Boost devices” on page 59](#) provides details on creating pools.

5. Ensure that the Client resource for the source storage node specifies, in its **Clone storage node** attribute, the target storage node hostname. This setting is not required if the target storage node is on the NetWorker server.

If the **Clone storage node** attribute is *not* specified, then the NetWorker server becomes the storage node for the clone operation.

If the **Clone storage node** attribute lists a storage node for a non-Data Domain volume, and **Media type required** is *not* set to **Data Domain** in the target clone pool, then only regular clones may be stored on those volumes.

6. Ensure that the source DD Boost device is mounted and available on the source storage node.

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

7. Ensure that the target DD Boost device is labeled 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.
8. 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 scheduled clone operations, in the **Write clone data to pool** attribute of the Clone resource, select *newclonepool*.
 - For immediate clone operations for a group, in the **Clone pool** attribute of the Group resource, select *newclonepool*.
 - For clones of entire volumes, [“Cloning by pools” on page 76](#) provides details.
9. Ensure that Data Domain encryption is either enabled or disabled on *both* the source and target Data Domain devices and that there is sufficient bandwidth available for the clone operation.

Scheduling a clone operation

To configure NMC for a clone operation to be performed at a scheduled time, for either a clone-controlled replication (also known as optimized clone or DD format clone) or a regular (non-deduplicated) clone:

1. Ensure the requirements for cloning are met. [“Clone requirements” on page 73](#) provides details.
2. From the **NetWorker Administration** window, click **Configuration**.
3. In the browser tree, click **Clones**, and from the **File** menu, select **New** to open the **Create Clone** window. [Figure 11 on page 75](#) shows the General settings.
4. In the **Name** attribute, create a unique name to identify your scheduled Clone resource. Type any additional information in the **Comment** attribute.
5. Set the **Browse Policy** and **Retention Policy** attributes if you want the clones to use different policies than the original backup.
6. In the **Write clone data to pool** attribute, specify the clone pool that targets the devices on which the clones will be stored.
7. If diagnostic mode is enabled, values may be assigned to the **Storage node to WRITE save sets** and the **Storage node to READ save sets** attributes. These attributes enable multiple storage nodes to be used as targets or sources.
8. To skip invalid save sets, select **Continue on save set error**. If *not* selected (default setting), an invalid save set will stop the clone operation and generate an error message.
9. The **Limit number of save set clones** attribute has a default value of 1. This allows only *one* clone to be made to the target pool for each save set.

This setting of 1 is useful, for example, if you need to manually restart an interrupted clone operation, and you want to produce only the missed clones but do *not* want to reproduce clones that were already successful.

This attribute limits the number of clones of each save set that may be produced in the clone pool for a scheduled clone operation. A zero (0) setting allows an unlimited number of duplicate clones.

Note: The NetWorker software allows only one unique save set clone to be stored on a single volume. If multiple clones of the same save set are created, each will be stored on a different volume in the pool.

10. Configure the schedule for the clone operation. [Figure 11 on page 75](#) shows the settings for a schedule.

Typically, to reduce the use of resources, you would schedule clone operations to be performed after the backup window has completed:

- Select **Enable** to enable the clone operation to run at its scheduled times.
- In the **Start Time** attribute, click the spin arrows, or type values, to set the start time.
- Select either **Weekly by day** or **Monthly by day** to display either a calendar week or a calendar month. In the calendar, select the days on which you want the clone operation to run.
- To repeat the clone operation within a day, specify an **Interval** time in hours. For example, if the start time is 6 a.m., and the interval is 6 hours, then the clone operation will run at 6 a.m., 12 p.m., and 6 p.m.

If the **Limit the number of save set clones** value is set, then the repeat clone operation will fail after the limit is reached.

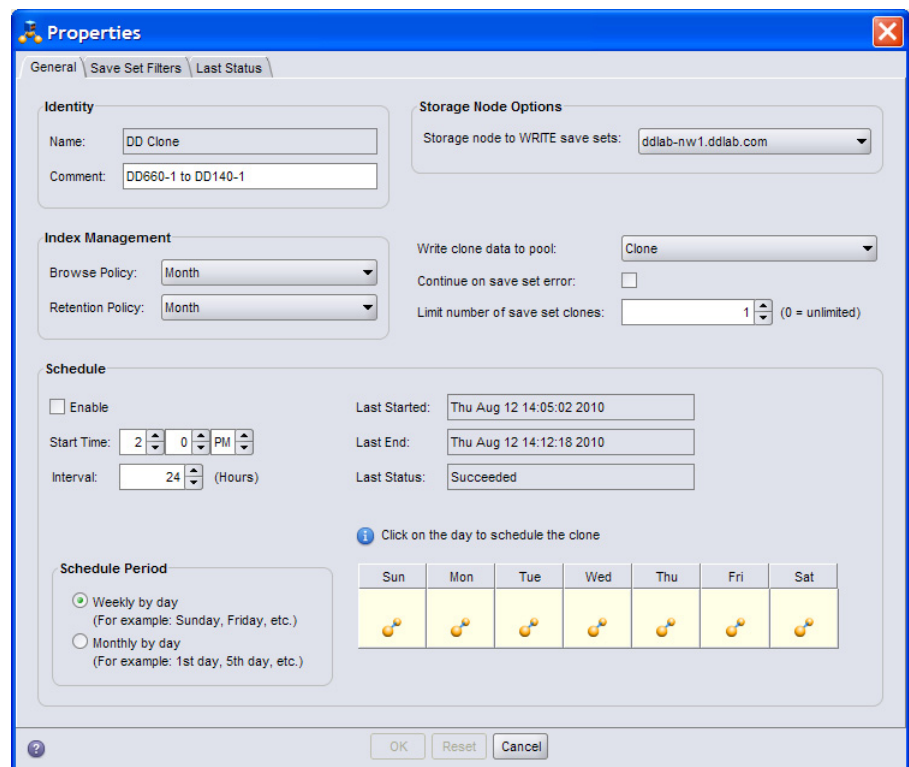


Figure 11 Clone properties, general settings

11. Select the **Save Set Filters** tab and specify the source save sets to include in your scheduled clone operation.

There are two main options as follows:

- Select **Clone save sets that match selections** and limit the save sets to be cloned by specifying the following details:
 - save groups
 - NetWorker clients
 - pools
 - save set levels
 - save set name, as specified in the NetWorker Client resource
 - save sets from the past number of days, weeks, months, or years
- Select **Clone specific save sets** and type the names of the source save sets to be cloned, either by save set ID (ssid) or by clone ID (clonid). Use a separate line for each save set name.

To find save set and clone IDs, in the **NetWorker Administration** window, click **Media**, select **save sets** in the browser tree, and specify values for your search. Alternatively, you may use the NetWorker **mminfo** command. The *EMC NetWorker Administration Guide* provides details.

12. Click **Preview Save Set Selection** to review save set filter settings.
13. Select **OK** to save the scheduled clone operation.
14. After the clone operation completes, verify the cloned data on the target devices. If required, test selected save sets to ensure they can be restored to the client hosts.

[Chapter 7, “Monitoring, Reporting, and Troubleshooting,”](#) provides details on the verification of NetWorker operations and the monitoring of clone operations, in the **NMC Clones** window, similar to a group.

Cloning by pools

To copy save sets from Data Domain storage to a device, a special pool may be specified. 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 DD Boost device.
- ◆ To copy the existing save sets from one DD Boost device to another DD Boost device, typically at a remote location for disaster recovery purposes.

Note: Selection of the volumes, pools, or save sets for a clone operation must ensure that all appropriate save sets are identified and successfully cloned. Failure to clone the entire backup set may result in a failure to perform a recovery from that clone copy. Expiration or deletion of the primary save sets should be performed only when you are confident that all save sets have been successfully cloned.

To specify a clone pool for the source volumes:

1. Ensure that cloning requirements are met. [“Clone requirements” on page 73](#) provides details.
2. Create a clone pool.
[“Create pools to target DD Boost devices” on page 59](#) provides details.
3. In NetWorker Administration window, click the **Media** view.
4. Click **Disk Volumes** and select volumes to clone.
5. Right-click and select **Clone**.
6. In the **Source volumes to clone** window, in the **Target clone media pool** list box, select the name of the clone pool, for example, *newclonepool*.
7. Once the pool is defined, the clone copies may be configured and run as a scheduled clone operation or manually started. [“Scheduling a clone operation” on page 74](#) provides details.

Starting a scheduled clone operation on demand

You can manually start a scheduled clone operation at any time without affecting the regularly-scheduled start time. There are two methods as follows:

Manually start a clone operation from the Configuration window

To manually start a scheduled clone operation:

1. From the Administration window, click **Configuration**.
2. In the browser tree, select **Clones**.
3. Right-click a clone resource in the right panel listings and select **Start**.

Manually start a clone operation from the Monitoring window

To manually start a scheduled clone operation:

1. From the Administration window, click **Monitoring**.
2. Select the **Clones** tab.
3. Right-click a clone resource and select **Start**.

CHAPTER 6

Recovering Data

This chapter includes the following sections:

- ◆ [Recovering deduplicated data](#) 80
- ◆ [Disaster recovery](#)..... 80

Recovering deduplicated data

Provided that the NetWorker and Data Domain systems have been properly configured, the recovery of deduplicated data from DD Boost devices is basically the same as for the recovery of non-deduplicated data, except that the backup is stored in two components in different places:

- ◆ Deduplicated client backup data is stored on DD Boost devices, which are on the Data Domain system.
- ◆ Backup metadata, which allows the backups to be browsed and selected for recovery and also determines how long data should be retained, is stored on the NetWorker server.

Recovery requirements

The requirements for recovery from DD Boost devices are as follows:

- ◆ All of the deduplicated data must be available on the Data Domain or Data Domain Archiver system and must be within its retention period before the data can be recovered.
- ◆ Both the Data Domain system and the NetWorker storage node must be online during the recovery of deduplicated data.

Supported NetWorker recovery procedures

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

- ◆ NetWorker software can be used to browse the media index and recover a file or save set.
- ◆ Directed recoveries are supported for supported NetWorker clients and storage nodes.
- ◆ Attempts to recover expired backup data can use NetWorker **scanner** program to rebuild a media index. However, this index provides access to only the surviving backup metadata for deduplication backups.

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

Disaster recovery

For the purpose of this guide, a disaster is any loss of data in which the computing environment required to restore that data is not available. Disaster recovery (DR) 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:

- ◆ Hardware or debilitating software failures
- ◆ Computer viruses that corrupt the computing system

- ◆ Infrastructure interruptions, inconsistencies, or loss of services, such as communications or network connections that result in damage to the computing environment

The *EMC NetWorker Disaster Recovery Guide* provides information on how to recover the NetWorker server, and a NetWorker client or storage node, from various types of disasters.

Potential losses

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

- ◆ Data Domain server where the deduplicated client backups are stored
- ◆ NetWorker storage node where the deduplication metadata for the backups is stored
- ◆ NetWorker server where the storage information (media indexes) for the backups is stored

Disaster recovery requirements

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

The two sites may be configured 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 media indexes cloned from the primary NetWorker server

[“Disaster recovery environment” on page 33](#) shows an example of a simple disaster recovery environment.

Disaster recovery scenarios

The procedures used to recover from disaster will vary depending on the circumstances, which could include:

- ◆ The deployment of the disaster recovery environment
- ◆ Which systems are affected by the disaster
- ◆ Time required to successfully recover from the disaster

The *EMC NetWorker Disaster Recovery Guide* provides information on how to recover the NetWorker server, and a NetWorker client or storage node, from various types of disasters.

CHAPTER 7

Monitoring, Reporting, and Troubleshooting

This chapter includes the following sections:

- ◆ [Monitoring Data Domain events, statistics, and logs](#) 84
- ◆ [Backup, recovery, and clone reports.....](#) 88
- ◆ [Troubleshooting](#) 91
- ◆ [Replace a failed or old storage node.....](#) 93

Monitoring Data Domain events, statistics, and logs

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

View backup statistics, log, and alerts in NetWorker

The NetWorker server window of NMC provides a comprehensive view of the backup status, log, and alerts for connected Data Domain systems:

Note: In some logs and other notifications, Client Direct operations are listed as direct file access (DFA), direct file save, or DIRECT_FILE operations.

1. Ensure that SNMP is configured for your Data Domain system.
[“Configure SNMP for Data Domain” on page 86](#) provides details.
2. In the NMC window for your NetWorker server, click the **Devices** view.
3. In the folder tree, select **Data Domain Systems**.
4. Select a Data Domain system in the folder tree and view its backup information.

[Figure 12 on page 85](#) shows an example:

- The **Status** table displays the following backup statistics:
 - Pre-Compression (/backup: pre-comp) — indicates the amount of space that would have been used for backup if the data had *not* been deduplicated and compressed. The NetWorker software tracks this value as the size of backups.
 - /backup: post-comp — indicates three values, the total capacity of the Data Domain system, the amount of disk space actually used, and the amount of space available.
 - /ddvar — indicates log file space used on the Data Domain filesystem.
 - Compression (Reduction) — is represented by two values:

$$\frac{\text{Pre-Compression}}{\text{Post-comp Used}}$$

$$(1 - \frac{\text{Post-comp Used}}{\text{Pre-Compression}}) \times 100\%$$
- The **Log** table displays a chronological list of events that occur during backup operations.
- The **Alerts** table displays messages for operational issues that might require administrative attention. 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**.

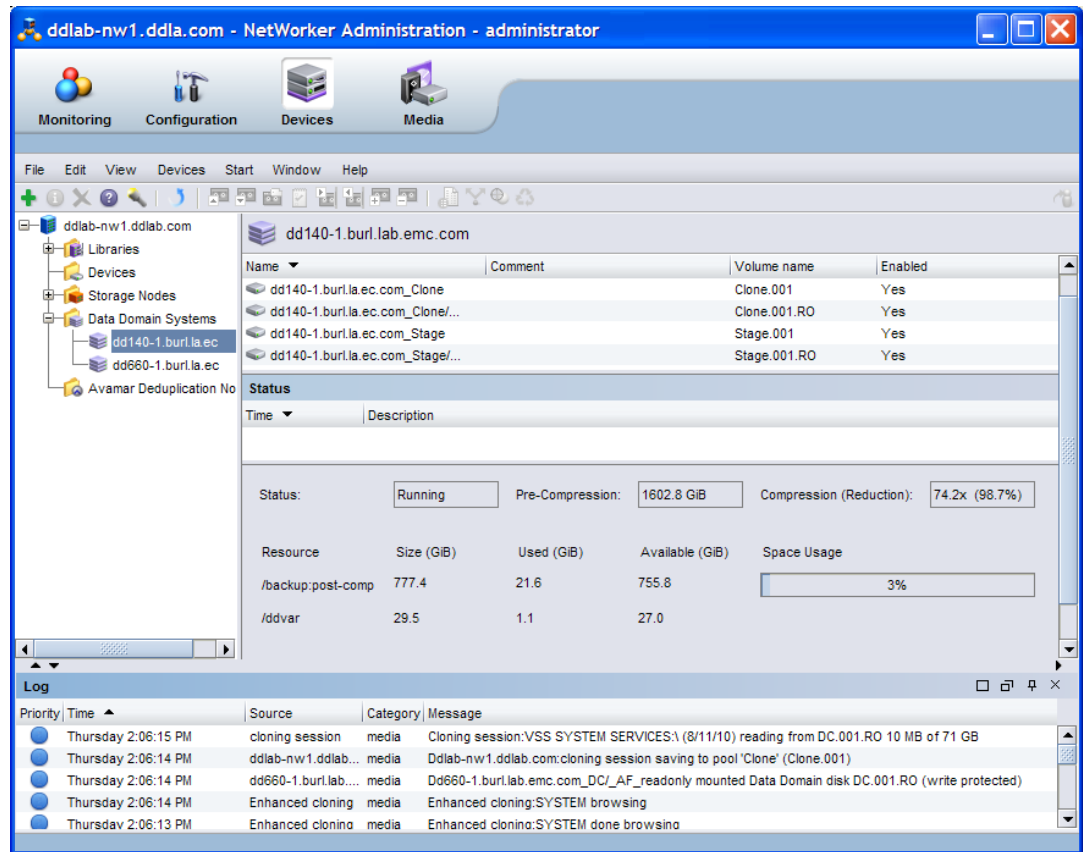


Figure 12 NetWorker Administration showing DD Boost devices

View backup statistics in NMC

To view the storage statistics for backups on a connected Data Domain system in NMC:

1. In the NMC **Enterprise** view, select a Data Domain host.
2. The right-hand table shows the storage statistics on the selected system.

View backup alerts (SNMP traps) in NMC

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

To view backup alerts on a connected Data Domain system in NMC:

1. Ensure that SNMP is configured for your Data Domain system.
[“Configure SNMP for Data Domain” on page 86](#) provides details.
2. In NMC, select the **Events** view.

The backup alerts (SNMP traps) are listed in chronological order.

Note: The same alert messages are also listed in the NetWorker **Alerts** table.

Delete individual messages from NMC Events and NetWorker Alerts

Individual messages can be deleted from the NetWorker Alerts and NMC Events tables by removing them from the NMC Events table. The two views show the same messages.

To delete unwanted messages:

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**.

The selected messages are removed from the NMC Events table and the NetWorker Alerts table.

Configure SNMP for Data Domain

To configure NMC to monitor Data Domain alerts (SNMP traps):

Note: On NMC servers with HP-UX operating systems, status monitoring (SNMP) in the Enterprise view is supported, but event monitoring (SNMPTRAP) in the Events view is *not* supported.

1. Provided you have viewing privileges, the NMC **Enterprise** view should list the Data Domain systems as network hosts. [“Add a host Data Domain system to NMC Enterprise view” on page 54](#) provides details.
2. Ensure that SNMP is enabled on the Data Domain system and Data Domain system is configured to send traps to the NMC server. [“Configure the Data Domain system for NetWorker” on page 50](#) provides details.
3. In the NMC **Enterprise** view left panel, right-click the Data Domain system to be monitored and select **Properties**.
4. In the **Properties** window, on the **Manage Data Domain** tab, select **Capture Events**.
If the box is *not* checked, NMC will monitor the status of the DD Boost devices, but will *not* monitor Data Domain SNMP traps, required to monitor event.
5. 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. [Figure 13 on page 87](#) shows an example.
6. Type a value for the **SNMP Process Port**. The default value is 162. This setting should agree with the firewall setting on the Data Domain system. [“Firewall requirements” on page 31](#) provides details.
7. Select the appropriate **SNMP Traps** that you want to monitor. Some traps are pre-selected. [Figure 13 on page 87](#) shows an example for Data Domain 4.8 alerts. Other versions may differ.

8. Click **OK**.

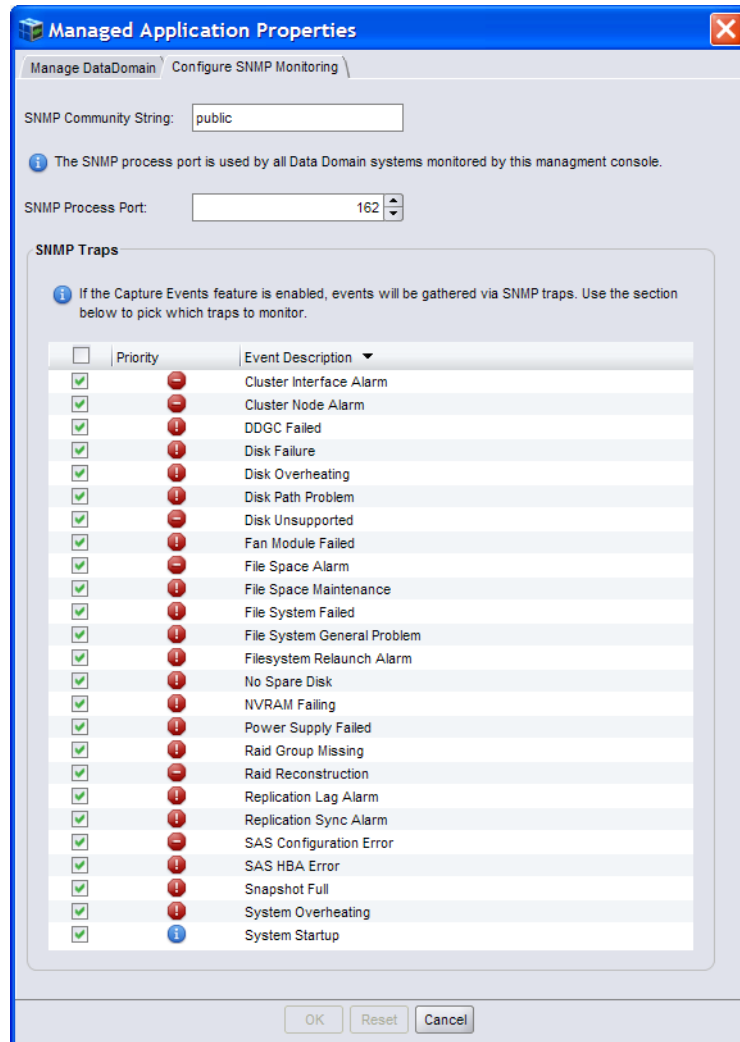


Figure 13 Data Domain alerts to monitor

Backup, recovery, and clone reports

Statistical reports of NetWorker Data Domain backup, recovery, and clone activities are available from the NMC Reports view.

Configure a report

To configure and display a Data Domain backup or clone report:

1. From the **NetWorker Management Console** window, click **Reports**.
2. Expand the **Reports** folder, expand the **Data Domain Statistics** folder, and then select a summary or statement report type to view.

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 by selecting the item parameters and clicking the **Remove** (◀), **Add** (▶), **Remove All** (◀◀), or **Add All** (▶▶) buttons as required.

If **Save Time** values are not specified, the report will display all the available data.

[Table 1 on page 88](#) lists details of report configuration parameters. The specific parameters available depend on the type of report selected.

[Figure 14 on page 89](#) shows an example report configuration.

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

Table 1 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.	Selected client names
Save Set Name	Selects one or more save sets. Values are case-sensitive and wild cards cannot be used.	Selected save set names
Save Time	Limits the report to a specified time range. Note: The date/time format available depends on the language locale of the operating system.	Save time (range)

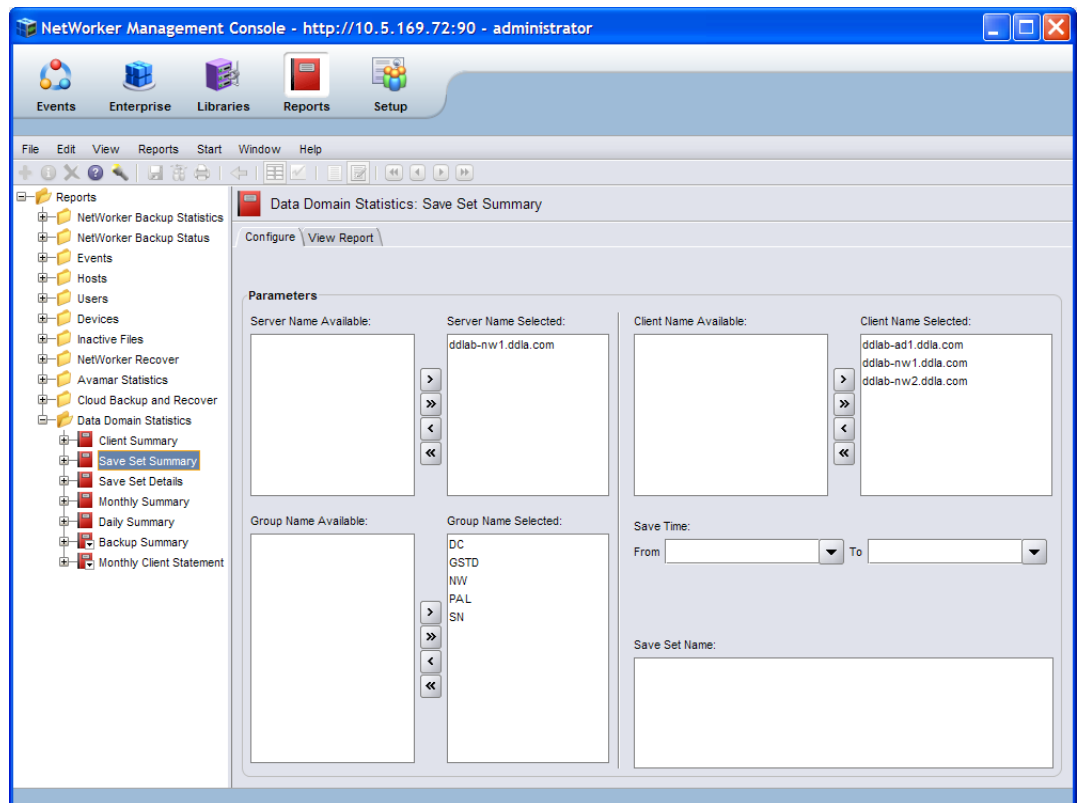


Figure 14 Report configuration

View a report

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 89 describes details of basic reports.
- ◆ “Table 2 on page 90 describes the basic reports available for Data Domain.” on page 90 describes details of drill-down reports.
- ◆ “Advanced Reporting” on page 91 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, time span, or attribute. You can modify the scope of a report by adjusting the parameters on the Configure tab.

Table 2 on page 90 describes the basic reports available for Data Domain.

Table 2 Data Domain basic reports

Report name	Purpose
Client Summary	For all or specified clients, provides the following statistics: <ul style="list-style-type: none"> • Amount of data: The amount of the data that would have been moved by using traditional NetWorker 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. • Number of Save Sets: The number of save sets included in the backup. • Number of Files: The number of files included in the backup.
Save Set Summary	For all or specified save sets, provides deduplication statistics on the following: <ul style="list-style-type: none"> • Amount of data: The amount of the data that would have been moved by using traditional NetWorker backup. • Target size: Size of the data after deduplication has taken place on the Data Domain System. • Deduplication ratio: Percentage of savings by using deduplication. • Number of Save Sets: The number of save sets included in the backup. • Number of Files: The number of files included in the save set.
Save Set Details	Displays details about each save set, including backup duration. The Save Set Details report provides statistics on the following: <ul style="list-style-type: none"> • 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 included 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 used in the top layer.

Reports can be run for groups, clients, or save sets. You can modify the scope of a report by adjusting the parameters on the Configure tab.

[Table 3 on page 91](#) lists the drill-down reports available for Data Domain Statistics.

Table 3 Data Domain statistics drill-down report

Report name	Purpose	Sequence
Backup Summary	Reports backup statistics over a period of time, starting from client summary down to individual save sets level.	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 recommended for larger environments where additional analysis with forecasts and trends are required.

Troubleshooting

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

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 used within the NetWorker software configuration.

Validate connections in both directions for both IP address and network names. If you use short names or aliases, then these should also be verified. Correct improper names by amending DNS entries or by populating the local hosts files:

- ◆ Use the NetWorker **nslookup** command to verify that network names resolve to the appropriate IP address on storage nodes and servers.
- ◆ On the Data Domain system, you may use the **net hosts** command.

[“Host naming guidelines” on page 29](#) provides suggestions for names.

Network connection issues

The network connections for a Data Domain system can be tested by using the **net lookup** command through an SSH telnet session, which requires sysadmin or administrator permissions. The Data Domain system can also show the current network configuration by using the **net show** and other network related commands, available through the Data Domain interface (http://DDR_IP_ADDRESS). Log in and go to the specific Data Domain system. Then select the Hardware > Network tabs to access the commands.

Diagram and verify all relevant network connections. A typical Data Domain network configuration should provide a minimum of two network connections, one of which is dedicated to administration and the other for backup data only. Although this is not a firm requirement, it is strongly recommended as a best practice. Make effective use of 10 GbE connectivity or multiple backup connections that are aggregated or “teamed” by using the **ifgroup** command on the Data Domain system.

[“Network connectivity and firewall” on page 30](#) 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

If the Data Domain filesystem becomes inaccessible or has been disabled and re-enabled, for example for service or testing, the devices are left in an unmounted state and backup operations for the devices will elicit the following error message.

“Waiting for 1 writable volume(s) to backup pool.”

To ensure the device is mounted:

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. In the **Enabled** column, ensure the device is enabled. If not enabled, right-click the device and select **Enable/Disable**.

NFS service errors

The NetWorker software requires the Data Domain NFS service to be enabled in order to access DD Boost devices. If NFS is not enabled, an error such as the following results, typically while attempting 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.”

[“Configure the Data Domain system for NetWorker” on page 50](#) describes how to enable NFS access.

Backup fails for older NetWorker application modules

The **Data Domain backup** attribute in the Client resource is intended to ensure that *only* DD Boost devices will be used for backups, even if the configured pool contains a mix of other device types, although pools with mixed devices are *not* recommended.

NetWorker application modules support backup to DD Boost devices, except some older modules do *not* support the **Data Domain backup** and **Target Pool** or **Pool** attributes in the Client resource and these attributes must *not* be set.

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.

Replace a failed or old storage node

If a storage node fails or is replaced with a new storage node, the data that is stored on the associated DD Boost devices may be recovered on a different or replacement 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 should not be affected and a complete recovery can be expected.
- ◆ If the volumes were mounted but *not* reading or writing data, complete recovery is still likely.
- ◆ If the devices were reading or writing at the time of the disruption, then data loss or corruption are more likely to have occurred and complete recovery cannot be assured.

Provided the volume structure of the devices are intact, if you service the failed storage node, or if you replace it with a new storage node that uses the same name, then the NetWorker server should be able to continue its operations with the existing devices with minimal impact.

However, if the replacement storage node has a different name or if you use the NetWorker server directly instead of a storage node, then you need to remove the devices from the NetWorker application and then re-create them as follows:

1. Ensure the following requirements:
 - 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 media database entries as the original.
2. From NMC, run the NetWorker application, select the **Devices** view, and select **Devices** in the navigation tree.

3. For each affected original remote (storage node-based) DD Boost device, right-click the device, select **Properties**, and record the following information:
 - **General** tab:
 - Name
 - Device Access Information
 - **Operations** tab:
 - Volume Name
 - Volume Pool
4. 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, **Devices** tree, right-click and unmount each affected device that is mounted. Mounted devices have a Volume Name.
 - b. In the **Media** view, **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.
5. 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 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) that were associated with the failed storage node.
 When you leave this page, a message box notifies you that the devices were previously associated with a different storage node. Confirm your selection.
 - d. On the **Configure Pool Information** page, specify the media pool for the devices and *unselect* the **Label and Mount** option. The devices must be manually mounted on the new storage node later in this procedure.

IMPORTANT

Ensure the Label and Mount checkbox is *not* selected. If the volume is re-labeled all the data will be *lost*. This action cannot be undone.

- 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 itself, which will handle the devices

The storage node must be running on the appropriate network and its hostname must be resolvable by DNS.

- f. Complete the wizard.
6. Manually mount each new device:
 - a. From the NMC window for your NetWorker server, click **Devices**.
 - b. In the navigation tree, select the Data Domain system.
 - c. In the right panel, right-click each device to be mounted, and select **Mount**.

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

If no errors are reported from this procedure, then the device and its volume should be available for use. Backup and recovery operations may require further configuration depending on the original settings and the purpose of the device recovery.

APPENDIX A

Upgrading from NetWorker 7.6 SP1

This appendix includes the following sections:

- ◆ [Data Domain storage folder structure and session changes](#) 98
- ◆ [Planning the upgrade](#)..... 98

Data Domain storage folder structure and session changes

Earlier versions of NetWorker software used a different storage structure on Data Domain systems. Due to the changes in this structure, special procedures are required for upgrade to the NetWorker 8.0 environment.

Previously, with NetWorker 7.6 SP1, DD Boost devices were created at the level of a Data Domain storage unit (SU) and there was one SU folder for each DD Boost device. There was no recommended limit on the number of these device SUs.

NetWorker 7.6 SP2 and later creates DD Boost devices as sub-folders of SU folders. There is no recommended limit to the number of DD Boost devices that may be created within each SU folder. For best performance, no more than 14 active SU folders should be used. However, increasing the number of DD Boost devices can impact Data Domain performance and maintenance. By default, the NMC Device Configuration Wizard names each SU folder after the short hostname of the NetWorker server that creates the devices.

Note: An exception where more storage units might be configured could be for service providers to allow quota enforcements, or for reporting on actual capacity used and deduplication ratios for each application group or end-customer. However, this practice should not be configured where it is not explicitly required.

The maximum number of backup sessions (save streams) enabled for each DD Boost device was increased to 60:

- ◆ NetWorker 7.6 SP1 used DD Boost library 2.2.2, which enforced a maximum of 10 sessions on each DD Boost device (SU). Due to this restriction on sessions, some environments might have been configured with a large number of DD Boost devices.
NetWorker 7.6 SP1 with DD Boost library 2.2.2 is supported by DD OS 4.8 and 4.9.
- ◆ NetWorker 7.6 SP2 and later uses DD Boost library 2.3.1 and later, which raises the number of sessions per device (nsrmm process) to a maximum of 60. Hence, fewer devices can handle the same sessions load as the previous devices.

NetWorker 7.6 SP2 with DD Boost library 2.3.1 is supported by DD OS 4.9 and 5.0 and later.

Planning the upgrade

DD Boost devices that were created as SU folders by NetWorker 7.6 SP1 are not modified in any way by upgrade to 7.6 SP2 and later and remain available without any changes. These old devices appear in the NetWorker Administration program's Devices view, but they do not appear in the NMC Device Configuration Wizard.

With DD OS 5.0 and later, it is recommended to use no more than 14 SU folders to avoid impairment of performance. These 14 SU folders can handle the DD Boost devices for 14 NetWorker datazones. However, an existing DD OS 4.9 system might already have more than 14 DD Boost devices that were configured at the level of Data Domain SU folders.

[Table 4 on page 99](#) provides details of the supported versions and upgrades.

Table 4 Supported upgrade options

NetWorker version	DD OS 4.8	DD OS 4.9	DD OS 5.0	DD OS 5.1
NW 7.6 SP1	Upgrade (optional) to DD OS 4.9	Upgrade to NetWorker 7.6 SP2 (except if DSP will be used)	Not supported	Not Supported
	Upgrade to DD OS 4.9 then upgrade to NetWorker 7.6 SP2	Upgrade to DD OS 5.0 then upgrade to NetWorker 7.6 SP2		
NW 7.6 SP2	Not supported	Upgrade to DD OS 5.0 as required.	Supported	Supported
NW 7.6 SP3	Not supported	Not supported	Supported	Recommended
NW 8.0	Not supported	Not supported	Supported	Redommended

The *EMC NetWorker Data Domain Deduplication Devices Release 7.6 SP2 Integration Guide* provides details of the upgrade procedures.

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	Feature of a NetWorker resource. It is a setting or information that the resource provides.
authorization code	Unique code that in combination with an associated enabler code unlocks the software for permanent use on a specific host computer. See also "license key."

B

backup	<ol style="list-style-type: none"> 1. Duplicate of database or application data, or entire computer system, stored separately from the original, which can be used to recover the original if it is destroyed or damaged. 2. Operation that saves data to a volume for use as a backup.
backup group	See "group."
backup volume	See "volume."
Boost	An optimized library and communication layer with a special Data Domain API that allows NetWorker to interact with the Data Domain system.
bootstrap	Save set that is essential for NetWorker disaster recovery procedures with the NetWorker server. The bootstrap consists of three components that reside on the NetWorker server: the media database, the resource database, and a server index.
browse policy	NetWorker policy that specifies the period of time during which backup entries are retained in the client file index. The index makes the associated backed-up data readily accessible for recovery by desktop users. See "retention policy."

C

client	Computer, workstation, or application server whose data can be backed up and restored with the NetWorker or NMDA software.
Client Direct	Client Direct, also known as direct file access (DFA), is a NetWorker feature that enables clients to send backup data over an IP network directly to AFTD or DD Boost storage devices, bypassing the NetWorker storage node. The storage node manages the devices used by the backup clients but does not handle the backup data.
client file index	Database maintained by the NetWorker server that tracks every database object, file, or filesystem 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	Duplicate copy of backed-up data, which is indexed and tracked by the NetWorker server. Single save sets or entire volumes can be cloned.
clone volume	Exact duplicate of a backup or archive volume. NetWorker software can track four types of volumes (backup, archive, backup clone, and archive clone). Save sets of these different types may not be intermixed on one volume. Clone volumes may be used in exactly the same way as the original backup or archive volume.
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.
Console server	See “NetWorker Management Console (NMC).”

D

database	<ol style="list-style-type: none"> 1. A collection of data arranged for ease and speed of update, search, and retrieval by computer software. 2. An 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	An optimized library and communication framework with a special Data Domain API that allows the NetWorker software to define and interact with storage devices on the Data Domain system.
DD Boost device	A logical storage device created on a Data Domain system, which 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 replaced by metadata pointers 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	<ol style="list-style-type: none"> 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 business operations and data in the event of hardware failure or software corruption.
distributed segment processing (DSP)	Part of the DD Boost interface that 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: <ul style="list-style-type: none"> ◆ Evaluation enablers or temporary enablers, such as a Powerlink entitlement for NetWorker, 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. <p>See also "license key."</p>
---------------------	--

G

group	Client computer or group of clients that are configured to back up files during a NetWorker scheduled backup, according to a single designated schedule or set of conditions.
--------------	---

H

host	Computer on a network.
host ID	Eight-character alphanumeric number that uniquely identifies a computer.
hostname	Name or address of a physical or virtual host computer that is connected to a network.

I

IFGROUP	A feature on a Data Domain system and other network devices used to aggregate two or more network interfaces to appear as a single network interface.
----------------	---

L

label	Electronic header on a volume used for identification by NetWorker or other data mover 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 magnetic tape, optical disk, or filesystem, to which backup data is written. See also "volume."

media index	Database that contains indexed entries of storage volume location and the lifecycle status of all data and volumes managed by the NetWorker server. Also known as media database.
media pool	See “pool.”
MTree	A logical partition of the namespace in a Data Domain 5.0 and later file system that can be used to group a set of files for management purposes, shortened from “managed tree.” MTrees can be referred to as storage units and is typically associated with a single NetWorker datazone.
multiple session backup and restore	Method of backing up or restoring multiple parallel streams of data simultaneously between a database and multiple media devices. Also known as multistripe.
multiplex	To simultaneously write data from more than one save set to the same storage device.
N	
NetWorker Data Domain device	See “DD Boost device.”
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 database).
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 starting from the root directory and working down the directory tree. A relative pathname indicates how to find a file starting from the current location.
policy	Set of defined rules for client backups that can be named and applied to multiple groups. Groups have dataset, schedule, and retention policies.
pool	<ol style="list-style-type: none"> 1. NetWorker feature that assigns specific backup data to be stored on selected media volumes. 2. Collection of NetWorker backup volumes to which specific data has been backed up.

R

recover	To restore data files from a backup volume to a client disk and apply transactional (redo) logs to the data to make it consistent with a given point in time.
remote device	<ol style="list-style-type: none"> 1. Storage device that is attached to a storage node that is separate from the NetWorker server. 2. Storage device located 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 also "clone."
resource	Software component that describes details of the NetWorker server or its clients. Clients, devices, schedules, groups, and policies are all NetWorker resources. Each resource has configurable attributes that define its properties.
resource database	NetWorker database of information about each configured resource.
restore	To retrieve individual data files from backup media and copy the files to disk, without applying transaction logs. See also "recover."
retention policy	NetWorker setting that determines the minimum period of time that backup data is retained on a volume available for recovery. After this time is exceeded, the data is eligible to be overwritten. See also "browse policy."
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	Group of files or a filesystem that has been backed up on storage media by using the NetWorker software.
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.
ssid	See "save set ID (ssid)."
storage device	See "device."
storage node	Storage device physically attached to a computer other than the NetWorker server, whose backup operations are administered from the controlling 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 program to report an error or status message.

V

virtual tape library (VTL) Software emulation of a physical tape library storage system.

- volume**
1. Unit of physical storage medium, such as a magnetic tape, optical disk, or filesystem used to store data. Backup data must be stored on a backup volume and cannot be stored on an archive volume or a clone volume.
 2. Identifiable unit of data storage that may reside on one or more computer disks.

volume ID (volid) Internal identification that NetWorker software assigns to a backup volume.

volume name Name that you assign to a backup volume when it is labeled. [See also "label."](#)

INDEX

A

- Add New Host Wizard 54
- AFTD disk configurations 27, 43
- authorization code 20

B

- backup
 - redirect 40
- backup failure
 - older NetWorker application modules 93
- bootstrap 65

C

- cascaded replication configuration 34
- CIFS formats 27, 43
- cleaning storage space 27
- client
 - configuration wizard 62
 - non-wizard configuration 63
- Client Direct
 - configuration 64
 - description 16
- clone
 - auto 72
 - by pool 76
 - feature 70
 - immediate 72
 - manually started 77
 - native non-deduplicated 71
 - nsrclone 72
 - operations 18
 - pool 70
 - regular 71
 - requirements 70, 73
 - scheduled 72, 74
 - to native format 18
- clone-controlled replication
 - cascade configuration 34
 - feature 18
 - for disaster recovery 33
 - process 71
 - scheduling 74
 - shared datazones 36
- compression not supported 18
- configuring
 - Data Domain for NetWorker 50
 - NetWorker for Data Domain 51

D

- daily summary report 90
- Data Domain
 - backup attribute 41, 62, 65
 - replication 42, 70

- Retention Lock not supported 18
- server 18

- Data Domain Archiver 25
- Data Domain backup storage 24
- DD Boost
 - enabling 51
 - feature 16
 - structure 17
- DD Boost devices
 - creating 52
 - creation by property windows 54
- DDBoost
 - username 50
- deduplication
 - feature 16
 - ratio 28
- device
 - creating 52
 - deactivate 66
 - delete 67
 - disable 66
 - erase data 67
 - identity 55
 - mounting 61
 - non-wizard configuration 55
 - removal 27
- Device Configuration Wizard 52
- direct file access (DFA) 16
- disaster recovery configuration 33
- distributed segment processing 16, 51
- DPA advanced reports 91
- dynamic drive sharing 59

E

- EMC online support website 9
- evaluation period 20

F

- filesystem clean command 27
- filesystem show space command 27
- firewall requirements 31
- full device 27

I

- ifgroup
 - multiple storage nodes 34, 38
 - NIC usage 31
- immediate clone 72

L

- label
 - device error 92

- template 58
- legacy save sets migration 43
- license requirements 20

M

- max nsrmmmd count 57
- max sessions 57
- memory requirement, storage node 26
- migration
 - legacy save sets 42, 43
 - scenarios 44
 - tape to local devices 45
 - tape to remote devices 45
 - VTL to local devices 48
 - VTL to remote devices 47
- monitoring
 - backup operations 84
 - backup statistics 85
 - devices 86
- mounting a device 61
- MTree 17, 24
- multiple datazones 36

N

- naming guidelines 29
- native non-deduplicated format 43
- NDMP deduplication backup 65
- net hosts add command 29
- network connectivity 31
- NetWorker
 - application modules 20
 - backup reports 90, 91
 - client 19
 - server 19
 - storage node 19
- NetWorker Management Console
 - description 18
 - portal 19
 - server 19
- NFS formats 27, 43
- NFS service errors 92
- NFS services 50, 53
- NIC connections 31
- nsrclone
 - command 74
 - script 72

O

- openstorage license 50
- optimized clone 18, 71
- OST username 50

P

- pool
 - associated with device 58
 - configuration 64
 - removing 27

R

- read-only devices 16
- recover
 - deduplicated data 80
 - disaster 80
- redirect backups 40
- regular clone 71
- remote device 56
- replication license 50
- replication, Data Domain 70
- reports
 - advanced 91
 - basic 89
 - configuring 88
 - drill-down 90
 - reviewing 89

S

- SAN support 26
- Save Set Details
 - NetWorker Backup Statistics Report 90
- Save Set Summary 90
 - NetWorker Backup Statistics Report 90
- save streams 25
- seed
 - as a data migration step 46
 - before redirected backups begin 40, 43
 - by native Data Domain replication 42
 - diminishing value of 45
 - new device with legacy data 42
 - not required on existing system 46
- sessions 25
- shared datazones 36
- short names, preferred over long names 29
- snmp
 - add ro-community command 55
 - add trap-host command 55
 - traps 51, 55
- statistics, monthly backup report 90
- storage node 26
 - dedicated 38
 - replace 93
- storage space, cleaning 27
- storage unit
 - maximum on Data Domain 24
 - structure 17, 24
- supported versions and upgrades 99

T

- target
 - devices 60
 - sessions 57
- throughput, maximize 31
- troubleshooting
 - connectivity issues 91
 - device access errors 92
 - IP addresses 30
 - network connections 92
 - network issues 92

no writable volume 92

U

upgrade to latest NetWorker 98

V

version requirements 18

virtual tape library 27, 43

volume name 58

Volume Shadow Copy Service 64

volume,erase 67

W

wizard

 Add New Host 54

 Client Configuration 62

 device configuration 52

