



# EMC Data Domain Operating System

Version 5.4

## Administration Guide

302-000-072

REV. 06

**EMC<sup>2</sup>**

Copyright © 2009-2014 EMC Corporation. All rights reserved. Published in USA.

Published September, 2014

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.

EMC<sup>2</sup>, EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

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

EMC Corporation  
Hopkinton, Massachusetts 01748-9103  
1-508-435-1000 In North America 1-866-464-7381  
[www.EMC.com](http://www.EMC.com)

# CONTENTS

	<b>Preface</b>	<b>11</b>
<b>Chapter 1</b>	<b>Introducing the EMC Data Domain System</b>	<b>15</b>
	About EMC Data Domain Systems.....	16
	EMC Data Domain System Features.....	16
	Data Integrity.....	16
	Data Compression.....	17
	Restore Operations.....	17
	EMC Data Domain Replicator.....	17
	Multipath and Load Balancing.....	18
	System Access.....	18
	Licensed Features.....	18
	How EMC Data Domain Systems Integrate into the Storage Environment.....	19
	Backup Software Requirements.....	21
<b>Chapter 2</b>	<b>Getting Started</b>	<b>23</b>
	About DD System Manager.....	24
	Using DD System Manager.....	24
	Logging In and Out of DD System Manager.....	24
	About the DD System Manager Interface.....	25
	Using the Configuration Wizard.....	28
	Using the CLI.....	28
	Logging into the System Using the CLI.....	28
	Finding Online Help for Commands.....	29
<b>Chapter 3</b>	<b>Managing Data Domain Systems</b>	<b>31</b>
	About Managing Data Domain Systems.....	32
	Managing System Availability.....	32
	Adding a System to DD System Manager.....	33
	Removing a System from DD System Manager.....	33
	Rebooting a System.....	34
	Powering a Data Domain System On or Off.....	34
	Working with Upgrade Images.....	35
	Viewing the Upgrade Package List.....	35
	Obtaining Upgrade Packages.....	35
	Upgrading a Data Domain System.....	36
	Removing an Upgrade Image.....	37
	Managing System Licenses.....	37
	Displaying Licenses.....	37
	Adding Licenses.....	37
	Removing Licenses.....	38
	Managing System Storage.....	38
	Viewing System Storage Information.....	38
	Physically Locating a Disk.....	42
	Configuring Storage.....	42
	Managing Network Connections.....	43
	Configuring Network Interfaces.....	43

	Configuring Network Settings.....	57
	Configuring Routes.....	61
	Managing Access to the System.....	63
	Managing Administrator Access.....	63
	Managing Local User Access to the System.....	68
	Managing NIS Servers and Workgroups.....	75
	Managing Windows Servers and Workgroups.....	77
	Managing General Configuration Settings.....	80
	Working with Email Settings.....	80
	Working with Time and Date Settings.....	82
	Working with System Properties.....	83
	Working with SNMP.....	83
	Managing Reporting and Logging.....	90
	Managing Autosupport Reporting.....	90
	Managing Support Bundles.....	92
	Managing Log Files.....	93
	Managing Remote System Power with IPMI.....	97
	Getting Started with IPMI.....	98
	Configuring IPMI for a Managed System.....	98
	Logging Into a Remote System for IPMI Power Management.....	102
	Managing Remote System Power After Login.....	102
<b>Chapter 4</b>	<b>Monitoring Data Domain Systems</b>	<b>105</b>
	About Monitoring Data Domain Systems.....	106
	Monitoring Using the DD Network Summary.....	106
	Viewing the DD Network Status.....	106
	About the System Summary Statistics.....	106
	About the Space Usage Statistics.....	107
	About Individual-System Statistics.....	107
	Monitoring a Single System.....	108
	Viewing the System Status Summary.....	108
	Viewing System Details.....	109
	About the Fibre Channel View.....	109
	About the Physical Resources View.....	109
	About the Access Groups View.....	110
	Monitoring Chassis Status.....	110
	Fans.....	111
	Temperature.....	111
	Power Supply.....	112
	PCI Slots.....	112
	NVRAM.....	112
	Working with Alerts.....	113
	Working with the Current Alerts Tab.....	113
	Working with the Alerts History Tab.....	115
	Working with the Notification View.....	116
	Configuring the Daily Alert Summary Distribution List.....	118
	Viewing Active Users.....	119
	Viewing System Statistics.....	119
	About the Performance Statistics Graphs.....	120
	Working with Reports.....	120
	Types of Reports.....	121
	Creating a Report.....	124
	Viewing Saved Reports.....	125
	Printing Saved Reports.....	125
	Deleting Saved Reports.....	125

	Renaming Saved Reports.....	126
	Viewing the Task Log.....	126
<b>Chapter 5</b>	<b>Working with the File System</b>	<b>127</b>
	About the File System.....	128
	How the File System Stores Data.....	128
	How the File System Reports Space Usage.....	128
	How the File System Uses Compression.....	129
	How the File System Implements Data Integrity.....	130
	How the File System Reclaims Storage Space with File System Cleaning.....	130
	Supported Interfaces.....	131
	Supported Backup Software.....	131
	Data Streams Sent to a Data Domain System.....	131
	File System Limitations.....	133
	Monitoring File System Usage.....	134
	Accessing the File System View.....	134
	Managing File System Operations.....	140
	Performing Basic Operations.....	140
	Performing Cleaning.....	143
	Modifying Basic Settings.....	144
	Fast Copy Operations.....	146
	Perform a Fast Copy Operation.....	146
<b>Chapter 6</b>	<b>Managing Encryption of Data at Rest</b>	<b>147</b>
	How Encryption of Data at Rest Works.....	148
	Configuring Encryption.....	148
	About Key Management.....	149
	Rectifying Lost or Corrupted Keys.....	150
	Key Manager Support.....	150
	Working with the RSA DPM Key Manager.....	151
	How the Cleaning Operation Works.....	153
	Working with the Embedded Key Manager.....	153
	Key Manager Setup.....	154
	RSA DPM Key Manager Encryption Setup.....	154
	Changing Key Managers after Setup.....	157
	Checking Settings for Encryption of Data at Rest.....	157
	Enabling and Disabling Encryption of Data at Rest.....	157
	Enable Encryption of Data at Rest.....	157
	Disable Encryption of Data at Rest.....	158
	Locking and Unlocking the File System.....	158
	Locking the File System.....	158
	Unlock the File System.....	159
	Changing the Encryption Algorithm.....	159
	Managing the Encryption Passphrase.....	160
	Changing the Encryption Passphrase.....	161
<b>Chapter 7</b>	<b>Working with DD Retention Lock</b>	<b>163</b>
	About DD Retention Lock Software.....	164
	DD Retention Lock Protocol.....	165
	DD Retention Lock Flow.....	165
	Supported Data Access Protocols.....	166
	Enabling DD Retention Lock on an MTree.....	167

	Enabling DD Retention Lock Governance on an MTree.....	167
	Enabling DD Retention Lock Compliance on an MTree.....	168
	Client-Side Retention Lock File Control.....	170
	Setting Retention Locking on a File.....	171
	Extending Retention Locking on a File.....	173
	Identifying a Retention-Locked File.....	173
	Specifying a Directory and Touching Only Those Files.....	174
	Reading a List of Files and Touching Only Those Files.....	174
	Deleting or Expiring a File.....	174
	Using ctime or mtime on Retention-Locked Files.....	174
	System Behavior with DD Retention Lock.....	175
	DD Retention Lock Governance.....	175
	DD Retention Lock Compliance.....	176
<b>Chapter 8</b>	<b>Working with MTrees</b>	<b>185</b>
	About MTrees.....	186
	Quotas.....	186
	About the MTree Overview Panel.....	186
	About the Summary View.....	187
	About the Space Usage View.....	190
	About the Daily Written View.....	191
	Monitoring MTree Usage.....	191
	Managing MTree Operations.....	192
	Create an MTree.....	192
	Configure and Enable/Disable MTree Quotas.....	193
	Delete an MTree.....	194
	Undelete an MTree.....	195
	Renaming an MTree.....	195
	Replicating a System with Quotas to One Without.....	195
<b>Chapter 9</b>	<b>Working with Snapshots</b>	<b>197</b>
	About Snapshots.....	198
	Monitoring Snapshots and Their Schedules.....	199
	About the Snapshots View.....	199
	Managing Snapshots.....	200
	Create a Snapshot.....	200
	Modify a Snapshot Expiration Date.....	201
	Rename a Snapshot.....	201
	Expiring a Snapshot.....	201
	Managing Snapshot Schedules.....	202
	Create a Snapshot Schedule.....	202
	Modify a Snapshot Schedule.....	203
	Delete a Snapshot Schedule.....	204
	Recover Data from a Snapshot.....	204
<b>Chapter 10</b>	<b>Working with CIFS</b>	<b>205</b>
	CIFS Overview.....	206
	Performing CIFS Setup.....	206
	Prepare Clients for Access to Data Domain Systems.....	206
	Enabling CIFS Services.....	207
	Naming the CIFS Server.....	207
	Setting Authentication Parameters.....	207
	Disable CIFS Services.....	210

	Working with Shares.....	210
	Creating Shares on the Data Domain System.....	211
	Modify a Share on a Data Domain System.....	212
	Creating a Share from an Existing Share.....	213
	Disable a Share on a Data Domain System.....	213
	Enable a Share on a Data Domain System.....	213
	Delete a Share on a Data Domain System.....	213
	Performing MMC Administration.....	214
	Connecting to a Data Domain System from a CIFS Client.....	214
	Display CIFS Information .....	215
	Managing Access Control.....	215
	Accessing Shares from a Windows Client.....	216
	Provide Domain Users Administrative Access.....	216
	Allow Access from Trusted Domain Users.....	216
	Allowing Administrative Access to a Data Domain System for Domain Users.....	216
	Restrict Administrative Access from Windows.....	217
	File Access.....	217
	Monitoring CIFS Operation.....	220
	Display CIFS Status.....	220
	Display CIFS Configuration.....	220
	Display CIFS Statistics.....	222
	Performing CIFS Troubleshooting.....	222
	Display Clients Current Activity.....	222
	Set the Maximum Open Files on a Connection.....	223
	Data Domain System Clock.....	223
	Synchronizing from a Windows Domain Controller.....	223
	Synchronize from an NTP Server.....	224
<b>Chapter 11</b>	<b>Working with NFS</b>	<b>225</b>
	About NFS.....	226
	Managing NFS Client Access to the Data Domain System.....	226
	Enable NFS Services.....	226
	Disable NFS Services.....	227
	Create an Export.....	227
	Modify an Export.....	228
	Creating an Export from an Existing Export.....	229
	Delete an Export.....	229
	Displaying NFS Information.....	230
	View NFS Status.....	230
	View NFS Exports.....	230
	View Active NFS Clients.....	230
<b>Chapter 12</b>	<b>Working with DD Boost</b>	<b>231</b>
	About Data Domain DD Boost Software.....	232
	Managing DD Boost with DD System Manager.....	232
	Set or Modify a DD Boost User Name.....	233
	Enable DD Boost.....	234
	Disable DD Boost.....	234
	View DD Boost Storage Unit.....	234
	Delete a Storage Unit.....	237
	Clear DD Boost Statistics.....	238
	DD Boost Options.....	238
	About Interface Groups.....	240

	Create Interface Groups.....	241
	Delete an Interface Group.....	241
	Enable/Disable an Interface Group.....	241
	Modify an Interface Group's Name/Interfaces.....	242
	Delete a Client from the Interface Group.....	242
	Modify a Client's Name or Interface Group.....	242
	Destroy DD Boost.....	243
	Managing Fibre Channel Transport.....	243
	Set Fibre Channel Server Name.....	243
	Create Access Group.....	244
	Delete Access Groups.....	244
	Monitoring DD Boost.....	245
	About the DD Boost Tabs.....	245
	Checking Interface Groups and Clients.....	247
<b>Chapter 13</b>	<b>Working with DD Virtual Tape Library</b>	<b>249</b>
	About EMC Data Domain Virtual Tape Library.....	250
	Planning a VTL.....	250
	Limitations of the VTL Feature.....	251
	Number of Supported Tape Drives.....	252
	Number of Supported Data Streams.....	252
	About Tape Barcodes.....	252
	About LTO Tape Drive Compatibility.....	253
	About the DD System Manager VTL View.....	254
	Setting Up a VTL.....	255
	Working with the VTL Service Operations.....	255
	Viewing the VTL Service Information Panel.....	255
	Working with Libraries.....	257
	Working with a Library.....	260
	Working with Tape Slots and CAPs.....	273
	Working with the Vault.....	274
	Working with Vault Pools.....	275
	Working with Storage Pools.....	276
	Viewing Storage Pools Information.....	276
	Creating Storage Pools.....	277
	Converting a Directory Pool to an MTree Pool.....	278
	Renaming Storage Pools.....	279
	Deleting Storage Pools.....	279
	Replicating Storage Pools.....	279
	Working with a Single Storage Pool.....	280
<b>Chapter 14</b>	<b>Working with SCSI Target</b>	<b>281</b>
	About SCSI Target.....	282
	Working with Access Groups.....	283
	Viewing Access Groups Information.....	283
	Configuring an Access Group.....	283
	Deleting an Access Group.....	286
	Working with an Access Group.....	287
	Viewing Access Group Information.....	287
	Configuring the NDMP Device TapeServer Group.....	288
	Working with Physical Resources.....	289
	About Endpoints.....	289
	About Initiators.....	294
	Setting a Loop ID.....	296



	FC Link Monitoring .....	296
<b>Chapter 15</b>	<b>Working with DD Replicator</b>	<b>297</b>
	About EMC Data Domain Replicator .....	298
	Replication Types .....	298
	Directory Replication .....	299
	MTree Replication .....	300
	Collection Replication .....	300
	Supported Replication Topologies .....	301
	One-to-One Replication .....	301
	Bi-Directional Replication .....	301
	One-to-Many Replication .....	302
	Many-to-One Replication .....	302
	Cascaded Replication .....	303
	Using Encryption of Data at Rest with Replication .....	304
	Encryption on the Wire .....	304
	Low Bandwidth Optimization .....	305
	Bandwidth Delay Settings .....	305
	About the Replication View .....	305
	Replication Status .....	305
	Summary View .....	306
	DD Boost View .....	308
	Topology View .....	309
	Performance View .....	309
	Advanced Settings View .....	309
	Preparing to Configure Replication .....	310
	Limitations on Number of Contexts .....	310
	Configuring Replication .....	310
	Creating a Replication Pair .....	311
	Enabling and Disabling a Replication Pair .....	314
	Deleting a Replication Pair .....	314
	Converting a Directory Replication Pair to an MTree .....	314
	Changing Host Connection Settings .....	315
	Managing Bandwidth with Throttling .....	316
	Changing Network Settings .....	317
	Resynchronizing Data in a Replication Pair .....	318
	Resyncing a Directory, MTree, or Pool Replication Pair .....	319
	Recovering Data from a Replication Pair .....	320
	Recovering Directory Pool Data .....	320
	Recovering Collection Replication Pair Data .....	320
	Recovering Directory Replication Pair Data .....	321
	Aborting a Replication Pair Recovery .....	321
	Replication Seeding .....	321
	Monitoring Replication .....	322
	Checking Replication Status .....	322
<b>Chapter 16</b>	<b>Working with DD Extended Retention</b>	<b>325</b>
	About DD Extended Retention Software .....	326
	Getting Started .....	327
	Accessing Data .....	328
	Supported Replication Types .....	328
	Licenses .....	330
	Using the DD System Manager .....	331
	Data Domain Provided Hardware .....	332

Customer-Provided Infrastructure.....	334
Initial Setup.....	334
Initial Configuration.....	335
Configuring Data Movement.....	337
Data Movement Policy.....	337
Administration.....	339
Avoiding Common Sizing Errors.....	339
Cleaning and Snapshots.....	339
Expanding an Active or Archive Tier.....	339
Deleting a Retention Unit.....	340
Expanding an Archive Unit.....	340
Reclaiming Space in the Retention Tier.....	341
Upgrading Data Domain Systems for Extended Retention.....	341
Changing Retention Tier Compression.....	342
Upgrades and Recovery.....	343
Upgrading to DD OS 5.4.....	343
Upgrading a Data Domain System Controller with the DD Extended Retention Software Option Enabled.....	343
Replacing Data Domain Systems.....	343
Replication Recovery.....	344
Recovering a System with the DD Extended Retention Software Option Enabled.....	344
Recovering from System Failures.....	345
Reusing a Retention Unit.....	345

# 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 technical support professional if a product does not function properly or does not function as described in this document.

---

## Note

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

---

## Purpose

This guide explains how to manage the EMC Data Domain® systems with an emphasis on procedures using the EMC Data Domain System Manager (DD System Manager), a browser-based graphical user interface (GUI). If an important administrative task is not supported in DD System Manager, the Command Line Interface (CLI) commands are described.

---

## Note

- ◆ DD System Manager was formerly known as the Enterprise Manager.
- ◆ In some cases, a CLI command may offer more options than those offered by the corresponding DD System Manager feature. See the *EMC Data Domain Operating System Command Reference Guide* for a complete description of a command and its options.

---

## Audience

This guide is for system administrators who are familiar with standard backup software packages and general backup administration.

## Related documentation

The following Data Domain system documents provide additional information:

- ◆ Installation and setup guide for your system, for example, *EMC Data Domain DD 2500 Storage System, Installation and Setup Guide*
- ◆ *EMC Data Domain Operating System USB Installation Guide*
- ◆ *EMC Data Domain Operating System DVD Installation Guide*
- ◆ *EMC Data Domain Operating System Release Notes*
- ◆ *EMC Data Domain Operating System Initial Configuration Guide*
- ◆ *EMC Data Domain Product Security Guide*
- ◆ *EMC Data Domain Operating System Command Reference Guide*
- ◆ *EMC Data Domain Operating System MIB Quick Reference*
- ◆ *EMC Data Domain Operating System Offline Diagnostics Suite User's Guide*
- ◆ Hardware overview guide for your system, for example, *EMC Data Domain DD4200, DD4500, and DD7200 Systems, Hardware Overview*

- ◆ Field replacement guides for your system components, for example, *Field Replacement Guide, Data Domain DD4200, DD4500, and DD7200 Systems, IO Module and Management Module Replacement or Upgrade*
- ◆ *EMC Data Domain, System Controller Upgrade Guide*
- ◆ *EMC Data Domain Expansion Shelf, Hardware Guide* (for shelf model ES20 or ES30)
- ◆ *EMC Data Domain Boost for OpenStorage Administration Guide*
- ◆ *EMC Data Domain Boost for Oracle Recovery Manager Administration Guide*
- ◆ *EMC Data Domain Boost SDK Programmer's Guide*
- ◆ *Statement of Volatility for the Data Domain DD2500 System*
- ◆ *Statement of Volatility for the Data Domain DD4200, DD4500, or DD7200 System*

If you have the optional RSA Data Protection (DPM) Key Manager, see the latest version of the *RSA Data Protection Manager Server Administrator's Guide*, available with the RSA Key Manager product.

### Special notice conventions used in this document

EMC uses the following conventions for special notices:



Indicates a hazardous situation which, if not avoided, will result in death or serious injury.

---



Indicates a hazardous situation which, if not avoided, could result in death or serious injury.

---



Indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.

---



Addresses practices not related to personal injury.

---

### Note

Presents information that is important, but not hazard-related.

---

### Typographical conventions

EMC uses the following type style conventions in this document:

#### Table 1 Typography

<b>Bold</b>	Indicates interface element names, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Highlights publication titles listed in text
Monospace	Indicates system information, such as: <ul style="list-style-type: none"> <li>• System code</li> <li>• System output, such as an error message or script</li> </ul>

**Table 1** Typography (continued)

	<ul style="list-style-type: none"> <li>• Pathnames, filenames, prompts, and syntax</li> <li>• Commands and options</li> </ul>
<i>Monospace italic</i>	Highlights a variable name that must be replaced with a variable value
<b>Monospace bold</b>	Indicates text for user input
[ ]	Square brackets enclose optional values
	Vertical bar indicates alternate selections—the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

**Where to get help**

The following topics describe how to get more product information and contact technical support.

**Product information**

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

**EMC Data Domain product documentation**

To view documentation for EMC Data Domain products, go to EMC Online Support and click Support by Product below the Search box. Type **Data Domain** in the Find a Product box, wait for those words to appear in the list of matches below the box, and click the words. Then click »». In the list of categories under the Search box, click Documentation.

- ◆ The Product choices let you filter results by Data Domain system model number, such as DD990, or by DD OS software release.
- ◆ The Content Type choices let you filter results by category. Click More under Content Type to see all of the categories. The categories that contain end-user and compatibility documentation are:
  - Manuals and Guides, for the software and hardware manuals for your system, and for integration guides that explain how to use EMC Data Domain systems with backup software and other products
  - Release Notes, for specific versions of the EMC Data Domain Operating System and EMC Data Domain products
  - Compatibility Document, for guides that show which EMC and third-party components are compatible

**Technical support**

Go to EMC Online Support and click Service Center. You will see several options for contacting EMC Technical Support. 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: [DPAD.Doc.Feedback@emc.com](mailto:DPAD.Doc.Feedback@emc.com).



# CHAPTER 1

## Introducing the EMC Data Domain System

This chapter includes:

- ◆ [About EMC Data Domain Systems](#)..... 16
- ◆ [EMC Data Domain System Features](#)..... 16
- ◆ [How EMC Data Domain Systems Integrate into the Storage Environment](#)..... 19

## About EMC Data Domain Systems

EMC Data Domain systems are disk-based inline deduplication appliances and gateways that provide data protection and disaster recovery (DR) in the enterprise environment.

All systems run the EMC Data Domain Operating System (DD OS), which provides both a command-line interface (CLI) for performing all system operations, and the EMC Data Domain System Manager (DD System Manager) graphical user interface (GUI) for configuration, management, and monitoring.

---

### Note

DD System Manager was formerly known as the Enterprise Manager.

---

Systems consist of appliances that vary in storage capacity and data throughput. Systems are typically configured with expansion shelves that add storage space.

---

### Note

Legacy Data Domain Gateway systems store all data on qualified third-party storage arrays through a Fibre Channel interface. See the list of qualified arrays in the *Storage Array Compatibility List* at <https://support.emc.com>.

---

## EMC Data Domain System Features

The following sections describe how Data Domain systems ensure data integrity and provide multiple levels of data compression, reliable restoration, data replication, and multipath configuration.

- ◆ [Data Integrity on page 16](#)
- ◆ [Data Compression on page 17](#)
- ◆ [Restore Operations on page 17](#)
- ◆ [EMC Data Domain Replicator on page 17](#)
- ◆ [Multipath and Load Balancing on page 18](#)
- ◆ [System Access on page 18](#)
- ◆ [Licensed Features on page 18](#)

### Data Integrity

The DD OS Data Invulnerability Architecture™ protects against data loss from hardware and software failures.

- ◆ When writing to disk, the DD OS creates and stores checksums and self-describing metadata for all data received. After writing the data to disk, the DD OS then recomputes and verifies the checksums and metadata.
- ◆ An append-only write policy guards against overwriting valid data.
- ◆ After a backup completes, a validation process examines what was written to disk and verifies that all file segments are logically correct within the file system and that the data is identical before and after writing to disk.
- ◆ In the background, the online verify operation continuously checks that data on the disks is correct and unchanged since the earlier validation process.



- ◆ Storage in most Data Domain systems is set up in a double parity RAID 6 configuration (two parity drives). Additionally, most configurations include a hot spare in each enclosure, except the DD1xx series systems, which have eight disks. Each parity stripe has block checksums to ensure that data is correct. Checksums are constantly used during the online verify operation and while data is read from the Data Domain system. With double parity, the system can fix simultaneous errors on as many as two disks.
- ◆ To keep data synchronized during a hardware or power failure, the Data Domain system uses NVRAM (non-volatile RAM) to track outstanding I/O operations. An NVRAM card with fully charged batteries (the typical state) can retain data for a period of hours, which is determined by the hardware in use.
- ◆ When reading data back on a restore operation, the DD OS uses multiple layers of consistency checks to verify that restored data is correct.

## Data Compression

Using Global Compression, a Data Domain system eliminates redundant data from each backup image and stores only unique data.

Duplicate data is stored only once. The storage of unique data is invisible to backup software.

DD OS data compression is independent of data format. Data can be structured, such as databases, or unstructured, such as text files. Data can derive from file systems or from raw volumes.

Typical compression ratios are 20-to-1, on average, over many weeks. This ratio assumes there are weekly full backups and daily incremental backups. A backup that includes many duplicate or similar files (files copied several times with minor changes) benefits the most from compression.

Depending on backup volume, size, retention period, and rate of change, the amount of compression can vary. The best compression happens with backup volume sizes of at least 10 MiB (MiB is the base 2 equivalent of MB).

To take full advantage of multiple Data Domain systems, a site that has more than one Data Domain system should consistently backup the same client system or set of data to the same Data Domain system. For example, if a full back up of all sales data goes to Data Domain system A, the incremental backups and future full backups for sales data should also go to Data Domain system A.

## Restore Operations

With disk backup using the Data Domain system, incremental backups are always reliable and can be easily accessed. Furthermore, with a Data Domain system, you can perform full backups more frequently without the penalty of storing redundant data. With tape backups, a restore operation may rely on multiple tapes holding incremental backups. Also, the more incremental backups a site has on multiple tapes, the more time-consuming and risky the restore process. One bad tape can kill the restore.

From a Data Domain system, file restores create little or no contention with backup or other restore operations. Unlike tape drive backups, multiple processes can access a Data Domain system simultaneously. A Data Domain system allows your site to offer safe, user-driven, single-file restore operations.

## EMC Data Domain Replicator

The EMC Data Domain Replicator sets up and manages the replication of backup data between two Data Domain systems. After replication is started, the source Data Domain

system automatically sends any new backup data to the destination Data Domain system.

A Replicator pair deals with a complete data set, a directory, or an MTree from a source Data Domain system that is sent to a destination Data Domain system. An individual Data Domain system can be a part of multiple replication pairs and can serve as a source for one or more pairs and a destination for one or more pairs.

## Multipath and Load Balancing

Multipath configuration and load balancing is supported on Data Domain systems that have at least two HBA ports. In a multipath configuration on a Data Domain system, each of two HBA ports on the system is connected to a separate port on the backup server. On a Data Domain gateway, each of two HBA ports are connected to a separate port on the array that the gateway uses as a backup destination. For more on multipath configuration, see the *EMC DD OS Initial Configuration Guide* and the *EMC DD OS Command Reference Guide*.

## System Access

The DD OS provides the following ways to access the system for configuration and management:

- ◆ CLI—A Data Domain system has a complete command set available to users in a command-line interface. Commands perform initial system configuration and changes to individual system settings as well as display system and operation status. The command-line interface is available through a serial console or through Ethernet connections using SSH or Telnet.
- ◆ DD System Manager—A browser-based graphical user interface that is available through Ethernet connections. Use DD System Manager to perform initial system configuration, make configuration changes after initial configuration, display system and component status, and generate reports and charts. DD System Manager also provides centralized management for one or multiple systems.

---

### Note

Some older systems support access using a keyboard and monitor attached directly to the system.

---

## Licensed Features

A license is required to operate each of the following features on a Data Domain system. Consult with your EMC BRS Data Domain representative for more information and to purchase licensed features.

**Table 2** Features Requiring Licenses

Feature/License Name	Description
EMC Data Domain ArchiveStore	Licenses Data Domain systems for archive use, such as file and email archiving, file tiering, and content and database archiving.
EMC Data Domain Boost	Enables the use of a Data Domain system with the following applications: EMC Avamar, EMC NetWorker, Oracle RMAN, Quest vRanger, Symantec Veritas NetBackup (NBU), and

**Table 2** Features Requiring Licenses (continued)

Feature/License Name	Description
	Backup Exec. The managed replication feature of DD Boost also requires the DD Replicator license.
EMC Data Domain Encryption	Allows data on system drives or external storage to be encrypted while being saved, and then locked before moving it to another location.
EMC Data Domain Expansion Storage	Allows external shelves to be added to the Data Domain system for additional capacity.
EMC Data Domain Extended Retention (formerly DD Archiver)	Licenses the Extended Retention storage feature. See <a href="#">About DD Extended Retention Software on page 326</a> for additional information.
EMC Data Domain I/OS (for IBM i operating environments)	An I/OS license is required when VTL is used to backup systems in the IBM i operating environment. Apply this license before adding virtual tape drives to libraries.
EMC Data Domain NDMP Tape Server	Enables the use of a Data Domain system as a virtual tape library that supports backups of NAS devices over Ethernet/IP networks.
EMC Data Domain Replicator	Adds DD Replicator for replication of data from one Data Domain system to another. A license is required on each system.
EMC Data Domain Retention Lock Compliance Edition	Meets the strictest data retention requirements from regulatory standards such as SEC17a-4.
EMC Data Domain Retention Lock Governance Edition	Protects selected files from modification and deletion before a specified retention period has expired.
EMC Data Domain Shelf Capacity	Enables a Data Domain system to expand the active tier storage capacity beyond the entry capacity defined for that system.
EMC Data Domain Virtual Tape Library (VTL)	Enables the use of a Data Domain system as a virtual tape library over a Fibre Channel network.
Gateway Expanded Storage Level 2	Enables gateway systems to support additional usable capacity.
Gateway Expanded Storage Level 3	Enables gateway systems to support additional capacity greater than Expanded Storage Level 2 usable capacity.

See [Managing System Licenses on page 37](#) for instructions on viewing and installing licenses.

## How EMC Data Domain Systems Integrate into the Storage Environment

EMC Data Domain systems integrate easily into existing data centers:

- ◆ All Data Domain systems can be configured as storage destinations for leading backup and archiving applications using NFS, CIFS, Boost, or VTL protocols.
- ◆ Consult the compatibility matrices at <https://support.emc.com> for information on the applications that work with the different configurations.

- ◆ The Data Domain gateway series uses external disk arrays for storage. Data Domain gateways work with Data Domain arrays and are also qualified with storage systems from several leading enterprise storage providers.
- ◆ Multiple backup servers can share one Data Domain system.
- ◆ One Data Domain system can handle multiple simultaneous backup and restore operations.
- ◆ Multiple Data Domain systems can be connected to one or more backup servers.


For use as a backup destination, a Data Domain system can be configured either as a disk storage unit with a file system that is accessed through an Ethernet connection or as a virtual tape library (VTL) that is accessed through a Fibre Channel connection. The VTL feature enables Data Domain systems to be integrated into environments where backup software is already configured for tape backups, minimizing disruption.

Configuration is performed both in the DD OS, as described in the relevant sections of this guide, and in the backup application, as described in the backup application’s administrator guides and in Data Domain application-related guides and tech notes.

- ◆ All backup applications can access a Data Domain system as either an NFS or a CIFS file system on the Data Domain disk device.
- ◆ The following applications work with a Data Domain system using the DD Boost interface: EMC Avamar, EMC NetWorker, Oracle RMAN, Quest vRanger, Symantec Veritas NetBackup (NBU), and Backup Exec.

The following figure shows a Data Domain system integrated into an existing basic backup configuration.

**Figure 1** Data Domain System Integrated into a Storage Environment



1. Primary storage
2. Ethernet
3. Backup server
4. SCSI/Fibre Channel
5. Gigabit Ethernet or Fibre Channel
6. Tape system
7. Data Domain system
8. Management
9. NFS/CIFS/VTL/DD Boost

**Figure 1** Data Domain System Integrated into a Storage Environment (continued)

10. Data Verification
11. Data Domain file system
12. Global compression
13. RAID

As shown in Figure 1, data flows to a Data Domain system through an Ethernet or Fibre Channel connection. Immediately, the data verification processes begin and are continued while the data resides on the Data Domain system. In the file system, the DD OS Global Compression™ algorithms dedupe and compress the data for storage. Data is then sent to the disk RAID subsystem. When a restore operation is required, data is retrieved from Data Domain storage, decompressed, verified for consistency, and transferred via Ethernet to the backup servers using Ethernet (for NFS, CIFS, DD Boost), or using Fiber Channel (for VTL and DD Boost).

## Backup Software Requirements

This section provides information needed to set up a Data Domain system as a storage destination for an application. It includes:

- ◆ [Application Compatibility Matrices and Integration Guides on page 21](#)
- ◆ [Viewing Data Domain Application-Related Documents on page 21](#)
- ◆ [Generic Application Configuration Guidelines on page 21](#)

### Application Compatibility Matrices and Integration Guides

The EMC support Web site provides compatibility matrices and integration documents on how to integrate Data Domain systems as storage destinations with qualified backup applications. Integration is generally easy and straightforward. The integration guides provide specific parameters and limitations that must be understood and followed for the applications to work with Data Domain systems.

To locate these documents, search at <https://support.emc.com/documentation>.

### Viewing Data Domain Application-Related Documents

#### Procedure

1. Log into the EMC Support portal at <https://support.emc.com/documentation>.
2. Use the search function to locate documentation related to the application.

### Generic Application Configuration Guidelines

The DD OS accommodates relatively large streams of sequential data from backup software and is optimized for high throughput, continuous data verification, and high compression. It also accommodates the large numbers of smaller files in nearline storage (DD ArchiveStore).

Data Domain system performance is best when storing data from applications that are not specifically backup software under these circumstances:

- ◆ Data is sent to the Data Domain system as sequential writes (no overwrites).
- ◆ Data is neither compressed nor encrypted before being sent to the Data Domain system.



# CHAPTER 2

## Getting Started

This chapter includes:

- ◆ [About DD System Manager](#)..... 24
- ◆ [Using DD System Manager](#)..... 24
- ◆ [Using the Configuration Wizard](#)..... 28
- ◆ [Using the CLI](#)..... 28

## About DD System Manager

DD System Manager is a browser-based graphical user interface, available through Ethernet connections, for managing up to 20 systems (depending on the model) at any location. DD System Manager provides a single, consolidated management interface that allows for configuration and monitoring of many system features and system settings.

DD System Manager provides real-time graphs and tables that allow you to monitor the status of system hardware components and configured features.

Additionally, a command set that performs all system functions is available to users at the command-line interface (CLI). Commands configure system settings and provide displays of system hardware status, feature configuration, and operation.

The command-line interface is available through a serial console or through an Ethernet connection using SSH or Telnet.

---

### Note

Some older systems support access using a keyboard and monitor attached directly to the system.

---

## Using DD System Manager

This section describes how to log into and out of DD System Manager and describes its graphical user interface. It includes:

- ◆ [Logging In and Out of DD System Manager on page 24](#)
- ◆ [About the DD System Manager Interface on page 25](#)

## Logging In and Out of DD System Manager

### Procedure

1. Open a web browser and enter the IP address or hostname to connect to DD System Manager. It must be:
  - A fully qualified domain name (for example, `http://dd01.emc.com`)
  - A hostname (`http://dd01`)
  - An IP address (`http://10.5.50.5`)

---

### Note

DD System Manager uses HTTP port 80 and HTTPS port 443. If your Data Domain system is behind a firewall, you may need to enable port 80 if using HTTP, or port 443 if using HTTPS to reach the system. The port numbers can be easily changed if security requirements dictate.

---

2. For HTTPS secure login, click **Login using enhanced security**.  
When HTTPS is used, the browser warns if a certificate is not authorized.
3. Enter a username and password (assigned during the initial configuration).  
See the *EMC DD OS Initial Configuration Guide* for details.
4. Click **Login**.



The Summary view appears in the Information panel. For details on this view, see [DD Network Summary View on page 25](#).

5. To log out, click the **Log Out** link in the DD System Manager banner.

## About the DD System Manager Interface

This section describes the main views of DD System Manager and its components. The following topics are covered:

- ◆ [DD Network Summary View on page 25](#)
- ◆ [Single System View on page 25](#)
- ◆ [Navigation Panel on page 26](#)
- ◆ [Banner on page 26](#)
- ◆ [Information Panel on page 27](#)
  - [Tab Bar on page 27](#)
  - [More Tasks Menu on page 27](#)
  - [Help Buttons and Menus on page 27](#)
- ◆ [View End User License Agreement \(EULA\) on page 28](#)

### DD Network Summary View

After you log into DD System Manager, the default DD Network Summary view (see Figure 1) appears (if the default view does not appear, select DD Network in the navigation panel). This view presents a status overview of all systems managed by DD System Manager and summarizes key operating information. A tally of alerts and charts of disk space enable you to identify problems. See the section [Monitoring Using the DD Network Summary on page 106](#) for more information about this view.

Click **+** in the navigation panel to display the systems that DD System Manager is managing.

### Single System View

To display information about a single system, select the system in the Navigation panel. The Status Summary view displays important data about the selected system and displays a set of tabs at the top of the Information panel. You can use these tabs to configure and monitor the selected system. For more information on understanding this view, see [Monitoring a Single System on page 108](#).

### Page Elements

This figure shows the principal elements in DD System Manager pages.

**Figure 2** DD System Manager Page Components

1. Banner
2. Navigation Panel
3. Information Panel

## Navigation Panel

The Navigation panel, always visible on the left edge of the page, displays a hierarchal tree of the systems managed by the DD System Manager and the **Reports** and **Task Log** buttons.

- ◆ Click the top-level **DD Network** icon to display the global Summary page.
- ◆ Click the **Add** or **Remove** icons to add or remove a system managed by the DD System Manager. See [Adding a System to DD System Manager on page 33](#) and [Removing a System from DD System Manager on page 33](#) for details.
- ◆ Expand the DD Network and select a system in the tree to open the Status Summary view, where tabs allow you to configure and monitor the selected system.
- ◆ Click **Reports** to open a report generator tool and provide access to saved reports for the selected system. Reports for file system and replication usage statistics can be generated. See [Working with Reports on page 120](#) for more information.
- ◆ Click **Task Log** to show a history of tasks that have been performed on the system you are logged into. See [Viewing the Task Log on page 126](#) for more information.

## Banner

The DD System Manager banner appears above the Navigation and Information panels and displays:

- ◆ Management station host name.
- ◆ Selected system host name.
- ◆ DD OS version
- ◆ Selected system model number.
- ◆ User name and role for the current logged in user.
- ◆ **Log Out** icon. Click to log out of the current session.
- ◆ **Refresh** icon. Click to refresh the DD System Manager display.
- ◆ **Help** icon. Click to view the top-level online help. See [Help Buttons and Menus on page 27](#) for details.

## Information Panel

The Information panel displays information about the selected item in the Navigation panel (either the DD Network or a selected system).

At the top of the Information panel is a bar with information about the system or group selected in the Navigation panel, such as the full system name, uptime, model number, and the DD OS version number.

### Tab Bar

When you select a single system in the Navigation panel, the Tab bar appears. Its tabs provide access to the configuration and monitoring tools for the system. Many of the these tabs have their own set of tabs. The top-level set of tabs are as follows:

- ◆ Status—displays important information about the system. Subtabs include Summary, Alerts, Active Users, and Stats.
- ◆ Data Management—contains subtabs for File System, MTree, Quota, Snapshots, CIFS, NFS, VTL, and DD Boost.
- ◆ Replication—provides data replication monitoring and management tools.
- ◆ Hardware—provides tabs for monitoring health and statistics of hardware for Storage, Network, Fibre Channel, and Chassis.
- ◆ System Settings—provides tabs for Licenses, Access Management, and General Configuration.
- ◆ Maintenance—provides tabs for System, Support, Logs, and IPMI.

### Working with Table View Options

Many of the views with tables of items contain controls for filtering, navigating, and sorting the information in the table.

How to use common table controls:

- ◆ Click the diamond icon in a column heading to reverse the sort order of items in the column.
- ◆ Click the ◀ and ▶ arrows at the bottom right of the view to move forward or backward through the pages. To skip to the beginning of a sequence of pages, click ◀◀. To skip to the end, click ▶▶.
- ◆ Use the scroll bar to view all items in a table.
- ◆ Enter text in the **Filter By** box to search for or prioritize the listing of those items.
- ◆ Click **Update** to refresh the list.
- ◆ Click **Reset** to return to the default listing.

### More Tasks Menu

Some pages have a More Tasks menu at the top of the view that contains commands related to the current view.

### Help Buttons and Menus

Help is available globally and from individual panes:

- ◆ Help icon—This icon is always visible on the right side of the DD System Manager banner. Click to display online help, which is derived from this guide. The Help window includes navigation icons that show the guide contents, index, favorites, search field, and an option to send to printer. Use the directional arrows to page through the sections of the book.

- ◆ Context-sensitive help—Most windows and individual views have a Help icon, represented by a question mark (?). Click the icon to open online help for the current window. The tools described for the Help icon above are also available.

## View End User License Agreement (EULA)

To view the End User License Agreement at any time, select EULA from the **More Tasks** menu on the Maintenance page.

## Using the Configuration Wizard

The Configuration Wizard guides you through the initial configuration of your system. There are two wizards, one uses DD System Manager and the other uses the CLI. See the *EMC DD OS Initial Configuration Guide* for more information.

## Using the CLI

The *EMC DD OS Command Reference Guide* provides information for using the commands to accomplish administration tasks.

Online help is available and provides the complete syntax for each command. To display CLI help, type the `help` command.

Any Data Domain system command that accepts a list, such as a list of IP addresses, accepts entries separated by commas, by spaces, or both.

The Tab key can be used:

- ◆ to complete a command entry when that entry is unique. Tab completion is supported for all keywords. For example, entering `sys` Tab `sh` Tab `st` Tab displays the command `system show stats`.
- ◆ to show the next available option, if you do not enter any characters before pressing the Tab key.
- ◆ to show partial matched tokens or to complete a unique entry, if you enter characters before pressing the Tab key.

## Logging into the System Using the CLI

After the initial configuration, use the SSH or Telnet (if enabled) utilities to access the system remotely and to use the CLI.

- ◆ From a serial console, use the communication settings: 9600 baud, 8 data bits, no parity, and 1 stop bit.
- ◆ From a directly attached keyboard and monitor, log into the Data Domain system at the login prompt.
- ◆ From a remote machine over an Ethernet connection, use SSH or Telnet to connect to the Data Domain system.

For SSH, use the following command (with the hostname you chose for the Data Domain system at initial configuration) and provide the `sysadmin` password:

```
# ssh -l sysadmin hostname
```

```
Data Domain OS 5.4.0.0-19899
```

```
Password:
```

## Finding Online Help for Commands

There are several ways to find help for commands:

- ◆ To list Data Domain system commands, enter a question mark (?), or type the command `help` at the prompt.
- ◆ To list the options for a command, enter the command with no options at the prompt.
- ◆ To find a keyword used in a command option when you do not remember which command to use, enter a question mark (?) or the `help` command followed by the keyword.  
For example, the question mark followed by the keyword `password` displays all Data Domain system command options that include `password`. If the keyword matches a command, such as `net`, an explanation of that command appears.
- ◆ To display a detailed explanation of a command, enter the `help` command followed by the command's name.
- ◆ Use these keyboard shortcuts:
  - Up and down arrow keys to move through a displayed command.
  - The `q` key to quit/exit.
  - A slash character (/) followed by a pattern to use as search criteria. Matches are highlighted.



# CHAPTER 3

## Managing Data Domain Systems

This chapter includes:

◆ About Managing Data Domain Systems.....	32
◆ Managing System Availability.....	32
◆ Working with Upgrade Images.....	35
◆ Managing System Licenses.....	37
◆ Managing System Storage.....	38
◆ Managing Network Connections.....	43
◆ Managing Access to the System.....	63
◆ Managing General Configuration Settings.....	80
◆ Managing Reporting and Logging.....	90
◆ Managing Remote System Power with IPMI.....	97

## About Managing Data Domain Systems

---

### Note

When processing a heavy load, a system might be less responsive than normal. In this case, management commands issued from either DD System Manager or the CLI might take longer to complete. When the duration exceeds allowed limits, a timeout error is returned, even if the operation completed.

---

The DD System Manager controls individual systems, which are listed in the Navigation panel and are referred to as managed systems.

---

### Note

- ◆ A managed system should be managed by one management system at one time.
- ◆ If you are an admin on the management system you become a global admin, which means that you can configure and monitor all managed systems.
- ◆ If you are a user on the management system you become a global user, which means that you can monitor all managed systems.

---

This table recommends the maximum number of systems and user sessions that can be managed by DD System Manager:

**Table 3** Maximum Number of Systems and Users Managed by DD System Manager

System Model	Maximum Active Users	Maximum Logged In Users	Maximum Systems
4 GB models <sup>a</sup>	5	10	8
8 GB models <sup>b</sup>	10	15	12
16 GB and greater models <sup>c</sup>	10	20	20

a. Includes DD120, DD140, DD510, and DD530

b. Includes DD565, DD610, and DD630

c. Includes DD580, DD660, DD670, DD690, DD860, DD880, DD890, DD990, DD580g, DD690g, and DD880g

## Managing System Availability

The topics in this section include how to:

- ◆ [Adding a System to DD System Manager on page 33](#)
- ◆ [Removing a System from DD System Manager on page 33](#)
- ◆ [Rebooting a System on page 34](#)
- ◆ [Powering a Data Domain System On or Off on page 34](#)



## Adding a System to DD System Manager

---

### Note

Make sure the system being added is running a DD OS version that is compatible with the DD System Manager. DD System Manager supports the management of systems running the previous version, the current version, and the next version when it becomes available. To support replication, DD System Manager supports the addition of systems running the previous two versions, the current version and the next two versions as they become available. For Release 5.4, DD System Manager supports management of versions 5.2 to 5.5, and it supports the addition of systems for replication for DD OS Version 5.1 to 5.6.

---

### Procedure

1. Click the Add icon (+) on the Navigation panel.
  2. In the Add System dialog box, enter the hostname or IP address of the system to be added in the System Name box.
- 

### Note

A system should be added to and managed by only one DD System Manager.

---

3. In Administration Credentials, enter the sysadmin user name in the **User Name** box, followed by the password.
  4. Optionally, click **Advanced** to enter a proxy IP address (or system name) of a system that cannot be reached directly. If configured, enter a custom port instead of the default port 3009.
  5. Click **OK**.
- 

### Note

If the system is unreachable after adding it to DD System Manager, ensure the following:

- If a hostname (either a fully-qualified domain name (FQDN) or non-FQDN) is entered, make sure it is resolvable on the managed system. Either configure a domain name for the managed system, ensure a DNS entry for the system exists, or ensure a IP address to hostname mapping is defined).
  - If an IP address or hostname is entered, ensure there is a route from the managing system to the system being added.
- 
6. If the system certificate has not been verified, the Verify Certificate dialog box shows details about the certificate. Check the system credentials. Click **OK** if you trust the certificate, or click **Cancel**.

## Removing a System from DD System Manager

This topic describe how to remove a system (other than the system that hosts DD System Manager) from DD System Manager.

---

**Note**

Removing a system removes it from the DD Network list. It does not delete any replication context configured to or from that system.

---

**Procedure**

1. Click the **X** (remove) icon on the navigation panel.
  2. In the Remove System(s) dialog box:
    - To remove all systems, select **System**.
    - To remove one or more systems, select the system.
  3. Click **OK**.
- 

**Note**

If only the DD System Manager host system is present, clicking the X icon results in a message stating that no removable systems are found on DD System Manager.

---

## Rebooting a System

Some configuration changes, such as changing the time zone, require that you reboot the system.

**Procedure**

1. Select a system in the navigation panel.
2. Select **Maintenance** › **System**.
3. From the More Tasks menu, select **Reboot System**.
4. Click **OK** to confirm.

## Powering a Data Domain System On or Off

This topic describes how to power on or power off a Data Domain system.

**Procedure**

1. Power on any expansion shelves before powering on the Data Domain controller. The ES30 powers on when plugged in. Wait approximately three minutes after all expansion shelves are turned on.
- 

**Note**

The controller is the chassis and any internal storage. A *Data Domain system* refers to the controller and any optional external storage.

---

2. Plug in the power cord for your controller, and if the controller has a power button, press the controller's power button (as shown in the *Installation and Setup Guide* for your Data Domain system).
3. To shut down power to a Data Domain system, use the `system poweroff` CLI command.

This command automatically performs an orderly shut down of DD OS processes and is available to administrative users only.

**Results****⚠ CAUTION**

Do not use the chassis power switch to power off the system. Doing so prevents remote power control using IPMI. Use the `system poweroff` command instead. The `system poweroff` command shuts down the system and turns off the power.

---

**Note**

The IMPI Remote System Power Down feature does not perform an orderly shutdown of the DD OS. Use this feature only if the `system poweroff` command is unsuccessful.

---

## Working with Upgrade Images

DD System Manager provides a link to the Data Domain Support Web site. When you upgrade your DD OS, download the image from the Support site to a local system. Then upload the image to the system to be upgraded.

This section covers the following topics:

- ◆ [Viewing the Upgrade Package List on page 35](#)
- ◆ [Obtaining Upgrade Packages on page 35](#)
- ◆ [Upgrading a Data Domain System on page 36](#)
- ◆ [Removing an Upgrade Image on page 37](#)

## Viewing the Upgrade Package List

The **Maintenance** > **System** view provides a list of upgrade images (.rpm files) currently stored on the Data Domain controller. The Upgrade Packages Available on the Data Domain system list displays the following information for each image:

**Table 4** Upgrade Package List Items

Item	Description
File Name	The name of the .rpm file stored on the system.
Size	The size of the .rpm file.
Last Modified	The date the .rpm file was last changed (for example, if the image was modified).

## Obtaining Upgrade Packages

This topic describes how to connect to the Data Domain Support Web site and obtain an upgrade package.

**Procedure**

1. Expand **DD Network** in the navigation panel, and select a system.
2. Select **Maintenance** > **System**.
3. In the Upgrade Packages Available on the Data Domain System list, select Upload Upgrade Package.

A maximum of five upgrade packages can be uploaded using DD System Manager. To continue with this procedure, remove the excess images (see [Removing an Upgrade Image on page 37](#)).

There are no restrictions, other than space limitations, if you use FTP or NFS to copy an upgrade package to a system. FTP is disabled by default. To use NFS, `/ddvar` needs to be exported and mounted from an external host).

4. To obtain an upgrade package from the Data Domain Support site, click the Data Domain Support Portal link. Log into the site and use the Download Software link to navigate to the image recommended for your system by Support personnel. Save the upgrade image to a local computer with access to the system.
5. In the Upload Upgrade Package dialog box, click **Browse** to open the File Upload dialog box. Navigate to the system with the file, select the file, and click **Open**.
6. Click **OK**.

Progress on the status of the upload appears. Upon successful completion of the upload, the new .rpm is listed in Upgrade Packages Available on the Data Domain list.

## Upgrading a Data Domain System

This topic describes how to perform a system upgrade with an image that resides on the Data Domain system.

---

### Note

System upgrade files use the .rpm file extension. This topic assumes that you are updating only DD OS. If you make hardware changes, such as adding, swapping, or moving interface cards, you must update the DD OS configuration to correspond with the hardware changes.

---

### Procedure

1. After reading the Release Notes for instructions for this DD OS upgrade, log into the system where the upgrade is to be performed.

---

### Note

For most releases, upgrades are permitted from up to two prior major release versions. For Release 5.4, upgrades are permitted from up to three prior major release versions.

You cannot select a system from the DD Network list while logged in on another system.

- 
2. From the Upgrade Packages Available on this Data Domain System list, select the image to use for the upgrade.
  3. Click **Perform System Upgrade**.
  4. In the Upgrade System dialog box, verify the version of the target image (the upgrade image). Click **OK**.
  5. After installation completes, the system automatically reboots unless the DD OS installation also included firmware upgrades, in which case the system powers down.
    - If the system powers down, you must power cycle (AC power) the system. Unplug all of the power cables for 30 seconds and then plug them back in. The system powers on and reboots.
    - If your system does not automatically power on and has a power button on the front panel, press the button.

6. Log into the system.

## Removing an Upgrade Image

This topic describes how to remove a system upgrade image that resides on the Data Domain system.

### Procedure

1. From the Upgrade Packages Available on this Data Domain System list, select the image to remove. One image can be removed at a time.
2. Click **Remove Upgrade Package**.

## Managing System Licenses

Optional features for the Data Domain system are licensed and must be purchased separately. See [Licensed Features on page 18](#) for a list of licensed products. The following procedures describe how to display and enable licenses.

- ◆ [Displaying Licenses on page 37](#)
- ◆ [Adding Licenses on page 37](#)
- ◆ [Removing Licenses on page 38](#)

## Displaying Licenses

### Procedure

1. In the navigation panel, expand **DD Network** and select a managed system.
2. Select **System Settings** > **Licenses**.

The Feature Licenses list displays license keys and feature names.

## Adding Licenses

This topic describes how to add a feature license with DD Service Manager.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings** > **Licenses**.

The Feature Licenses list displays license keys and feature names.

3. Click **Add Licenses**.
4. In the Add Licenses Key dialog box, type or paste one or more license keys in the License Key box. Type each key on its own line, or separate each key by a space or comma (and they will be automatically placed on a new line).
5. Click **Add**.

The added licenses display in the Added license list.

Any errors are listed in the error license list. Select a license with an error to edit it. Select **Retry Failed License(s)** to retry the key. Otherwise, click **Done** to ignore the errors and return to the Feature Licenses list.

## Removing Licenses

This topic describes how to remove one or more feature licenses with DD Service Manager.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings** > **Licenses**.  
The Feature Licenses list displays license keys and feature names.
3. In the Feature Licenses list, select one or more licenses to remove and click **Delete Selected Licenses**.
4. In the Warning dialog box verify the license(s) to delete and click **OK**.  
The licenses are removed from the license list.

## Managing System Storage

The Storage view organizes the Data Domain system storage so that you can view disks by usage type (Active, Archive, Failed, and so on), operational status, and location. The system is automatically scanned and inventoried so that the status and inventory are shown for all enclosures, disks, and RAID groups.

The following topics are covered:

- ◆ [Viewing System Storage Information on page 38](#)
- ◆ [Physically Locating a Disk on page 42](#)
- ◆ [Configuring Storage on page 42](#)

## Viewing System Storage Information

This topic describes how to view system storage information with DD System Manager.

### Procedure

1. Select a system in the navigation panel.
2. Select **Hardware** > **Storage**.  
The Storage view Status area shows the current status of the storage, such as Operational or Non-Operational, and any active alerts, which can be clicked to view alert details.

### Results

The Status area presents the following information.

**Table 5** Storage System Status

Item	Description
Status	Storage system status can be one of the following: <ul style="list-style-type: none"> <li>• Normal—System operational (green). All disks in the system are in good condition.</li> <li>• Warning—System operational (yellow). The system is operational, but there are problems that need to be</li> </ul>

**Table 5** Storage System Status (continued)

Item	Description
	<p>corrected. Warnings may result from a degraded RAID group, presence of foreign storage, or failed or absent disks.</p> <ul style="list-style-type: none"> <li>• Error—System non-operational (red). The system is not operational.</li> </ul>
Operational Drives	<p>Count of drives operating normally:</p> <ul style="list-style-type: none"> <li>• Total—Total number of drives operating.</li> <li>• In-Use—Number of drives the system is using.</li> <li>• Spare—Number of spare drives (that can be activated if an in-use disk fails).</li> </ul>
Non-Operational Drives	<p>Count of drives that are not operating normally.</p> <ul style="list-style-type: none"> <li>• Total—Total number of non-operational drives.</li> <li>• Failed—Number of known failed drives.</li> <li>• Absent—Number of slots without drives.</li> <li>• Foreign—Number of foreign or unsupported drives.</li> <li>• Unknown—Number of new disks in a shelf.</li> </ul>

Below the Status area are tabs that organize how the storage inventory is presented.

## Storage Overview

The Overview area displays information for all disks in the selected Data Domain system organized by type. The categories that display are dependent on the type of storage configuration in use. The Overview section lists the storage that is found, and can include the following sections. You can expand each of these sections to display detailed information:

- ◆ **Active Tier**  
Disks in the Active Tier are currently marked as usable by the Data Domain file system. Sections are organized by Disks in Use and Disks Not in Use.
- ◆ **Retention Tier**  
If the optional EMC Data Domain Extended Retention (formerly DD Archiver) license is installed, this section shows the disks that are configured for Extended Retention storage. Sections are organized by Disks in Use and Disks Not in Use. For more information, see the *EMC Data Domain Extended Retention Administration Guide*.
- ◆ **Usable Disks and Enclosures**  
For systems with optional enclosures, this section shows the disks and enclosures that can be added to the system.
- ◆ **Fail/Foreign/Absent Disks (Excluding Systems Disks)**  
Shows the disks that are in a failed state; these cannot be added to the system Active or Retention tiers.
- ◆ **Systems Disks**  
Shows the disks where the DD OS resides when the Data Domain controller does not contain data storage disks.

Each section heading displays a summary of the storage configured for that section. The Summary shows tallies for the total number of disks, disks in use, spare disks, reconstructing spare disks, available disks, and known disks.

Clicking the plus (+) icon for a section shows information about the status and content of the storage that is present.

Sections with the Disks In Use section show the Disk Group Status tallies and a table with the following information:

**Table 6** Disk Group Status

Item	Description
Disk Group	The name of the disk group that was created by the file system (for example, dg1).
Status	The status of the disk (for example Normal, Warning).
Disks Reconstructing	The disks that are undergoing reconstruction, by disk ID (for example, 1.11).
Total Disks	The total number of usable disks (for example, 14).
Disks	The disk IDs of the usable disks (for example, 2.1-2.14).

Sections with a Disks Not in Use section show the Disks Status tallies and a table with the following information:

**Table 7** Disk Status

Item	Description
Disk	The disk identifier. It can be: <ul style="list-style-type: none"> <li>• The enclosure and disk number (in the form <i>Enclosure.Slot</i>).</li> <li>• A gateway disk (<i>devn</i>).</li> <li>• A LUN.</li> </ul>
Status	The status of the disk, for example In Use, Available, Spare.
Size	The data storage capacity of the disk when used in a Data Domain system. <sup>a</sup>
Manufacturer/Model	The manufacturer’s model designation. The display may include a model ID or RAID type or other information depending on the vendor string sent by the storage array.
Firmware	The firmware level used by the third-party physical-disk storage controller.
Serial Number	The manufacturer’s serial number for the disk.

a. The Data Domain convention for computing disk space defines one gibibyte as 230 bytes, giving a different disk capacity than the manufacturer’s rating.

## Status View

The Status view shows the Disks Status table and the Reconstructing table.

The following table describes the entries in the Disks Status table.



**Table 8** Disks Status

Item	Description
Total	The total number of inventoried disks in the Data Domain system (including enclosures and gateway storage).
In Use	The number of disks currently in use by the file system.
Spare	The number of spare disks (available to replace failed disks).
Spare (reconstructing)	The number of disks that are in the process of data reconstruction (spare disks replacing failed disks).
Available	The number of disks that are available for allocation to an Active or Extended Retention storage tier.
Known	The number of known unallocated disks.
Unknown	The number of unknown unallocated disks.
Failed	The number of failed disks.
Foreign	The number of foreign disks.
Absent	The number of absent disks.

The following table describes the entries in the Reconstructing table.

**Table 9** Disk Reconstruction Status

Item	Description
Disk	Identifies disks that are being reconstructed. Disk labels are of the format <i>enclosure.disk</i> . Enclosure 1 is the Data Domain system, and external shelves start numbering with enclosure 2. For example, the label 3.4 is the fourth disk in the second shelf.
Disk Group	Shows the RAID group (dg#) for the reconstructing disk.
Tier	The name of the tier where the failed disk is being reconstructed.
Time Remaining	The amount of time before the reconstruction is complete.
Percentage Complete	The percentage of reconstruction that has been completed.

When a spare disk is available, the Data Domain file system automatically replaces a failed disk with a spare and begins the reconstruction process to integrate the spare into the RAID disk group. The disk use displays `Spare` and the status becomes `Reconstructing`. Reconstruction is performed on one disk at a time.

## Disks View

You can select how the disks are viewed: All Disks, by tier, or by disk group. The following table describes the entries in the Disks view.

**Table 10** System Disks Status

Item	Description
Disk	The disk identifier, which can be: <ul style="list-style-type: none"> <li>The enclosure and disk number (in the form <i>Enclosure.Slot</i>).</li> </ul>

**Table 10** System Disks Status (continued)

Item	Description
	<ul style="list-style-type: none"> <li>• A gateway disk (dev<math>n</math>).</li> <li>• A LUN.</li> </ul>
Status	The status of the disk (for example In Use, Spare).
Manufacturer/Model	The manufacturer’s model designation. The display may include a model ID or RAID type or other information depending on the vendor string sent by the storage array.
Firmware	The firmware level used by the third-party physical-disk storage controller.
Serial Number	The manufacturer’s serial number for the disk.

## Physically Locating a Disk

This topic describes how to use DD Service Manager to determine the location of a disk within a system.

### Procedure

1. Select the system in the navigation panel.
2. Select **Hardware > Storage > Disks**.
3. Select a disk from the **Disks** table and click **Beacon**.

---

### Note

You can select one disk at a time.  
The Beaconsing Disk dialog box appears, and the LED light on the disk begins flashing.

---

4. Click **Stop** to stop the LED beaconsing.

## Configuring Storage

---

### Note

Additional storage requires the appropriate license or licenses, and the Data Domain system must have enough installed memory to support it. Error messages display if more licenses or memory is needed.

---

### Procedure

1. Expand **DD Network** in the navigation panel, and select a system.
2. Select **Hardware > Storage**.
3. In the Overview tab, click **Configure Storage**.
4. In the Configure Storage dialog box, select the storage to be added from the **Available Storage** list.
5. Select the appropriate Tier Configuration (Archive or Active) from the menu.

---

**Note**

The two bars show the portion of licensed capacity (used and remaining) for each shelf model (ES20 and ES30).

---

6. Select the checkbox for the Shelf to be added.
  7. Click the **Add to Tier** button.
  8. Click **OK** to add the storage.
- 

**Note**

To remove an added shelf, select it in the Tier Configuration list, click **Remove from Configuration**, and click **OK**.

---

## Managing Network Connections

The following topics describe how to manage network interfaces, general network settings, and network routes.

### Configuring Network Interfaces

This section provides an overview of the types of connections, physical and virtual, and how they are used to create VLANs, IP aliases, and bonded interfaces for Data Domain systems.

This section includes the following tasks:

- ◆ [Viewing Interface Information on page 43](#)
- ◆ [Filtering the Interfaces Table on page 45](#)
- ◆ [Physical Interface Names and Limitations on page 46](#)
- ◆ [General Interface Configuration Guidelines on page 46](#)
- ◆ [Configuring Physical Interfaces on page 47](#)
- ◆ [Virtual Interface Configuration Guidelines on page 48](#)
- ◆ [Creating Virtual Interfaces on page 50](#)
- ◆ [Modifying a Virtual Interface on page 53](#)
- ◆ [Configuring a VLAN on page 54](#)
- ◆ [Modifying a VLAN Interface on page 55](#)
- ◆ [Configuring an IP Alias on page 55](#)
- ◆ [Modifying an IP Alias Interface on page 56](#)
- ◆ [Registering a DDNS on page 56](#)
- ◆ [Destroying an Interface on page 57](#)
- ◆ [Viewing an Interface Hierarchy in the Tree View on page 57](#)

### Viewing Interface Information

The Interfaces view allows you to manage and configure virtual interfaces, DHCP, DDNS, and IP addresses, and to display network information and status.

**Note**

The command-line interface (CLI) supports IPv6 for basic Data Domain network and replication commands, but not for backup and Extended Retention (*archive*) commands. CLI commands manage the IPv6 addresses. You can view IPv6 addresses using the DD System Manager, but you cannot manage IPv6 with the DD System Manager.

Collection, directory, and MTree replication are supported over IPv6 networks, which allows you to take advantage of the IPv6 address space. Simultaneous replication over IPv6 and IPv4 networks is also supported. Managed File Replication using DD Boost is not supported on IPv6 networks.

There are some restrictions for interfaces that have IPv6 addresses. For example, the minimum MTU is 1280. If you try to set the MTU lower than 1280 on an interface with an IPv6 address, an error message appears and the interface is removed from service. An IPv6 address can affect an interface even though it is on a VLAN attached to the interface and not directly on the interface.

**Procedure**

1. In the navigation panel, select the system to view or configure.
2. Select **Hardware > Network**.

The Network view appears, displaying the Interfaces, Settings, and Routes tabs.

The Interfaces table shows the following information:

**Table 11** Interfaces

Item	Description
Interface	The name of each interface associated with the selected system. Physical interfaces names are described in <a href="#">Physical Interface Names and Limitations on page 46</a> . Virtual interface names are described in <a href="#">Guidelines for Configuring Virtual Interfaces on page 49</a> .
Enabled	Whether the interface is enabled. <ul style="list-style-type: none"> <li>• Select <b>Yes</b> to enable the interface and connect it to the network.</li> <li>• Select <b>No</b> to disable the interface and disconnect it from the network.</li> </ul>
DHCP	Indicates if the interface is configured with an IP address from a DHCP (Dynamic Host Configuration Protocol) server (Yes/No).
IP Address	IP address associated with the interface. The address used by the network to identify the interface. If the interface is configured through DHCP, an asterisk appears after this value.
Netmask	Netmask associated with the interface. Uses the standard IP network mask format. If the interface is configured through DHCP, an asterisk appears after this value.
Link	Whether the interface currently has a live Ethernet connection (Yes/No).
Additional Info	Additional settings for the interface. For example, the bonding mode.
Intelligent Platform Management Interface (IPMI)	
Yes/No	Indicates if IPMI health and management monitoring is configured for the interface.

**Table 11** Interfaces (continued)

Item	Description
View IPMI Interfaces	Links to the <b>Maintenance &gt; IPMI</b> configuration tab.

3. Select an interface in the table to populate the Interface Details area.

The Interface Details area shows the following information:

**Table 12** Interface Details

Item	Description
Interface Name	Name of the selected interface.
Hardware Address	The MAC address of the selected interface. For example, 00:02:b3:b0:8a:d2.
Cable	Shows whether the interface is Copper or Fiber.  <b>Note</b> Some interfaces must be up before the cable status is valid.
MTU	MTU (Maximum Transfer Unit) value assigned to the interface. See <a href="#">About MTU Size Values on page 48</a> .
Autonegotiate	When this feature displays <code>Enabled</code> , the interface automatically negotiates Speed and Duplex settings. When this feature displays <code>Disabled</code> , then Speed and Duplex values must be set manually.
Duplex	Used in conjunction with the Speed value to set the data transfer protocol. Options are <code>Unknown</code> , <code>Full</code> , <code>Half</code> .
Speed	Used in conjunction with the Duplex value to set the rate of data transfer. Options are <code>Unknown</code> , <code>10 Mb/s</code> , <code>100 Mb/s</code> , <code>1000 Mb/s</code> , <code>10 Gb/s</code> .  <b>Note</b> Auto-negotiated interfaces must be set up before speed, duplex, and supported speed are visible.
Supported Speeds	Lists all of the speeds that the interface can use.

## Filtering the Interfaces Table

The Interfaces table can be filtered by either:

- ◆ Interface Name—Enter an interface name and click `Update` to filter the Interface view.
- ◆ Interface Type— Select an interface type and click `Update` to filter to Interface view. The value `All` displays physical, virtual (Failover and Aggregate), VLAN, and IP Alias interfaces.

### Procedure

1. Enter a value in the **Interface Name** field, or select a value from the **Interface Type** menu.

Filters support wildcards, such as `eth*`, `veth*`, or `eth0*`

2. Click **Update**.
3. To return the interfaces table to the default listing, click **Reset**.

## Physical Interface Names and Limitations

The physical interface names vary on different Data Domain systems and option cards, and some interfaces have limitations:

- ◆ For most systems the physical interface name format is `ethxy`, where `x` is the slot number for an on-board port or an option card and `y` is an alphanumeric string. For example, `eth0a`.
- ◆ For most on-board NIC vertical interfaces, the top interface is named `eth0a` and the bottom interface is `eth0b`.
- ◆ For most on-board NIC horizontal interfaces, the left interface as viewed from the rear, is named `eth0a` and the right is named `eth0b`.
- ◆ DD990 systems have four on-board interfaces: two on the top and two on the bottom. The top-left interface is `eth0a`, the top-right is `eth0b`, the bottom-left is `eth0c`, and the bottom-right is `eth0d`.
- ◆ DD2500 systems have six on-board interfaces. The four on-board 1G Base-T NIC ports are `ethMa` (top left), `ethMb` (top right), `ethMc` (bottom left), and `ethMd` (bottom right). The two on-board 10G Base-T NIC ports are `ethMe` (top) and `ethMf` (bottom).
- ◆ DD4200, DD4500, and DD7200 systems have one on-board Ethernet port, which is `ethMa`.
- ◆ For systems ranging between DD120 and DD990, the physical interface names for IO modules start at the top of the module or at the left side. The first interface is `ethxa`, the next is `ethxb`, the next is `ethxc`, and so forth.
- ◆ The port numbers on the horizontal DD2500 IO modules are labeled in sequence from the end opposite the module handle (left side). The first port is labeled 0 and corresponds to physical interface name `ethxa`, the next is 1/`ethxb`, the next is 2/`ethxc`, and so forth.
- ◆ The port numbers on the vertical DD4200, DD4500, and DD6200 IO modules are labeled in sequence from the end opposite the module handle (bottom). The first port is labeled 0 and corresponds to physical interface name `ethxa`, the next is 1/`ethxb`, the next is 2/`ethxc`, and so forth.

## General Interface Configuration Guidelines

Consider the following guidelines when configuring physical and virtual interfaces:

- ◆ When supporting both backup and replication traffic, EMC recommends using different interfaces for each traffic type so that neither traffic type impacts the other.
- ◆ When replication traffic is expected to be less than 1 Gb/s, EMC recommends against using 10 GbE interfaces for replication traffic because 10 GbE interfaces are optimized for faster traffic.
- ◆ On DD4200, DD4500, and DD7200 systems that use IPMI, EMC recommends that interface `ethMa` be reserved for IPMI traffic and system management traffic (using protocols such as HTTP, Telnet, and SSH). Backup data traffic should be directed to other interfaces. For more information, see [Managing Remote System Power with IPMI on page 97](#).
- ◆ For additional guidelines, see [Guidelines for Configuring Virtual Interfaces on page 49](#).

## Configuring Physical Interfaces

### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware > Network > Interfaces**.
3. Select an interface to configure.
4. Click **Configure**.
5. In the Configure Interface dialog box, determine how the interface IP address is to be set:
  - Use DHCP to assign the IP address—in the IP Settings area, select **Obtain IP Address using DHCP**.  
Setting a physical interface to DHCP automatically enables the interface.
  - Specify IP Settings manually—in the IP Settings area, select **Manually configure IP Address**.  
The **IP Address** and **Netmask** fields become active.
    - a. Enter an IP address.  
The Internet Protocol (IP) address is the numerical label assigned to the interface. For example, 192.168.10.23.
    - b. Enter a netmask address.  
The netmask is the subnet portion of the IP address that is assigned to the interface. If the interface is configured through DHCP, an asterisk appears after this value.  
The format is typically 255.255.255.0. If you do not specify a netmask, the selected system uses the netmask format determined by the TCP/IP address class (A,B,C) you are using.
6. Specify Speed/Duplex settings.  
The combination of speed and duplex settings define the rate of data transfer through the interface. Select one of these options:
  - **Autonegotiate Speed/Duplex** — Select this option to allow the network interface card to autonegotiate the line speed and duplex setting for an interface. Autonegotiation is not supported on the following DD2500, DD4200, DD4500, and DD7200 IO modules:
    - Dual Port 10GbE SR Optical with LC connectors (using SFPs)
    - Dual Port 10GbE Direct Attach Copper (SFP+ cables)
    - Quad port 2 port 10GbE Copper (RJ45) / 2 port 10GbE SR Optical
  - **Manually configure Speed/Duplex** — Select this option to manually set an interface data transfer rate. Select the speed and duplex from the menus.
    - Duplex options are half-duplex, full-duplex, and unknown.
    - Speed options listed are limited to the capabilities of the hardware device. Options are 10 Mb, 100 Mb, 1000 Mb (1 Gb), 10 Gb, and unknown. The 10G Base-T hardware supports only the 100 Mb, 1000 Mb and 10 Gb settings.
    - Half-duplex is only available for 10 Mb and 100 Mb speeds.
    - 1000 Mb and 10 Gb line speeds require full-duplex.
    - Optical interfaces require the Autonegotiate option.

- On DD2500, DD4200, DD4500, and DD7200 10GbE IO modules, copper interfaces support only the 10 Gb speed setting.
  - The default setting for 10G Base-T interfaces is Autonegotiate Speed/Duplex. If you manually set the speed to 1000 Mb or 10 Gb, you must set the Duplex setting to Full.
7. Specify the MTU (Maximum Transfer Unit) size for the physical (Ethernet) interface. See [About MTU Size Values on page 48](#).

Do the following:

- Click the **Default** button to return the setting to the default value.
  - Ensure that all of your network components support the size set with this option.
8. Optionally, select **Dynamic DNS Registration**.

Dynamic DNS (DDNS) is the protocol that allows machines on a network to communicate with and register their IP address on a Domain Name System (DNS) server.

The DDNS must be registered to enable this option.

---

**Note**

This option disables DHCP for this interface.

9. Click **Next**.

The Configure Interface Settings summary page appears. The values listed reflect the new system and interface state, which are applied after you click Finish.

10. Click **Finish** and **OK**.

## About MTU Size Values

**NOTICE**

Incorrect MTU size can affect the system's network performance.

---

Supported values for setting the maximum Transfer Unit (MTU) size for the physical (Ethernet) interface range from 350 to 9000. For 100 Base-T and gigabit networks, 1500 is the standard default.

**Note**

Although 9000 is the maximum size, to be backwardly compatible, the DD OS accepts up to 9014, but sets the size to 9000 when it is greater than 9000 and less than or equal to 9014.

The minimum MTU for IPv6 interfaces is 1280. The interface fails if you try to set the MTU lower than 1280.

---

## Virtual Interface Configuration Guidelines

Before you create a virtual interface, become familiar with the applicable guidelines:

- ◆ [Guidelines for Configuring Virtual Interfaces on page 49](#)
- ◆ [Guidelines for Configuring a Virtual Interface for Failover on page 49](#)
- ◆ [Guidelines for Configuring a Virtual Interface for Link Aggregation on page 50](#)



## Guidelines for Configuring Virtual Interfaces

The following considerations apply to both failover and aggregate virtual interfaces. When you create a virtual interface:

- ◆ The *virtual-name* must be in the form `vethx` where *x* is a number. The recommended maximum number is 99 because of name size limitations.
- ◆ You can create as many virtual interfaces as there are physical interfaces.
- ◆ Each interface used in a virtual interface must first be disabled. An interface that is part of a virtual interface is seen as disabled for other network configuration options.
- ◆ After a virtual interface has been destroyed, the physical interfaces associated with it remain disabled. You must manually re-enable the physical interfaces.
- ◆ The number and type of cards installed determines the number of Ethernet ports available.
- ◆ Each physical interface can belong to one virtual interface.
- ◆ A system can have multiple mixed failover and aggregation virtual interfaces, subject to the restrictions above.
- ◆ Virtual interfaces must be created from identical physical interfaces. For example, all copper, all optical, all 1 Gb, or all 10 Gb. However, 1 Gb interfaces support bonding a mix of copper and optical interfaces. This applies to virtual interfaces across different cards with identical physical interfaces, except for Chelsio cards. For Chelsio cards, only failover is supported, and that is only across interfaces on the same card.
- ◆ Failover and aggregate links improve network performance and resiliency by using two or more network interfaces in parallel, thus increasing the link speed and reliability over that of a single interface.
- ◆ Remove functionality is available using the **Configure** button. Click a virtual interface in the list of interfaces on the Interfaces tab and click **Configure**. From the list of interfaces in the dialog box, clear the checkbox for the interface to remove it from bonding (failover or aggregate), and click **Next**.
- ◆ On DD4200, DD4500, and DD6200 systems, the ethMa interface does not support failover or link aggregation.

## Guidelines for Configuring a Virtual Interface for Failover

Ethernet failover provides improved network stability and performance. A configurable Down Delay failover option allows you to configure a failover delay in 900 millisecond intervals. The failover delay guards against multiple failovers when a network is unstable.

The failover-enabled virtual interface represents a group of secondary interfaces, one of which can be specified as the primary. The system makes the primary interface the active interface whenever the primary interface is operational.

While planning, consider the following supported guidelines:

- ◆ A primary interface must be part of the failover. If a primary interface removal is attempted from a failover, an error message appears.
- ◆ When a primary interface is used in a failover configuration, it must be explicitly specified and must also be a bonded interface to the virtual interface. If the primary interface goes down and multiple interfaces are still available, the next interface is randomly selected.
- ◆ All interfaces in a virtual interface must be on the same physical network. Network switches used by a virtual interface must be on the same physical network.

- ◆ The recommended number of physical interfaces for failover is greater than one. You can, however, configure one primary interface and one or more failover interfaces, except with the following:
  - 10 Gb CX4 Ethernet card, which are restricted to one primary interface and one failover interface from the same card, and
  - 10 Gb single-port optical Ethernet cards, which cannot be used.
- ◆ On DD4200, DD4500, and DD7200 systems, the ethMA interface does not support link failover.

### Guidelines for Configuring a Virtual Interface for Link Aggregation

Link aggregation provides improved network performance and resiliency by using one or more network interfaces in parallel, thus increasing the link speed and reliability over that of a single interface. For example, you might enable link aggregation on virtual interface *veth1* to physical interfaces *eth1* and *eth2* in mode LACP (Link Aggregation Control Protocol) and hash XOR-L2L3.

When planning interface link aggregation, consider the following:

- ◆ Changes to disabled Ethernet interfaces flush the routing table. It is recommended that you make interface changes only during scheduled maintenance downtime. Afterwards, reconfigure the routing rules and gateways.
- ◆ Enable aggregation on an existing virtual interface by specifying the physical interfaces and mode and giving it an IP address.
- ◆ DD2500 on-board 10G Base-T interfaces ethMe and ethMf do not support link aggregation.
- ◆ On DD4200, DD4500, and DD7200 systems, the ethMA interface does not support link aggregation.

## Creating Virtual Interfaces

The following sections describe how to create virtual interfaces:

- ◆ [Creating a Virtual Interface for Failover on page 50](#)
- ◆ [Creating a Virtual Interface for Link Aggregation on page 52](#)

### Creating a Virtual Interface for Failover

This topic describes how to create a virtual interface for failover.

#### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware** > **Network** > **Interfaces**.
3. In the **Interfaces** table, disable the physical interface to which the virtual interface is to be added by clicking **No** in the **Enabled** column.
4. From the **Create** menu, select **Virtual Interface**.
5. In the Create Virtual Interface dialog box, specify a virtual interface name in the **veth** box.

Enter a virtual interface name in the form *veth<sub>x</sub>*, where *x* is a unique ID (typically one or two digits). A typical full virtual interface name with VLAN and IP Alias is *veth56.3999:199*. The maximum length of the full name is 15 characters. Special characters are not allowed. Numbers must be between 0 and 4094, inclusively.

6. Select **Failover** from the **Bonding Type** menu.

7. In the Select an interface to add to the failover configuration, select the corresponding to the interface and click **Next**. Virtual aggregate interfaces can be used for failover.

The Create virtual interface *veth\_name* dialog appears.

8. Enter an IP address, or enter 0 to specify no IP address.

The Internet Protocol (IP) address is the numerical label assigned to the interface. For example, 192.168.10.23.

9. Enter a netmask address.

The netmask is the subnet portion of the IP address that is assigned to the interface.

The format is typically 255.255.255.0. If you do not specify a netmask, the selected system uses the default netmask defined by the TCP/IP address class (A,B,C) you are using.

10. Specify the Speed/Duplex options.

The combination of speed and duplex settings defines the rate of data transfer through the interface. Select either:

- **Autonegotiate Speed/Duplex** to allow the network interface card to autonegotiate the line speed and duplex setting for an interface.
- **Manually configure Speed/Duplex** to manually set an interface data-transfer rate.
  - Duplex options are either half duplex or full duplex.
  - Speed options listed are limited to the capabilities of the hardware device. Options are 10 Base-T, 100 Base-T, 1000 Base-T (Gigabit), and 10,000 (10 Gb).
  - Half-duplex is available for 10 Base-T and 100 Base-T speeds only.
  - 1000 and 10000 line speeds require full-duplex.
  - Optical interfaces require the Autonegotiate option.
  - The copper interface default is 10 Gb. If a copper interface is set to 1000 or 10000 line speed, the duplex must be full-duplex.

11. Specify MTU settings. See [About MTU Size Values on page 48](#).

Do the following:

- Click the **Default** button to return the setting to the default value.
- Ensure that all of your network path components support the size set with this option.

12. Optionally, select Dynamic DNS Registration option.

Dynamic DNS (DDNS) is the protocol that allows machines on a network to communicate with and register their IP address on a Domain Name System (DNS) server.

The DDNS must be registered to enable this option. See [Registering a DDNS on page 56](#) for additional information.

---

**Note**

This option disables DHCP for this interface.

13. Click **Next**.

The Configure Interface Settings summary page appears. The values listed reflect the new system and interface state.

14. Complete the Interface, click **Finish** and **OK**.

## Creating a Virtual Interface for Link Aggregation

This topic describes how to create a virtual interface for link aggregation.

### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware** > **Network** > **Interfaces**.
3. In the Interfaces table, disable the physical interface where the virtual interface is to be added by clicking **No** in the **Enabled** column.
4. From the **Create** menu, select **Virtual Interface**.
5. In the Create Virtual Interface dialog box, specify a virtual interface name in the **veth** box.

Enter a virtual interface name in the form *vethx*, where *x* is a unique ID (typically one or two digits). A typical full virtual interface name with VLAN and IP Alias is *veth56.3999:199*. The maximum length of the full name is 15 characters. Special characters are not allowed. Numbers must be between 0 and 4094, inclusively.

6. Select **Aggregate** from the **Bonding Type** menu.

---

### Note

Registry settings can be different from the bonding configuration. When interfaces are added to the virtual interface, the information is not sent to the bonding module until the virtual interface is given an IP address and brought up. Until that time the registry and the bonding driver configuration are different.

---

7. From the **General** tab, specify the **Bonding Mode**.

Specify the mode that is compatible with the requirements of the system to which the interfaces are directly attached. Available modes are as follows:

- **Round-robin**  
Transmit packets in sequential order from the first available link through the last in the aggregated group.
- **Balanced**  
Data is sent over interfaces as determined by the hash method selected. This requires the associated interfaces on the switch to be grouped into an Ether channel (trunk) and given a hash via the Load Balance parameter.
- **LACP**  
Link Aggregation Control Protocol is similar to Balanced, except that it has a control protocol that communicates to the other end and coordinates which links within the bond are available to use. LACP provides a kind of heartbeat failover and must be configured at both ends of the link.

8. If you selected Balanced or LACP mode, specify a bonding hash type.

From the **General** tab, select from the **Bonding Hash** menu.

Options are: XOR-L2, XOR-L2L3, or XOR-L3L4.

XOR-L2 transmits through a bonded interface with an XOR hash of Layer 2 (inbound and outbound MAC addresses).

XOR-L2L3 transmits through a bonded interface with an XOR hash of Layer 2 (inbound and outbound MAC addresses) and Layer 3 (inbound and outbound IP addresses).

XOR-L3L4 transmits through a bonded interface with an XOR hash of Layer 3 (inbound and outbound IP addresses) and Layer 4 (inbound and outbound ports).

9. To select an interface to add to the aggregate configuration, select the checkbox that corresponds to the interface, and then click **Next**.

The Create virtual interface *veth\_name* dialog appears.

10. Enter an IP address, or enter 0 to specify no IP address.

The Internet Protocol (IP) address is the numerical label assigned to the interface. For example, 192.168.10.23.

11. Enter a netmask address.

The netmask is the subnet portion of the IP address that is assigned to the interface.

The format is typically 255.255.255.0. If you do not specify a netmask, the system uses the default netmask determined by the TCP/IP address class (A, B, C) you are using.

12. Specify Speed/Duplex options.

The combination of speed and duplex settings define the rate of data transfer through the interface. Select either:

- **Autonegotiate Speed/Duplex**  
Select this option to allow the network interface card to autonegotiate the line speed and duplex setting for an interface.
- **Manually configure Speed/Duplex**  
Select this option to manually set an interface data transfer rate.
  - Duplex options are half-duplex or full-duplex.
  - Speed options listed are limited to the capabilities of the hardware device. Options are 10 Base-T, 100 Base-T, 1000 Base-T (Gigabit), and 10,000 (10 Gb).
  - Half-duplex is only available for 10 Base-T and 100 Base-T speeds.
  - 1000 and 10000 line speeds require full-duplex.
  - Optical interfaces require the Autonegotiate option.
  - The 10 GbE copper NIC default is 10 Gb. If a copper interface is set to 1000 or 10000 line speed, duplex must be full-duplex.

13. Specify MTU Settings. See [About MTU Size Values on page 48](#).

Do the following:

- Click the **Default** button to return the setting to the default value.
- Ensure that all of your network components support the size set with this option.

14. Optionally, select Dynamic DNS Registration option.

Dynamic DNS (DDNS) is the protocol that allows machines on a network to communicate with, and register their IP address on, a Domain Name System (DNS) server.

The DDNS must be registered to enable this option. See [Registering a DDNS on page 56](#) for additional information.

15. Click **Next**.

The Configure Interface Settings summary page appears. The values listed reflect the new system and interface state.

16. Click **Finish** and **OK**.

## Modifying a Virtual Interface

This topic describes how to modify the settings for an existing virtual interface.

**Procedure**

1. In the navigation panel, select the system to configure.
2. Select **Hardware > Network > Interfaces**.
3. In the Interfaces column, select the interface and disable the virtual interface by clicking **No** in the **Enabled** column. Click **OK** in the warning dialog box.
4. In the **Interfaces** column, select the interface and click **Configure**.
5. In the **Configure Virtual Interface** dialog box, change the settings that are described in the procedures [Creating a Virtual Interface for Failover on page 50](#) or [Creating a Virtual Interface for Link Aggregation on page 52](#).
6. Click **Next** and **Finish**.

**Configuring a VLAN**

Create a new VLAN interface from either a physical interface or a virtual interface. The recommended total number is 80. You can create up to 100 interfaces (minus the number of aliases, physical and virtual interfaces) before the system prevents you from creating any more.

**Procedure**

1. In the navigation panel, select the system to configure.
2. Select **Hardware > Network > Interfaces**.
3. In the interfaces table, select the interface to which you want to add the VLAN.  
The interface you select must have an IP address before you can add a VLAN.
4. Click **Create** and select the **VLAN** option.
5. In the Create VLAN dialog box, specify a VLAN ID by entering a number in the **ID** field.  
The range of a VLAN ID is between 1 and 4094 inclusive.
6. Enter an IP address, or enter 0 to specify no IP address.  
The Internet Protocol (IP) address is the numerical label assigned to the interface. For example, 192.168.10.23.
7. Enter a netmask address.  
The netmask is the subnet portion of the IP address that is assigned to the interface. The format is typically 255.255.255.0. If you do not specify a netmask, the selected system uses the default netmask determined by the TCP/IP address class (A, B, or C) you are using.
8. Specify MTU Settings. See [About MTU Size Values on page 48](#).  
The VLAN MTU must be less than or equal to the MTU defined for the physical or virtual interface to which it is assigned. If the MTU defined for the supporting physical or virtual interface is reduced below the configured VLAN value, the VLAN value is automatically reduced to match the supporting interface. If the MTU value for the supporting interface is increased above the configured VLAN value, the VLAN value is unchanged.  
Do the following:
  - Click **Default** to return the setting to the default value.
  - Specify a specific MTU size. DD System Manager does not accept an MTU size that is larger than that defined for the physical or virtual interface to which the VLAN is assigned.

9. Specify Dynamic DNS Registration option.

Dynamic DNS (DDNS) is the protocol that allows machines on a network to communicate with, and register their IP address on, a Domain Name System (DNS) server.

The DDNS must be registered to enable this option. See [Registering a DDNS on page 56](#) for additional information.

10. Click **Next**.

The **Configure Interface Settings** summary page appears. The values listed reflect the new system and interface state.

11. Click **Finish** and **OK**.

## Modifying a VLAN Interface

This topic describes how to modify the settings for an existing VLAN interface.

### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware** > **Network** > **Interfaces**.
3. In the **Interfaces** column, select the checkbox of the interface and disable the VLAN interface by clicking **No** in the **Enabled** column. Click **OK** in the warning dialog box.
4. In the Interfaces column, select the checkbox of the interface and click **Configure**.
5. In the **Configure VLAN Interface** dialog box, change the settings that are described in the procedures [Configuring a VLAN on page 54](#).
6. Click **Next** and **Finish**.

## Configuring an IP Alias

Create a new IP alias interface from a physical interface, a virtual interface, or a VLAN.

The recommended total number of IP aliases, VLAN, physical, and virtual interfaces that can exist on the system is 80. Although up to 100 interfaces are supported, as the maximum number is approached, you might notice slowness in the display.

### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware** > **Network** > **Interfaces**.
3. Click the **Create** menu and select the IP Alias option.  
The Create IP Alias dialog box appears.
4. Specify an IP alias ID by entering a number in the **eth0a** field.  
Requirements are 1 to 4094 inclusive.
5. Enter an IP address.

The Internet Protocol (IP) address is the numerical label assigned to the interface. For example, 192.168.10.23.

6. Enter a netmask address.

The netmask is the subnet portion of the IP address that is assigned to the interface.

The format is typically 255.255.255.0. If you do not specify a netmask, the selected system uses the default netmask determined by the TCP/IP address class (A, B, or C) you are using.

7. Specify Dynamic DNS Registration option.

Dynamic DNS (DDNS) is the protocol that allows machines on a network to communicate with, and register their IP address on, a Domain Name System (DNS) server.

The DDNS must be registered to enable this option. See [Registering a DDNS on page 56](#) for additional information.

8. Click **Next**.

The Configure Interface Settings summary page appears. The values listed reflect the new system and interface state.

9. Click **Finish** and **OK**.

## Modifying an IP Alias Interface

This topic describes how to modify the settings for an existing virtual interface.

### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware > Network > Interfaces**.
3. In the **Interfaces** column, select the checkbox of the interface and disable the IP alias interface by clicking **No** in the **Enabled** column. Click **OK** in the warning dialog box.
4. In the **Interfaces** column, select the checkbox of the interface and click **Configure**.
5. In the Configure IP Alias dialog box, change the settings that are described in the procedure [Configuring an IP Alias on page 55](#).
6. Click **Next** and **Finish**.

## Registering a DDNS

DDNS (Dynamic DNS) is the protocol used by CIFS that allows machines on a network to communicate with, and register their IP address on, a DNS Server. You can do the following:

- ◆ Manually register (add) configured interfaces to the DDNS registration list.
- ◆ Remove interfaces from the DDNS registration list.
- ◆ Enable or disable DNS updates.
- ◆ Display whether DDNS registration is enabled or not.
- ◆ Display interfaces in the DDNS registration list.

### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware > Network > Interfaces**.
3. Click **DDNS Registration**.
4. In the DDNS Registration dialog box, to add an interface to the DDNS, click **Add**.  
The Add Interface dialog box appears.
  - a. Enter a name in the **Interface** field.
  - b. Click **OK**.
5. Optionally, to remove an interface from the DDNS:
  - a. Select the interface to remove, and click **Remove**.



- b. In the Confirm Remove dialog box, click **OK**.
6. Specify the DDNS Status.
  - Select **Enable** to enable updates for all interfaces already registered.
  - Click **Default** to select the default settings for DDNS updates.
  - Clear **Enable** to disable DDNS updates for the registered interfaces.
7. To complete the DDNS registration, click **OK**.

## Destroying an Interface

You can use DD System Manager to destroy or delete virtual, VLAN, and IP alias interfaces. When a virtual interface is destroyed, the system deletes the virtual interface, releases its bonded physical interface, and deletes any VLANs or aliases attached to the virtual interface. When you delete a VLAN interface, the OS deletes the VLAN and any IP alias interfaces that are created under it. When you destroy an IP alias, the OS deletes only that alias interface.

### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware** > **Network** > **Interfaces**.
3. Click the box next to the interface to destroy (Virtual or VLAN or IP Alias).
4. Click **OK** to confirm.

## Viewing an Interface Hierarchy in the Tree View

### Procedure

1. In the navigation panel, select the system to view.
2. Select **Hardware** > **Network** > **Interfaces**.
3. Click **Tree View**.
4. In the Tree View dialog box, select the plus or minus boxes to expand or contract the tree view that shows the hierarchy.
5. Click **Close** to exit this view.

## Configuring Network Settings

Use the **Hardware** > **Network** > **Settings** view to view and configure the network settings. Settings includes network parameters such as the hostname, domain name, search domains, host mapping, and DNS list.

See the following sections:

- ◆ [Viewing Network Settings Information on page 57](#)
- ◆ [Setting the DD System Manager Hostname on page 58](#)
- ◆ [Managing the Domain Search List on page 59](#)
- ◆ [Mapping Host Names to IP Addresses on page 59](#)
- ◆ [Set DNS IP Addresses on page 60](#)

## Viewing Network Settings Information

### Procedure

1. In the navigation panel, select the system to view or configure.

2. Select **Hardware > Network > Settings**.

The Settings view has Host Settings, Search Domain, and Host Mapping options.

**Results**

The Settings tab displays the following information.

**Table 13** Network Settings

Item	Description
Host Settings	
Host Name	The hostname of the selected system.
Domain Name	The fully qualified domain name associated with the selected system.
Search Domain List	
Search Domain	A list of search domains that the selected system uses. The system applies the search domain as a suffix to the hostname.
Hosts Mapping	
IP Address	IP address of the host to resolve.
Host Name	Hostnames associated with the IP address.
DNS List	
DNS IP Address	Current DNS IP addresses associated with the selected system. An asterisk (*) indicates that the IP addresses were assigned through DHCP.

**Setting the DD System Manager Hostname**

You can configure the DD System Manager hostname and domain name manually, or you can configure DD OS to automatically receive the host and domain names from a Dynamic Host Configuration Protocol (DHCP) server. One advantage to manually configuring the host and domain names is that you remove the dependency on the DHCP server and the interface leading to the DHCP server. To minimize the risk of service interruption, EMC recommends that you manually configure the host and domain names.

- ◆ Do not include an underscore in the hostname. It is incompatible with some browsers.

Changing the names of an active host could—

- ◆ cause a break in the current connection. If this happens, log back in and check the saved settings.
- ◆ disrupt replication and CIFS active directory authentication. If this happens, reconfigure CIFS authentication after you change the names.

**Procedure**

1. In the navigation panel, select the system to configure.
2. Select **Hardware > Network > Settings**.
3. Click **Edit** in the **Host Settings** area. The Configure Host dialog box appears.
4. To manually configure the host and domain names:

- a. Select **Manually configure the host**.  
The DNS IP address checkboxes become active.
  - b. Enter a hostname in the **Host Name** box.  
For example, `id##.yourcompany.com`
  - c. Enter a domain name in the **Domain Name** box.  
This is the domain name associated with your Data Domain system and, usually, your company's domain name. For example, `yourcompany.com`
  - d. Click **OK**.  
The system displays progress messages as the changes are applied.
5. To obtain the host and domain names from a DHCP server, select **Obtain Settings using DHCP** and click **OK**.  
At least one interface must be configured to use DHCP.

## Managing the Domain Search List

This topic describes how to add and delete entries in the domain search list.

### Procedure

1. In the Settings view, click **Edit** in the Search Domain List area.
2. To add a search domain using the Configure Search Domains dialog box:
  - a. Click Add (+).
  - b. In the Add Search Domain dialog box, enter a name in the **Search Domain** box.  
For example, `id##.yourcompany.com`
  - c. Click **OK**.  
The system adds the new domain to the list of searchable domains.
  - d. Click **OK** to apply changes and return to the Settings view.
3. To remove a search domain using the Configure Search Domains dialog box:
  - a. Select the search domain to remove.
  - b. Click Delete (X).  
The system removes the selected domain from the list of searchable domains.
  - c. Click **OK** to apply changes and return to the Settings view.

## Mapping Host Names to IP Addresses

Use the Hosts Mapping area to add a mapping that ties an IP address to a name. The following topics describe how to manage host maps:

- ◆ [Adding a Host Map on page 59](#)
- ◆ [Deleting a Host Map on page 60](#)

### Adding a Host Map

#### Procedure

1. In the Settings view, click **Add** in the Hosts Mapping area.
2. In the Add Hosts dialog box, enter the IP address of the host in the **IP Address** text boxes.

The Internet Protocol (IP) Address is the numerical label assigned to the interface, such as 192.168.10.23.

3. Click Add (+).

In the Add Host dialog box, enter a hostname in the **Host Name** box for the listed system, such as `id##.yourcompany.com`

4. Click **OK** to add the new hostname is added to the list of Host Names. Click **OK** to return to the Settings tab.

## Deleting a Host Map

### Procedure

1. In the Settings view, select the host mapping to delete.
2. Click **Delete** in the Hosts Mapping area and click **Delete** to confirm. Confirmation messages are displayed.
3. Click **Close** after the Completed message appears to return to the Settings tab.

## Set DNS IP Addresses

You can configure the DNS IP addresses manually, or you can configure DD OS to automatically receive IP addresses from a DHCP server. One advantage to manually configuring DNS IP addresses is that you remove the dependency on the DHCP server and the interface leading to the DHCP server. To minimize the risk of service interruption, EMC recommends that you manually configure the DNS IP addresses.

### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware > Network > Settings**.
3. Click **Edit** in the **DNS List** area.
4. To manually add a DNS IP address:
  - a. Select **Manually configure DNS list**.  
The DNS IP address checkboxes become active.
  - b. Click Add (+).
  - c. In the Add DNS dialog box, enter the DNS IP address to add.
  - d. Click **OK**.  
The system adds the new IP address to the list of DNS IP addresses.
  - e. Click **OK** to apply the changes.
5. To delete a DNS IP address from the list:
  - a. Select **Manually configure DNS list**.  
The DNS IP address checkboxes become active.
  - b. Select the DNS IP address to delete and click **Delete (X)**.  
The system removes the IP address from the list of DNS IP addresses.
  - c. Click **OK** to apply the changes.
6. To obtain DNS addresses from a DHCP server, select **Obtain DNS using DHCP** and click **OK**.

At least one interface must be configured to use DHCP.

## Configuring Routes

Routes determine the path taken to transfer data to and from the localhost (the Data Domain system) to another network or host.

Data Domain systems do not generate or respond to any of the network routing management protocols (RIP, EGRP/EIGRP, and BGP). The only routing implemented on a Data Domain system is based upon the internal route table, in which the administrator may define a specific network or subnet that a physical interface (or interface group) uses.

Data Domain systems use source-based routing, which means that outbound network packets that match the subnet of multiple interfaces are routed only over the physical interface from which they originated.

Set static routes multiple interfaces contain the same IPv6 subnets, and the connections are being made to IPv6 addresses with this subnet. Normally, static routes are not needed with IPv4 addresses with the same subnet, such as for backups. There are cases in which static addresses may be required to allow connections to work, such as connections from the Data Domain system to remote systems.

---

### Note

Routing for connections initiated from the Data Domain system, such as for replication, depends on the source address used for interfaces on the same subnet. To force traffic for a specific interface to a specific destination (even if that interface is on the same subnet as other interfaces), configure a static routing entry between the two systems: this static routing overrides source routing.

---

Configuring routes is described in these sections:

- ◆ [Viewing Route Information on page 61](#)
- ◆ [Setting the Default Gateway on page 62](#)
- ◆ [Creating Static Routes on page 62](#)
- ◆ [Deleting Static Routes on page 63](#)

## Viewing Route Information

### Procedure

1. In the navigation panel, select the system to view or configure.
2. Select **Hardware** > **Network** > **Routes**.

### Results

The Routes view presents the following information. IP routing tables show the destination, gateway, netmask, and other information for each route.

**Table 14** Routes

Item	Description
Static Routes	
Route Spec	Lists the route specification used to configure routes.
Dynamic Routes	List of dynamically assigned routes using network or host paths for data transmission.

**Table 14** Routes (continued)

Item	Description
Destination	The destination host/network where the network traffic (data) is sent.
Gateway	The address of the router in the DD network, or 0.0.0.0 if no gateway is set.
Genmask	The netmask for the destination net. Set to 255.255.255.255 for a host destination and 0.0.0.0 for the default route.
Flags	Possible flags include: U—Route is up, H—Target is a host, G —Use gateway, R —Reinstate route for dynamic routing, D—Dynamically installed by daemon or redirect, M —Modified from routing daemon or redirect, A —Installed by addrconf, C —Cache entry, and ! —Reject route.
Metric	The distance to the target (usually counted in hops). Not used by the DD OS, but might be needed by routing daemons.
MTU	Maximum Transfer Unit (MTU) size for the physical (Ethernet) interface. See <a href="#">About MTU Size Values on page 48</a> .
Window	Default window size for TCP connections over this route.
IRTT	Initial RTT (Round Trip Time) used by the kernel to estimate the best TCP protocol parameters without waiting on possibly slow answers.
Interface	Interface name associated with the routing interface.

## Setting the Default Gateway

You can configure the default gateway manually, or you can configure DD OS to automatically receive the default gateway IP addresses from a DHCP server. One advantage to manually configuring the default gateway is that you remove the dependency on the DHCP server and the interface leading to the DHCP server. To minimize the risk of service interruption, EMC recommends that you manually configure the default gateway IP address.

### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware > Network > Routes**.
3. Click **Edit** in the Default Gateway area.
4. To manually configure the default gateway address:
  - a. Select **Manually Configure**.
  - b. Enter the gateway address in the **Gateway** boxes.
  - c. Click **OK**.
5. To obtain the default gateway address from a DHCP server, select **Use DHCP value** and click **OK**.

At least one interface must be configured to use DHCP.

## Creating Static Routes

### Procedure

1. In the navigation panel, select the system to configure.

2. Select **Hardware** > **Network** > **Routes**.
3. Click **Create** in the Static Routes area.
4. In the **Create Routes** dialog box, select the interface that will host the static route, and click **Next**.
5. Specify the Destination. Select one of the following:
  - To specify a destination network, select **Network** and enter the network address and netmask for the destination network.
  - To specify a destination host, select **Host** and enter the hostname or IP address of the destination host.
6. Optionally, specify the gateway to use to connect to the destination network or host.
  - a. Select **Specify different gateway for this route**.
  - b. Enter the gateway address.
7. Review the configuration and click **Next**.  
The Create Routes Summary page appears. The values listed reflect the new configuration.
8. Click **Finish**.
9. After the process is completed, click **OK**.  
The new route specification is listed in the Route Spec list.

## Deleting Static Routes

### Procedure

1. In the navigation panel, select the system to configure.
2. Select **Hardware** > **Network** > **Routes**.
3. Select the Route Spec of the route specification to delete.
4. Click **Delete**.
5. Click **Delete** to confirm and then click **Close**.

The selected route specification is removed from the Route Spec list.

## Managing Access to the System

Access management includes viewing and configuring the services that provide administrator and user access to the system.

The tasks to manage access to the system include:

- ◆ [Managing Administrator Access on page 63](#)
- ◆ [Managing Local User Access to the System on page 68](#)
- ◆ [Managing NIS Servers and Workgroups on page 75](#)
- ◆ [Managing Windows Servers and Workgroups on page 77](#)

### Managing Administrator Access

The following tasks can be performed to manage administrator access:

- ◆ [Viewing Administrator Access on page 64](#)
- ◆ [Managing FTP Access on page 64](#)

- ◆ [Managing FTPS Access on page 65](#)
- ◆ [Managing HTTP and HTTPS Access on page 66](#)
- ◆ [Managing SSH and SCP Access on page 67](#)
- ◆ [Managing Telnet Access on page 68](#)

## Viewing Administrator Access

This topic describes how to view the configured information for administrator access.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management**.

### Results

The Access Management page has the Administrator Access, Local Users, NIS, and Windows tabs.

The Administrator Access view lists this information.

**Table 15** Administrator Access

Item	Description
Services	The name of a service/protocol that can access the system.
Enabled (Yes/No)	The status of the service. If the service is disabled, enable it by selecting it in the list and clicking <b>Configure</b> . Fill out the General tab of the dialog box. If the service is enabled, modify its settings by selecting it in the list and clicking <b>Configure</b> . Edit the settings in the General tab of the dialog box.
Allowed Hosts	The host or hosts that can access the service.
Service Options	The port or session timeout value for the service selected in the list.
<p><b>Note</b></p> <p>You cannot set options for DD OS 5.2 Data Domain systems that are managed by systems running DD OS 5.3 and later.</p>	
FTP/FTPS	Only the session timeout can be set.
HTTP port	The port number opened for the HTTP protocol (port 80, by default).
HTTPS port	The port number opened for the HTTPS protocol (port 443, by default).
SSH/SCP port	The port number opened for the SSH/SCP protocol (port 22, by default).
Telnet	No port number can be set.
Session Timeout	The amount of inactive time allowed before a connection closes. The default is Infinite, that is, the connection does not close. Use the Advanced tab of the dialog box to set a timeout in seconds.

## Managing FTP Access

You can enable either FTP or FTPS access. FTP access allows user names and passwords to cross the network in clear text, making FTP an insecure access method. FTPS is



recommended as a secure access method. When you enable either FTP or FTPS access, the other access method is disabled.

---

#### Note

LFTP clients that connect to a Data Domain system via FTPS or FTP are disconnected after reaching a set timeout limit. However the LFTP client uses its cached username and password to reconnect after the timeout while you are running any command.

---

#### Procedure

1. On the Access Management page, select **FTP** and click **Configure**.
  2. In the General tab of the Configure FTP Access dialog box, select **Allow FTP Access**.
  3. Determine how hosts connect:
    - To allow complete access, select **Allow all hosts** to connect.
    - To configure specific hosts, select **Limit Access** to the following systems, and do one of the following:
- 

#### Note

Hostnames can be a fully qualified hostname or an IP address.

---

- To add a host, click Add (+). Enter the hostname, and click **OK**.
  - To modify a hostname, select the hostname in the **Hosts** list and click Edit (pencil). Change the hostname and click **OK**.
  - To remove a hostname, select the hostname in the **Hosts** list and click Remove (-), and click **OK**.
4. Click **OK**. If FTPS is enabled, you are warned that it will be disabled and asked to click **OK** to proceed.
  5. To set a session timeout, select the Advanced tab, enter the timeout value in seconds, and click **OK**.
- 

#### Note

The session timeout default is Infinite, that is, the connection does not close.

---

## Managing FTPS Access

FTPS provides additional security over using FTP, such as support for the Transport Layer Security (TLS) and for the Secure Sockets Layer (SSL) cryptographic protocols.

---

#### Note

You can enable either FTP or FTPS access. When you enable one, the other is disabled.

FTPS does not show up as a service for DD systems that run DD OS 5.2, managed from a DD system running DD OS 5.3 or later.

**FTPS only:** When you issue the `get` command, matching versions of SSL have to be installed on the Data Domain system and compiled on the LFTP client to prevent this fatal error: `SSL_read: wrong version number lftp`. As a workaround, attempt to re-issue the `get` command on the same file.

---

#### Procedure

1. On the Access Management page, select **FTPS** and click **Configure**.

2. In the General tab of the Configure FTPS Access dialog box, select **Allow FTPS Access**.
3. Determine how hosts connect:
  - To allow complete access, select **Allow all hosts to connect**.
  - To configure specific hosts, select **Limit Access** to the following systems, and do one of the following:

---

**Note**

Hostnames can be a fully qualified hostname or an IP address.

- To add a host, click Add (+). Enter the hostname and click **OK**.
  - To modify a hostname, select the hostname in the **Hosts** list and click Edit (pencil). Change the hostname and click **OK**.
  - To remove a hostname, select the hostname in the **Hosts** list and click Remove (-).
4. Click **OK**. If FTP is enabled, you are warned that it will be disabled and asked to click **OK** to proceed.
  5. To set a session timeout, open the Advanced tab and enter the timeout value in seconds. Click **OK**.

---

**Note**

The session timeout default is Infinite, that is, the connection does not close.

## Managing HTTP and HTTPS Access

This topic describes how to provide access to the system using the HTTP and HTTPS protocols.

### Procedure

1. On the Access Management page, select **HTTP and/or HTTPS** and click **Configure**.
2. In the General tab of the **Configure HTTP/HTTPS Access** dialog box, select either or both Allow Access options.
3. Determine how hosts connect:
  - To allow complete access, select **Allow all hosts** to connect.
  - To configure specific hosts, select **Limit Access to the following systems**, and do one of the following:

---

**Note**

Hostnames can be a fully qualified hostname or an IP address.

- To add a host, click Add (+). Enter the hostname, and click **OK**.
  - To modify a hostname, select the hostname in the **Hosts** list and click Edit (pencil). Change the hostname and click **OK**.
  - To remove a hostname, select the hostname in the **Hosts** list and click Remove (-), and click **OK**.
4. To configure system ports and session timeout values, click the Advanced tab.
    - In the **HTTP Port** text entry box, enter the port for connection. Port 80 is assigned by default.
    - In the **HTTPS Port** text entry box, enter the port for connection. Port 443 is assigned by default.

- In the **Session Timeout** text entry box, enter the interval in seconds that must elapse before connection closes. The minimum is 60 seconds and the maximum is 31536000 seconds (one year).

---

#### Note

The session timeout default is Infinite, that is, the connection does not close.

---

5. Click **OK**.

## Managing SSH and SCP Access

You can use DD System Manager to enable administrator access using the SSH protocol, with or without SCP (secure copy). SCP requires SSH, so when SSH is disabled, SCP is automatically disabled.

---

#### Note

SCP does not show up as a service for DD systems that run DD OS 5.2, managed from a DD system running DD OS 5.3 or later.

---

#### Procedure

1. On the Access Management page, select **SSH** or **SCP**.
2. In the General tab of the Configure SSH/SCP Access dialog box, select **Allow SSH Access**.
3. If you want to also enable SCP access, select **Allow SCP Access**.
4. Determine how hosts connect:
  - To allow complete access, select **Allow all hosts to connect**.
  - To configure specific hosts, select **Limit Access** to the following systems, and do one of the following:

---

#### Note

Hostnames can be a fully qualified hostname or an IP address.

---

- To add a host, click Add (+). Enter the hostname, and click **OK**.
  - To modify a hostname, select the hostname in the Hosts list and click Edit (pencil). Change the hostname and click **OK**.
  - To remove a hostname, select the hostname in the **Hosts** list and click Remove (-), and click **OK**.
5. To configure system ports and session timeout values, click the **Advanced** tab.
    - In the **SSH/SCP Port** text entry box, enter the port for connection. Port 22 is assigned by default.
    - In the **Session Timeout** text entry box, enter the interval in seconds that must elapse before connection closes.

---

#### Note

The session timeout default is Infinite, that is, the connection does not close.

---

#### Note

Click **Default** to revert to the default value.

---

6. Click **OK**.

## Managing Telnet Access

This topic describes how to enable administrator Telnet access to the system.

---

### Note

Telnet access allows user names and passwords to cross the network in clear text, making Telnet an insecure access method.

---

### Procedure

1. On the Access Management page, select **Telnet** and click **Configure**.
  2. In the General tab of the Configure Telnet Access dialog box, select **Allow Telnet Access**.
  3. Determine how hosts connect:
    - To allow complete access, select **Allow all hosts** to connect.
    - To configure specific hosts, select **Limit Access** to the following systems, and do one of the following:
- 

### Note

Hostnames can be a fully qualified hostname or an IP address.

---

- To add a host, click Add (+). Enter the hostname, and click **OK**.
  - To modify a hostname, select the hostname in the **Hosts** list and click Edit (pencil). Change the hostname and click **OK**.
  - To remove a hostname, select the hostname in the **Hosts** list and click Remove (-), and click **OK**.
4. Click **OK**.
  5. To set a session timeout, open the Advanced tab and enter the timeout value in seconds. Click **OK**.
- 

### Note

The session timeout default is Infinite, that is, the connection does not close.

---

## Managing Local User Access to the System

The following sections describe the tasks to manage user access:

- ◆ [Viewing Local User Information on page 69](#)
- ◆ [Creating Local Users on page 71](#)
- ◆ [Modifying a Local User Profile on page 72](#)
- ◆ [Deleting a Local User on page 73](#)
- ◆ [Enabling and Disabling Local Users on page 73](#)
- ◆ [Enabling Security Authorization on page 74](#)
- ◆ [Changing User Passwords on page 74](#)
- ◆ [Modifying the Password Policy on page 74](#)

## Viewing Local User Information

This topic describes how to view the configuration for local users on a Data Domain system.

### Note

The user-authentication module uses Greenwich Mean Time (GMT). Therefore, the expiration date for disabling a user's account and password expiration dates should reflect GMT instead of local time.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management > Local Users** .

The Local Users view appears and shows the Local Users table and the Detailed Information area.

The Local Users table lists the following information.

**Table 16** Local Users

Item	Description
Name	The user ID, as added to the system.
Role	<p>Possible roles of users based on a set of privileges:</p> <ul style="list-style-type: none"> <li>• Admin role: Allows one to administer, that is, configure and monitor, the entire system.</li> <li>• User role: Allows one to monitor systems and perform the fastcopy operation.</li> <li>• Security role: In addition to the user role privileges, allows one to set up security-officer configurations and manage other security-officer operators.</li> <li>• Backup-operator role: In addition to the user role privileges, a back-up operator can create snapshots for MTrees, and import, export tapes and move tapes between elements in a virtual tape library. A user in the role of backup-operator can also add and delete <code>ssh</code> public keys for nonpassword-required log ins, such as for automated scripting. A back-up operator can add, delete, reset, and view CLI command aliases, synchronize modified files, wait for replication to complete on the destination system, and change a password for a user.</li> <li>• Data-access role: Intended for DD Boost authentication, an operator with this role <i>cannot</i> monitor or configure a Data Domain system.</li> </ul>
Status	<ul style="list-style-type: none"> <li>• Active—User access to the account is permitted.</li> <li>• Disabled—User access to the account is denied because the expiration date for the account has been reached or a locked account's password has not been renewed. Admin users can disable/enable users with admin or user roles, except the <code>sysadmin</code> user which cannot be disabled. Security officers can only disable/enable other security officers.</li> </ul>

**Table 16** Local Users (continued)

Item	Description
	<ul style="list-style-type: none"> <li>Locked—User access the account is denied because the password has expired.</li> </ul>
Disable Date	The date the account is set to be disabled.
Last Login From	The location where the user last logged in.
Last Login Time	The time the user last logged in.

**Note**

Users who have admin or security officer roles can view all users. Users with other roles can view only their own user accounts.

3. Select the user you want to view from the list of users.

Information about the selected user displays in the Detailed Information area.

The Detailed Information area displays the following information about the selected user:

**Table 17** Detailed Information

Item	Description
Password Last Changed	The date the password was last changed.
Minimum Days Between Change	The minimum number of days between password changes that you allow a user. Default is 0.
Maximum Days Between Change	The maximum number of days between password changes that you allow a user. Default is 99999.
Warn Days Before Expire	The number of days to warn the users before their password expires. Default is 7.
Disable Days After Expire	The number of days after a password expires to disable the user account. Default is Never.

**Note**

The default values are the initial default password policy values. A system administrator (admin role) can change them using the Modify Password Policy task.

**UID Conflicts: Local User and NIS User Accounts**

Local user accounts on a Data Domain system start with a UID of 500. When you set up a Data Domain system in an NIS environment, be aware of potential UID conflicts between local and NIS user accounts. To avoid such conflicts, during initial planning consider the size of potential local accounts when you define allowable UID ranges for NIS users.

**User Roles**

To enhance security, each user can be assigned a different role. Roles enable you to restrict system access to a set of privileges. A Data Domain system supports the following roles:

- ◆ Admin role: Allows one to administer, that is, configure and monitor, the entire Data Domain system.
- ◆ User role: Allows one to monitor Data Domain systems and perform the fastcopy operation.
- ◆ Security role: In addition to the user role privileges, allows one to set up security-officer configurations and manage other security-officer operators.
- ◆ Backup-operator role: In addition to the user role privileges, allows one to create snapshots, import and export tapes to a VTL library and move tapes within a VTL library.
- ◆ Data-access role: Intended for DD Boost authentication, an operator with this role cannot monitor or configure a Data Domain system.

## Creating Local Users

This topic describes how to create new users on a Data Domain system.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management > Local Users**.  
The Local Users view appears.
3. Click the **Create** button to create a new user.  
The Create User dialog box appears.
4. Enter the following information in the General Tab:

Item	Description
User	The user ID or name.
Password	The user password. Set a default password, and the user can change it later.
Verify Password	The user password, again.
Role	<p>The role assigned to the user:</p> <ul style="list-style-type: none"> <li>• Admin role: Allows one to administer, that is, configure and monitor, the entire system.</li> <li>• User role: Allows one to monitor systems and perform the fastcopy operation.</li> <li>• Security role: In addition to the user role privileges, allows one to set up security-officer configurations and manage other security-officer operators.</li> <li>• Backup-operator role: In addition to the user role privileges, allows one to create snapshots, import and export tapes to a VTL library and move tapes within a VTL library.</li> <li>• Data-access role: Intended for DD Boost authentication, an operator with this role cannot monitor or configure a system.</li> </ul>

---

**Note**

The default value for the minimum length of a password or minimum number of character classes required for a user password is 1. Allowable character classes include:

- Lowercase letters (a-z)
  - Uppercase letters (A-Z)
  - Numbers (0-9)
  - Special Characters (\$, %, #, +, and so on)
- 

**Note**

The available roles display based on user’s role. Only the Sysadmin user can create the first security officer. After the first security officer is created, only security officers can create or modify other security officers. Sysadmin is the default admin user and cannot be deleted or modified.

---

5. Enter the following information in the Advanced tab:

---

Item	Description
Minimum Days Between Change	The minimum number of days between password changes that you allow a user. Default is 0.
Maximum Days Between Change	The maximum number of days between password changes that you allow a user. Default is 99999.
Warn Days Before Expire	The number of days to warn the users before their password expires. Default is 7.
Disable Days After Expire	The number of days after a password expires to disable the user account. Default is Never.
Disable account on the following date	Check this box and enter a date (mm/dd/yyyy) when you want to disable this account. Also, you can click the calendar to select a date.

---

6. Click **OK**.

---

**Note**

Note: The default password policy can change if the admin user changes them from the Modify Password Policy task. The default values are the initial default password policy values.

---

## Modifying a Local User Profile

This topic describes how to change a user profile.

**Procedure**

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management > Local Users**.  
The Local Users view appears.
3. Click a user name from the list.
4. Click **Modify** to make changes to a user account.



The Modify User dialog box appears.

5. Enter the following information in the General tab:

Item	Description
User	The user ID or name.
Role	Select the role from the list.

6. Enter the following information in the Advanced tab:

Item	Description
Minimum Days Between Change	The minimum number of days between password changes that you allow a user. Default is 0.
Maximum Days Between Change	The maximum number of days between password changes that you allow a user. Default is 99999.
Warn Days Before Expire	The number of days to warn the users before their password expires. Default is 7.
Disable Days After Expire	The number of days after a password expires to disable the user account. Default is Never.

7. Click **OK**.

## Deleting a Local User

You can delete certain users based on your user role. If one of the selected users cannot be deleted, the Delete button is disabled. For example, Sysadmin cannot be deleted. Admin users cannot delete security officers. Security officers can delete, enable, and disable other security officers.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management > Local Users**.  
The Local Users view appears.
3. Click one or more user names from the list.
4. Click **Delete** to delete the user accounts.  
The Delete User dialog box appears.
5. Click **OK** and **Close**.

## Enabling and Disabling Local Users

This topic describes how to enable and disable local users.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management > Local Users**.  
The Local Users view appears.
3. Click one or more user names from the list.
4. Click either **Enable** or **Disable** to enable or disable user accounts.  
The Enable or Disable User dialog box appears.

5. Click **OK** and **Close**.

## Enabling Security Authorization

You need to use the Data Domain system's command-line interface (CLI) to manage the security authorization policy. Log in using the security officer credential.

---

### Note

The Retention Lock Compliance license must be installed.

For more information about the commands used in this procedure, see the *EMC DD OS Command Reference Guide*.

---

### Procedure

1. A security user is needed to invoke some of the commands for Retention Lock Compliance. To set up a security user, see [Creating Local Users on page 71](#).
2. Log in with the security role and enable the security officer authorization policy by entering: `# authorization policy set security-officer enabled`

## Changing User Passwords

This topic describes how to change user passwords.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Click **System Settings > Access Management > Local Users**.  
The Local Users view appears.
3. Click a user name from the list.
4. Click **Change Password** to change the user password.  
The Change Password dialog box appears.  
If prompted, enter your old password.
5. Enter the new password into the **New Password** box.
6. Enter the new password again into **Verify New Password** box.
7. Click **OK**.

## Modifying the Password Policy

This topic describes how to modify the password policy for all DD System Manager users.

### Procedure

1. In the navigation panel, expand **DD Network** and select a managed system.
2. Select **System Settings > Access Management**.
3. Select **More Tasks > Modify Password Policy**.  
The Modify Password Policy dialog box appears.
4. Enter the password policy information in the appropriate boxes. To select the default value, click **Default** next to each value.

Item	Description
Minimum Days Between Change	The minimum number of days between password changes that you allow a user. Default is 0.
Maximum Days Between Change	The maximum number of days between password changes that you allow a user. Default is 99999.
Warn Days Before Expire	The number of days to warn the users before their password expires. Default is 7.
Disable Days After Expire	The number of days after a password expires to disable the user account. Default is Never.
Minimum Length of Password	The minimum password length required. Default is 1.
Minimum Number of Character Classes	The minimum number of character classes required for a user password. Default is 1. Character classes include: <ul style="list-style-type: none"> <li>• Lowercase letters (a-z)</li> <li>• Uppercase letters (A-Z)</li> <li>• Numbers (0-9)</li> <li>• Special Characters (\$, %, #, +, and so on)</li> </ul>

5. Click **OK** to save the password settings.

## Managing NIS Servers and Workgroups

Use NIS workgroup management to configure NIS authentication, domain names, and NIS groups as described in the following topics:

- ◆ [Viewing NIS Information on page 75](#)
- ◆ [Enabling and Disabling NIS Authentication on page 76](#)
- ◆ [Configuring the NIS Domain Name on page 76](#)
- ◆ [Specifying NIS Authentication Servers on page 76](#)
- ◆ [Configuring NIS Groups on page 77](#)

### Viewing NIS Information

This topic describes how to view the NIS configuration on a managed system.

#### Procedure

1. In the navigation panel, expand **DD Network** and select a managed system.
2. Select **System Settings** > **Access Management** > **NIS** .

The NIS view appears.

#### Results

The NIS view lists the following information.

Item	Description
Status	The status of the service, either enabled or disabled.
Domain Name	The name of the domain for this service.

Item	Description
Authentication Server	
Server	The name of the server performing authentication.
Configured NIS Groups	
Group	The name of the NIS group.
Role	The role of the group (admin or user).

## Enabling and Disabling NIS Authentication

This topic describes how to enable or disable NIS authentication for a managed system.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings** > **Access Management** > **NIS**.  
The NIS view appears.
3. Click **Enable** to enable or **Disable** to disable NIS Authentication.  
The Enable or Disable NIS dialog box appears.
4. Click **OK**.

## Configuring the NIS Domain Name

### Procedure

1. Click **Edit** next to Domain Name to edit the NIS domain name.  
The Configure NIS Domain Name dialog box appears.
2. Enter the domain name in the **Domain Name** box.
3. Click **OK**.

## Specifying NIS Authentication Servers

### Procedure

1. Click **Edit** below **Authentication Servers** to configure the authentication server.
2. In the Configure NIS Authentication Server dialog box, select one of the following:
  - **Obtain NIS Servers from DHCP**
  - **Manually Configure**
    - a. To add an authentication server, click Add (+) Enter the server name, and click **OK**.
    - b. To modify an authentication server, select the checkbox of the authentication server name in the server list and click the edit icon (pencil). Change the server name, and click **OK**.
    - c. To remove an authentication server name, select the checkbox of the hostname in the server list, click the X icon, and click **OK**.
3. Click **OK**.

## Configuring NIS Groups

### Procedure

1. Click **Edit** in the Configured NIS Groups area to configure the NIS groups.
2. In the Configure Allowed NIS Groups dialog box, select an NIS group.
  - To add an NIS group, click Add (+). Enter the NIS group name and role, and click **Validate**. Click **OK** to exit the add NIS group dialog box. Click **OK** again to exit the **Configure Allowed NIS Groups** dialog box.
  - To modify an NIS group, select the checkbox of the NIS group name in the NIS group list and click Edit (pencil). Change the NIS group name, and click **OK**.
  - To remove an NIS group name, select the NIS group in the list and click Delete (X). Click **OK**.

## Managing Windows Servers and Workgroups

Use Windows workgroup management to configure Windows authentication, configure Active Directory, and assign group roles.

### Note

For more information about Active Directory, see *Data Domain System Integration into Microsoft's Active Directory (AD)* in the *EMC DD OS Initial Configuration Guide*.

The tasks to manage Windows workgroups include:

- ◆ [Viewing Windows Information on page 77](#)
- ◆ [Configuring Workgroup Authentication Parameters on page 78](#)
- ◆ [Configuring Active Directory Authentication Parameters on page 78](#)
- ◆ [Specifying Allowed Groups on page 79](#)
- ◆ [Modifying Group Names and Roles on page 79](#)
- ◆ [Deleting Groups on page 80](#)

## Viewing Windows Information

This topic describes how to view the configuration for Windows access.

### Procedure

1. In the navigation panel, expand **DD Network** and select a managed system.
2. Select **System Settings > Access Management > Windows**.

The Windows view appears.

The Windows view lists the following information.

Item	Description
Authentication	
Mode	The name type of authentication mode (Workgroup or Active Directory).
Workgroup/Active Directory Names	The name of the Workgroup or Active Directory.
CIFS Server Name	The name of the CIFS Server in use.

Item	Description
WINS Server	The name of the WINS Server in use.
Allowed Groups	
Windows Group	The name of the Windows group.
Role	The role of the group (admin, user, or backup-operator).

## Configuring Workgroup Authentication Parameters

This topic describes how to configure Workgroup authentication parameters.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management > Windows**.  
The Windows view appears.
3. Click **Configure**.  
The Configure Authentication dialog appears.
4. In the Mode list box, select **Workgroup**.  
The Workgroup mode joins a Data Domain system to a workgroup domain.
5. Optionally, clear the **Use Default** box and enter a name in the Workgroup Name box.
6. Click the **Advanced** tab to set additional information.
7. Optionally, clear the **Use Default** box and enter a name in the CIFS Server Name box.
8. Click **OK**.

## Configuring Active Directory Authentication Parameters

This topic describes how to set Active Directory authentication parameters. The system must meet all active-directory requirements, such as a clock time that differs no more than five minutes from that of the domain controller.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management > Windows**.  
The Windows view appears.
3. Select **Configure**.  
The Configure Authentication dialog box appears.
4. In the Mode list box, select **Active Directory**.  
The active-directory mode joins a system to an active-directory domain.
5. In the Realm Name box, type the full realm name for the system, such as **domain1.local**.
6. In the **Domain Joining Credential** area, type a user name and password. Specify either a user in a domain to be joined, or a user in a domain that is a trusted domain of your company. Specify the user as user@realmname. Use the complete realm name. Ensure that the user name has sufficient privileges to join the system to the trusted domain. The user name and password must be compatible with Microsoft

requirements for the Active Directory domain being joined. This user must have permission to create accounts in this domain.

7. Click the **Advanced** tab to set additional information.
8. Optionally, to set a CIFS server name, in the **CIFS Server Name** area:
  - Select the checkbox to use the default CIFS server name.
  - Clear the checkbox and type a name in the CIFS Server Name box.
9. In the Domain Controller area, determine how domain controllers are assigned:
  - For automatic assignment, select **Automatically assign Domain Controllers**. This is the default and recommended method.
  - To add specific domain controllers, select **Manually assign Domain Controllers** and type a controller name in one of the Domain Controllers boxes. Up to three controller names can be added. You can enter fully qualified domain names, hostnames, or IP addresses.
10. Optionally, to set Organizational Units, in the Organizational Unit area:
  - Select the checkbox to use the default Organizational Unit.
  - Clear the checkbox and type a name in the Organizational Unit box.

---

**Note**

The account is moved to the new Organizational Unit.

11. Click **OK**.

## Specifying Allowed Groups

This topic describes how to specify which Windows groups are permitted access to Windows files.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management > Windows**.  
The Windows view appears.
3. Click **Create** in the Allowed Groups area.  
The Create Windows Group dialog appears.
4. Enter a name in the Group box. The domain for the group must be specified. For example, domain\group name.
5. In the Role list box, select either **Admin** or **User**.
6. Click **OK**.

## Modifying Group Names and Roles

This topic describes how to modify the group names and roles configured for Windows groups.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management > Windows**.  
The Windows view appears.

3. Select a Windows Group from the list.
4. Click **Modify** in the Allowed Groups area.  
The Edit Windows Group dialog appears.
5. Edit the name in the Group box. The domain for the group must be specified. For example, domain\group name.
6. In the Role list box, select either **Admin** or **User**.
7. Click **OK**.

## Deleting Groups

You cannot delete default Windows groups, such as Domain Admins. If the default Windows group is selected, the **Delete** button is grayed out.

### Procedure

1. In the navigation panel, expand the DD Network and select a system.
2. Select **System Settings > Access Management > Windows**.  
The Windows view appears.
3. Select a Windows Group from the list.
4. Select **Delete** in the Allowed Groups area.  
The Delete Windows Group dialog appears.
5. Click **OK**.

## Managing General Configuration Settings

The **System Settings > General Configuration** area allows you to view and set system configuration parameters.

General configuration settings include:

- ◆ [Working with Email Settings on page 80](#)
- ◆ [Working with Time and Date Settings on page 82](#)
- ◆ [Working with System Properties on page 83](#)
- ◆ [Working with SNMP on page 83](#)

## Working with Email Settings

The procedures for working with email settings include:

- ◆ [Configuring Mail Server Settings on page 80](#)
- ◆ [Configuring the Autosupport Mailing List on page 81](#)
- ◆ [Testing the Alerts Email List on page 81](#)

## Configuring Mail Server Settings

This topic describes how to specify the mail server to which the OS sends email reports.

### Procedure

1. Select **System Settings > General Configuration > Mail Server**.
2. Select **More Tasks > Set Mail Server**.  
The Set Mail Server dialog box appears.



3. Enter the name of the mail server in the **Mail Server** box.
4. Click **OK**.

## Configuring the Autosupport Mailing List

This topic describes how to add, modify, and delete names in the autosupport mailing list. Autosupport emails are sent through the configured mail server (see [Configuring Mail Server Settings on page 80](#)) to all subscribers in the autosupport email list. After you configure the mail server and autosupport email list, it is a good practice to test the setup to ensure that autosupport messages reach the intended destinations.

### Procedure

1. Select **Maintenance** › **Support** › **Autosupport**.
2. Click **Configure**.  
The Configure Autosupport Subscribers dialog box appears.
3. To add a subscriber, do the following:
  - a. Click Add (+).  
The Email dialog box appears.
  - b. Enter the recipients email address in the Email box.
  - c. Click OK.
4. To delete a subscriber, do the following:
  - a. In the Configure Autosupport Subscribers dialog box, select the subscriber to delete.
  - b. Click Delete (X).
5. To modify a subscriber email address, do the following:
  - a. In the Configure Autosupport Subscribers dialog box, select the subscriber name to edit.
  - b. Click Modify (pencil icon).  
The Email dialog box appears.
  - c. Modify the email address as needed.
  - d. Click OK.
6. Click **OK** to close the Configure Autosupport Subscribers dialog box.  
The revised autosupport email list appears in the Autosupport Mailing List area.

## Testing the Alerts Email List

### Procedure

1. Click **Maintenance** › **Support** › **Autosupport**. In the **Vendor Support** area, verify that the Notification Status for vendors is Enabled. You cannot change the email address.
2. Select **Send Test Alert** from the More Tasks menu.  
The Send Test Alert dialog box appears.
3. In the **Notification Groups** area, select groups to send test emails and click **Next**.
4. Optionally, add or create additional email addresses.
5. Click **Send Now** and **OK**.

- Click **Status** › **Alerts** › **Notification**. In the **Alerts** area, verify that default is selected as a group. Verify that the **Subscribers Email List** contains the administrator and `autosupport-alert@autosupport.datadomain.com` addresses.

---

**Note**

You can create other groups; add, modify, and delete email subscribers to groups; and set class attributes. Class attributes refer to the name of a class, such as Hardware or File System. You set the severity of ranking, such as Warning or Critical, that is to trigger an alert for each class attribute.

---

**Results**

For more information, such as how to set up class attributes and groups, see [Working with the Notification View on page 116](#).

To test newly added alerts emails for mailer problems, enter: `autosupport test email email-addr`

For example, after adding the email address `djones@yourcompany.com` to the list, check the address with the command: `autosupport test email djones@yourcompany.com`

## Working with Time and Date Settings

The procedures for working with time and date settings include:

- ◆ [Viewing Time and Date Information on page 82](#)
- ◆ [Configuring Time and Date Settings on page 82](#)

### Viewing Time and Date Information

**Procedure**

- Select the system to be checked in the navigation panel.
- Select **System Settings** › **General Configuration** › **Time and Date Settings**.

**Results**

The Time and Date Settings page presents the current system date and time, and shows whether NTP is enabled or not, and the IP addresses or hostnames of configured NTP servers.

### Configuring Time and Date Settings

This topic describes how to configure time and date settings.

**Procedure**

- Select **System Settings** › **General Configuration** › **Time and Date Settings**.
- Select **More Tasks** › **Configure Time Settings**.

The Configure Time Settings dialog box appears.

- In the **Time Zone** list box, select the time zone where the Data Domain system resides.
- Set how system time is synchronized:
  - To manually set the time and date, select **None**, type the date in the **Date** box, and set the time in the **Time** list boxes.
  - To use NTP to synchronize the time, select NTP. Set how the NTP server is accessed:

- To use DHCP to automatically select a server, select **Obtain NTP Servers using DHCP**.
- To configure an NTP server IP address, select **Manually Configure**, add the IP address of the server, and click **OK**.

---

#### Note

Using time synchronization from an Active Directory domain controller might cause excessive time changes on the system if both NTP and the domain controller are modifying the time.

---

5. Click **OK**.
6. If you changed the time zone, you must reboot the system.
  - a. Select **Maintenance > System**.
  - b. From the More Tasks menu, select Reboot System.
  - c. Click OK to confirm.

## Working with System Properties

The procedures for working with system property settings include:

- ◆ [Viewing System Properties on page 83](#)
- ◆ [Configuring System Properties on page 83](#)

### Viewing System Properties

#### Procedure

1. Select the system to be checked in the navigation panel.
2. Select **System Settings > General Configuration > System Properties**.

The System Properties tab displays the system location, the administrator email address, and the administrator hostname.

### Configuring System Properties

This topic describes how to configure system properties that identify the managed system location, administrator, and host name.

#### Procedure

1. In the navigation panel, expand **DD Network** and select a managed system.
2. Select **System Settings > General Configuration > System Properties**.
3. Select **More Tasks > Set System Properties**.

The Set System Properties dialog box appears.

4. In the **Location** box, enter information about where the Data Domain system is located.
5. In the **Admin Email** box, enter the email address of the system administrator.
6. In the **Admin Host** box, enter the name of the administration server.
7. Click **OK**.

## Working with SNMP

The Simple Network Management Protocol (SNMP) is a standard protocol for exchanging network management information, and is a part of the Transmission Control Protocol/

Internet Protocol (TCP/IP) protocol suite. SNMP provides a tool for network administrators to manage and monitor network-attached devices, such as Data Domain systems, for conditions that warrant administrator attention.

To monitor Data Domain systems using SNMP, you will need to install the Data Domain MIB in your SNMP Management system. The Data Domain MIB can be obtained by following the instructions in [Downloading the SNMP MIB on page 86](#). The Data Domain MIB will allow SNMP queries for Data Domain-specific information.

DD OS also support the standard MIB-II so you can also query MIB-II statistics for general data such as network statistics. For full coverage of available data you should utilize both the Data Domain MIB and the standard MIB-II MIB.

Data Domain systems support SNMP V2C and/or SNMP V3. SNMP V3 provides a greater degree of security than V2C by replacing cleartext community strings as a means of authentication with user-based authentication using either MD5 or SHA1. As well, with SNMP V3, user authentication packets can be encrypted and their integrity verified with either DES or AES.

The default port that is open when SNMP is enabled is port 161. Traps are sent out through port 162.

- ◆ The *EMC DD OS Initial Configuration Guide* describes how to set up the Data Domain system to use SNMP monitoring.
- ◆ The *EMC DD OS Command Reference Guide* describes the full set of MIB parameters included in the Data Domain MIB branch.

The procedures for working with SNMP include:

- ◆ [Viewing SNMP Status and Configuration on page 84](#)
- ◆ [Enabling and Disabling SNMP on page 86](#)
- ◆ [Downloading the SNMP MIB on page 86](#)
- ◆ [Configuring SNMP Properties on page 86](#)
- ◆ [Managing SNMP V3 Users on page 86](#)
- ◆ [Managing SNMP V3 and V2C Trap Hosts on page 88](#)
- ◆ [Managing SNMP V2C Communities on page 89](#)

## Viewing SNMP Status and Configuration

### Procedure

1. Select the system to be checked in the navigation panel.
2. Select **System Settings** > **General Configuration** > **SNMP**.

The SNMP view shows the SNMP status, SNMP properties, SNMP V3 configuration, and SNMP V2C configuration.

### Status

The SNMP status displays the operational status of the SNMP agent on the system: Enabled or Disabled.

### SNMP Properties

Item	Description
SNMP System Location	The location of the Data Domain system being monitored.

Item	Description
SNMP System Contact	The person designated as the person to contact for the Data Domain system administration.

### SNMP V3 Configuration

Item	Description
SNMP Users	
Name	The name of the user on the SNMP manager with access to the agent for the Data Domain system.
Access	The access permissions for the SNMP user. This can be: <ul style="list-style-type: none"> <li>• Read-only</li> <li>• Read-write</li> </ul>
Authentication Protocols	The Authentication Protocol used to validate the SNMP user. This can be: <ul style="list-style-type: none"> <li>• MD5</li> <li>• SHA1</li> <li>• None</li> </ul>
Privacy Protocol	The encryption protocol used during the SNMP user authentication. This can be: <ul style="list-style-type: none"> <li>• AES</li> <li>• DES</li> <li>• None</li> </ul>
Trap Hosts	
Host	The IP address or domain name of the SNMP management host.
Port	The port used for SNMP trap communication with the host. For example, 162 is the default.
User	The user on the trap host authenticated to access the Data Domain SNMP information.

### SNMP V2C Configuration

Item	Description
<b>Communities</b>	
Community	The name of the community. For example, public, private, or localCommunity.
Access	The access permission assigned. This can be: <ul style="list-style-type: none"> <li>• Read-only</li> <li>• Read-write</li> </ul>
Hosts	The hosts in this community.

Item	Description
<b>Trap Hosts</b>	
Host	The systems designated to receive SNMP traps generated by the Data Domain system. If this parameter is set, systems receive alert messages, even if the SNMP agent is disabled.
Port	The port used for SNMP trap communication with the host. For example, 162 is the default.
Community	The name of the community. For example, public, private, or localCommunity.

## Enabling and Disabling SNMP

### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the Status area, click **Enable** or **Disable**.

## Downloading the SNMP MIB

### Procedure

1. From the **System Settings > General Configuration > SNMP** page, click **Download MIB file**.
2. In the Opening DATA\_DOMAIN.mib dialog box, select **Open**.
3. Click **Browse** and select a browser to view the MIB in a browser window.

---

### Note

If using the Microsoft Internet Explorer browser, enable Automatic prompting for file download.

---

4. Save the MIB or exit the browser.

## Configuring SNMP Properties

### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the SNMP Properties area, click **Configure**.  
The SNMP Configuration dialog box appears.
3. In the text fields, add an SNMP system location (a description of where the Data Domain system is located) and/or an SNMP system contact (for example, the email address of the system administrator for the Data Domain system).
4. Click **OK**.

## Managing SNMP V3 Users

The following topics describe how to manage SNMP V3 users:

- ◆ [Creating SNMP V3 Users on page 87](#)
- ◆ [Modifying SNMP V3 Users on page 87](#)
- ◆ [Removing SNMP V3 Users on page 87](#)

## Creating SNMP V3 Users

### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the SNMP Users area, click **Create**.  
The Create SNMP User dialog box appears.
3. In the **Name** text field, enter the name of the user on the SNMP manager who will have access to the agent for the Data Domain system. The name must be a minimum of eight characters.
4. Select either read-only or read-write access for this user.
5. To authenticate the user, select **Authentication**.
  - a. Select either the MD5 or the SHA1 protocol.
  - b. Enter the authentication key in the **Key** text field.
  - c. To provide encryption to the authentication session, select **Privacy**.
  - d. Select either the AES or the DES protocol.
  - e. Enter the encryption key in the **Key** text field.
6. Click **OK**.  
The newly added user account appears in the SNMP Users table.

## Modifying SNMP V3 Users

### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the **SNMP Users** area, select a checkbox for the user and click **Modify**.  
The Modify SNMP User dialog box appears. Add or change any of the following settings.
3. Select either read-only or read-write access for this user.
4. To authenticate the user, select **Authentication**.
  - a. Select either the MD5 or the SHA1 protocol.
  - b. Enter the authentication key in the Key text field.
  - c. To provide encryption to the authentication session, select **Privacy**.
  - d. Select either the AES or the DES protocol.
  - e. Enter the encryption key in the **Key** text field.
5. Click **OK**.  
The new settings for this user account appear in the SNMP Users table.

## Removing SNMP V3 Users

### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the SNMP Users area, select a checkbox for the user and click **Delete**.  
The Delete SNMP User dialog box appears.

---

**Note**

If the **Delete** button is disabled, the selected user is being used by one or more trap hosts. Delete the trap hosts and then delete the user.

---

3. Verify the user name to be deleted and click **OK**.
4. In the Delete SNMP User Status dialog box, click **Close**.  
The user account is removed from the SNMP Users table.

## Managing SNMP V3 and V2C Trap Hosts

The following topics describe how to manage SNMP trap hosts:

- ◆ [Creating SNMP V3 and V2C Trap Hosts on page 88](#)
- ◆ [Modifying SNMP V3 and V2C Trap Hosts on page 88](#)
- ◆ [Removing SNMP V3 and V2C Trap Hosts on page 88](#)

### Creating SNMP V3 and V2C Trap Hosts

#### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the SNMP V3 Trap Hosts or SNMP V2C Trap Hosts area, click **Create**.  
The Create SNMP [V3 or V2C] Trap Hosts dialog box appears.
3. In the **Host** text field, enter the IP address or domain name of the SNMP Host where traps will be sent.
4. In the **Port** text field, enter the port number for sending traps (port 162 is a common port).
5. Select the user (SNMP V3) or the community (SNMP V2C) from the drop-down menu.  
Alternately, select Create New User (SNMP V3) to add an SNMP user, or Create New Community (SNMP V2C) to add an SNMP community from the drop-down menu.
6. Click **OK**.

### Modifying SNMP V3 and V2C Trap Hosts

#### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the **SNMP V3 Trap Hosts** or **SNMP V2C Trap Hosts** area, select a Trap Host entry, and click **Modify**.  
The Modify SNMP [V3 or V2C] Trap Hosts dialog box appears. Modify any of the following items.
3. In the **Port** text field, enter the port number for sending traps (port 162 is a common port).
4. Select the user (SNMP V3) or the community (SNMP V2C) from the drop-down menu.
5. Click **OK**.

### Removing SNMP V3 and V2C Trap Hosts

#### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the **Trap Hosts** area (either for V3 or V2C, select a checkbox for the trap host and click **Delete**.



- The Delete SNMP [V3 or V2C] Trap Hosts dialog box appears.
3. Verify the host name to be deleted and click **OK**.
  4. In the Delete SNMP [V3 or V2C] Trap Hosts Status dialog box, click **Close**.
- The trap host entry is removed from the **Trap Hosts** table.

## Managing SNMP V2C Communities

---

### Note

The Community string is sent in cleartext and is very easy to intercept. If this occurs, the interceptor can retrieve information from devices on your network, modify their configuration, and possibly shut them down. Using the SNMP V3 Users configuration instead provides authentication and encryption to avoid this.

---

The following topics describe how to manage SNMP V2C communities:

- ◆ [Creating SNMP V2C Communities on page 89](#)
- ◆ [Modifying SNMP V2C Communities on page 89](#)
- ◆ [Deleting SNMP V2C Communities on page 90](#)

### Creating SNMP V2C Communities

#### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the Communities area, click **Create**.  
The Create SNMP V2C Community dialog box appears.
3. In the **Community** box, enter the name of a community who will have access to the agent on the Data Domain system.
4. Select either read-only or read-write access for this community.
5. In the Hosts area, select the checkbox of a host in the list, or:
  - a. Click **+** to add a host.  
The Host dialog box appears.
  - b. In the **Host** text field, enter the IP address or domain name of the host.
  - c. Click **OK**.  
The Host is added to the host list.
6. Click **OK**.  
The new community entry appears in the **Communities** table.

### Modifying SNMP V2C Communities

#### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the Communities area, select a checkbox for the community and click **Modify**.  
The Modify SNMP V2C Community dialog box appears. Add or change any of the following settings.
3. Select either read-only or read-write access for this community.
4. In the Hosts area, select the checkbox of a new host in the list, or:

- a. Click **+** to add a host.  
The Host dialog box appears.
  - b. In the **Host** text field, enter the IP address or domain name of the host.
  - c. Click **OK**.  
The Host is added to the host list.
5. Click **OK**.  
The modified community entry appears in the Communities table.

## Deleting SNMP V2C Communities

### Procedure

1. From the **System Settings > General Configuration** page, click **SNMP**.
2. In the **Communities** area, select a checkbox for the community and click **Delete**.  
The Delete SNMP V2C Communities dialog box appears.

---

### Note

If the **Delete** button is disabled, the selected community is being used by one or more trap hosts. Delete the trap hosts and then delete the community.

---

3. Verify the community name to be deleted and click **OK**.
4. In the Delete SNMP V2C Communities Status dialog box, click **Close**. The community entry is removed from the Communities table.

## Managing Reporting and Logging

The Data Domain system provides several types of reporting and logging, as described in the following sections:

- ◆ [Managing Autosupport Reporting on page 90](#)
- ◆ [Managing Support Bundles on page 92](#)
- ◆ [Managing Log Files on page 93](#)

### Managing Autosupport Reporting

The Autosupport feature sends to Data Domain Support a daily report that shows system identification information and consolidated output from a number of Data Domain system commands and entries from various log files. At the end of the report, extensive and detailed internal statistics and information are included to aid Data Domain in debugging system problems.

The time the email is sent can be scheduled; the default time is 06.00 A.M.

The procedures for managing autosupport reporting include:

- ◆ [Adding to the Autosupport Report Email List on page 90](#)
- ◆ [Reviewing Generated Autosupport Reports on page 91](#)
- ◆ [Verifying that Your Support/Alert Emails Are Received on page 91](#)

#### Adding to the Autosupport Report Email List

By default, autosupport reports are enabled and sent daily to Data Domain Customer Support. You may wish to add additional email addresses as recipients of autosupport

reports. To add to the autosupport mailing list, see [Configuring the Autosupport Mailing List on page 81](#).

## Reviewing Generated Autosupport Reports

### Procedure

1. Select the system to be checked in the navigation panel.
2. Select **Maintenance** › **Support** › **Autosupport**.

The Autosupport Reports page shows the autosupport report file name and file size, and the date the report was generated. Reports are automatically named. The most current report is autosupport, the previous day is autosupport.1, and the number increments as the reports move back in time.

3. Click the file name link to view the report using a text editor. If doing so is required by your browser, download the file first.

## Verifying that Your Support/Alert Emails Are Received

### Procedure

1. Open your web browser and enter your Data Domain system's IP address in the browser's address box. Wait for DD System Manager to display the login screen.
2. Enter your user name and password, and click **Login**.
3. Do one of the following:
  - If you used the DD System Manager Configuration Wizard, go to [Testing the Alerts Email List on page 81](#).
  - If you used the CLI Configuration Wizard, do the following:
    - a. Click **Maintenance** › **System**, and in the **More Tasks** list box, select **Launch Configuration Wizard**.
    - b. Skip each of the modules by clicking **No** or **Next** until you reach **System Settings**. Click **Yes** to configure system settings.
    - c. Verify the Administrator information is correct. Click **Next**.
    - d. The Email information in the Email/Location section must be correct. Verify the name of the mail server to be used to send outgoing Alert and Autosupport emails to recipients.
    - e. Verify that the **Send Alert Notification Emails to Data Domain** is selected.
    - f. Verify that the **Send Vendor Support Notification Emails to Data Domain** is selected. The address, which cannot be changed, is `autosupport@autosupport.datadomain.com`. This notification status is enabled by default. The mail server location is for your information only.
    - g. Click **Next** and verify that the information you entered is correct in the summary. If not, go back to the sections that need changing and edit them.
    - h. After approving the updated summary, click **Submit**.
    - i. Click **Next** to exit each of the remaining modules. Exit the Configuration Wizard.
4. Click **Maintenance** › **Support** › **Autosupport**. In the Vendor Support area, verify that the Notification Status for vendors is Enabled. You cannot change the email address.
5. Click **Status** › **Alerts** › **Notification**. In the **Alerts** area, verify that **default** is selected as a group. Verify that the Subscribers Email List contains the administrator and `autosupport-alert@autosupport.datadomain.com` addresses.

## Results

You can create other groups; add, modify, and delete email subscribers to groups; and set class attributes. Class attributes refer to the name of a class, such as Hardware or File System. You set the severity of ranking, such as Warning or Critical, that is to trigger an alert for each class attribute.

## Managing Support Bundles

The following tasks are used to manage support bundles:

- ◆ [Generating a Support Bundle on page 92](#)
- ◆ [Viewing the Support Bundles List on page 92](#)

### Generating a Support Bundle

When troubleshooting problems, Data Domain Customer Support may ask for a support bundle, which is a tar-g-zipped selection of log files with a README file that includes identifying autosupport headers. To create a support bundle, use the following procedure:

#### Procedure

1. Select a managed system in the navigation panel.
2. Select **Maintenance** > **Support**.
3. Select **More Tasks** > **Generate Support Bundle**.

---

#### Note

The system supports a maximum of 10 support bundles. If you attempt to generate an 11th support bundle, DD System Manager displays a warning that prompts you to click **OK** or **Cancel**. If you click **OK** to continue, the oldest support bundle, which is listed in the warning message, is deleted.

---

4. Click the link to download the bundle.
5. Email the file to Data Domain support at [support@datadomain.com](mailto:support@datadomain.com).

---

#### Note

If the bundle is too large to be emailed, use the EMC support site to upload the bundle. (Go to <https://support.emc.com>.)

---

### Viewing the Support Bundles List

#### Procedure

1. Select a managed system in the navigation panel.
2. Select **Maintenance** > **Support** > **Support Bundles**.

The Support Bundles list appears.

Listed are the support bundle file name, file size, and date the bundle was generated. Bundles are automatically named, where the most current bundle is `support-bundle.tar.gz`, and the previous bundle is `support-bundle.tar.gz.1`. The number increments as the reports move back in time.

3. Click the file name link and select a gz/tar decompression tool to view the ASCII contents of the bundle.

## Managing Log Files

The Data Domain system logs a system status message every hour. Log files can be bundled and sent to Data Domain Support to provide the detailed system information that aids in troubleshooting any system issues that may arise.

The Data Domain system log file entries contain messages from the alerts feature, autosupport reports, and general system messages. The log directory is `/ddvar/log`. Log messages from CIFS subsystem are logged only in `debug/cifs/cifs.log`. Examine this log file for CIFS issues.

---

### Note

Files in the `/ddvar` directory can be deleted based on the directory-level permissions. The directory must be writable in order for you to delete files within it.

---

Every Sunday at 0:45 a.m., the Data Domain system automatically opens new log files and renames the previous files with an appended number of 1 (one) through 9, such as `messages.1`. Each numbered file is rolled to the next number each week. For example, at the second week, the file `messages.1` is rolled to `messages.2`. If a file `messages.2` already existed, it rolls to `messages.3`. An existing `messages.9` is deleted when `messages.8` rolls to `messages.9`.

The procedures for working with log files include:

- ◆ [Viewing Log Files in DD System Manager on page 93](#)
- ◆ [Viewing the Log File List in the CLI on page 93](#)
- ◆ [Displaying a Log File in the CLI on page 94](#)
- ◆ [Understanding Log Messages on page 94](#)
- ◆ [Saving a Copy of Log Files on page 95](#)
- ◆ [Sending Log Messages to Another System on page 96](#)

### Viewing Log Files in DD System Manager

This topic describes how to display the log file list on a managed system and open log files.

#### Procedure

1. In the navigation panel, expand **DD Network** and select a managed system.
2. Select **Maintenance** > **Logs**.

The Logs list displays log file names and the size and generation date for each log file. Log files are automatically named. For more information on log files, see [Managing Log Files on page 93](#).

3. Click a log file name to view its contents. You may be prompted to select an application, such as Notepad.exe, to open the file.

### Viewing the Log File List in the CLI

To view the log files in the CLI, type:

```
log list
```

The basic log files are:

- ◆ `messages`—The system log, generated from Data Domain system actions and general system operations.

- ◆ `space.log`—Messages about disk space use by Data Domain system components and data storage, and messages from the cleaning process. A space use message is generated every hour. Each time the cleaning process runs, it creates about 100 messages. All the messages are in comma-separated format with tags that you can use to separate out the disk space or cleaning messages. You can use third-party software to analyze either set of messages. The tags are:
  - `CLEAN` for data lines from cleaning operations.
  - `CLEAN_HEADER` for lines that contain headers for the cleaning operations data lines.
  - `REPL` for data lines from replication operations.
  - `REPL_HEADER` for lines that contain headers for the replication data lines.
  - `SPACE` for disk space data lines.
  - `SPACE_HEADER` for lines that contain headers for the disk space data lines.
- ◆ `ssi_request`—Messages from DD System Manager when users connect with HTTPS.
- ◆ `debug/cifs/cifs.log`—CIFS-related activity.

## Displaying a Log File in the CLI

To view a log file in the CLI, use the `log view` command to view a file in the list (see previous section to list log files). With no argument, the command displays the current messages file. When viewing the log, use the up and down arrows to scroll through the file; use the `q` key to quit; enter a slash character (`/`) and a pattern to search through the file.

```
log view file_name
```

The display of the messages file is similar to the following. The last message in the example is an hourly system status message that the Data Domain system generates automatically. The message reports system uptime, the amount of data stored, NFS operations, and the amount of disk space used for data storage (%). The hourly messages go to the system log and to the serial console if one is attached.

```
# log view
Jun 27 12:11:33 localhost rpc.mountd: authenticated unmount
request from perfsun-g.datadomain.com:668 for /ddr/coll/segfs
(/ddr/coll/segfs)

Jun 27 12:28:54 localhost sshd(pam_unix)[998]: session opened
for user jsmith10 by (uid=0)

Jun 27 13:00:00 localhost logger: at 1:00pm up 3 days, 3:42,
52324 NFS ops, 84763 GiB data col. (1%)
```

### Note

GiB = Gibibytes = the binary equivalent of Gigabytes.

## Understanding Log Messages

In the log file is text similar to:

```
Jan 31 10:28:11 syrah19 bootbin: NOTICE: MSG-SMTOOL-00006: No
replication throttle schedules found: setting throttle to
unlimited.
```

The components of the message are:

```
DateTime Host Process [PID]: Severity: MSG-Module-MessageID: Message
```

Severity levels, in descending order, are Emergency, Alert, Critical, Error, Warning, Notice, Info, Debug.

### Procedure

1. View the log file. This can be done with the command `log view message` or the command `log view`, or from DD System Manager (see [Displaying a Log File in the CLI on page 94](#)).
2. Look up error messages in the Error Message Catalog for your DD OS version. To locate the Error Message Catalog, go to the EMC Online Support website at <https://support.emc.com>, enter *Error Message Catalog* in the search box, and click the search button.

The error message description looks similar to the following display.

```
ID: MSG-SMTOOL-00006 - Severity: NOTICE - Audience: customer
```

```
Message: No replication throttle schedules found: setting throttle to unlimited.
```

```
Description: The restorer cannot find a replication throttle schedule. Replication is running with throttle set to unlimited.
```

```
Action: To set a replication throttle schedule, run the replication throttle add command.
```

3. To resolve an issue, do the recommended action.

Based on the example message description, one could run the `replication throttle add` command to set the throttle.

## Saving a Copy of Log Files

To save a copy of log files, use FTP to move the files to another machine.

### Procedure

1. On the Data Domain system, use the `adminaccess show ftp` command to see whether FTP service is enabled. If the service is disabled, use the command `adminaccess enable ftp`.
2. On the Data Domain system, use the `adminaccess show ftp` command to see that the FTP access list has the IP address of your remote machine or a class-C address that includes your remote machine. If the address is not in the list, use the command `adminaccess add ftp ipaddr`.
3. On the remote machine, open a web browser.
4. In the **Address** box at the top of the web browser, use FTP to access the Data Domain system. For example:

```
ftp://Data Domain system_name.yourcompany.com/
```

**Note**

Some web browsers do not automatically ask for a login if a machine does not accept anonymous logins. In that case, add a user name and password to the FTP line. For example: `ftp://sysadmin:your-pw@Data Domain system_name.yourcompany.com/`

5. At the login pop-up, log into the Data Domain system as user `sysadmin`.
6. On the Data Domain system, you are in the directory just above the log directory. Open the log directory to list the messages files.
7. Copy the file that you want to save. Right-click the file icon and select **Copy To Folder** from the menu. Choose a location for the file copy.
8. If you want the FTP service disabled on the Data Domain system, after completing the file copy, use SSH to log into the Data Domain system as `sysadmin` and invoke the command `adminaccess disable ftp`.

## Sending Log Messages to Another System

Some log messages can be sent from the Data Domain system to other systems. DD OS uses `syslog` to publish log messages to remote systems.

A Data Domain system exports the following facility.priority selectors for log files. For information on managing the selectors and receiving messages on a third-party system, see your vendor-supplied documentation for the receiving system.

- ◆ `*.notice`—Sends all messages at the notice priority and higher.
- ◆ `*.alert`—Sends all messages at the alert priority and higher (alerts are included in `*.notice`).
- ◆ `kern.*`—Sends all kernel messages (kern.info log files).
- ◆ `local7.*`—Sends all messages from system startups (boot.log files).

The `log host` commands manage the process of sending log messages to another system.

The following topics describe how to manage sending messages to other systems:

- ◆ [Adding a Receiver Host on page 96](#)
- ◆ [Removing a Receiver Host on page 96](#)
- ◆ [Enabling Log Message Sending on page 97](#)
- ◆ [Disabling Log Message Sending on page 97](#)
- ◆ [Disabling Log Message Sending on page 97](#)

### Adding a Receiver Host

To add a system to the list that receives Data Domain system log messages, use the `log host add` command.

```
log host add host
```

For example, the following command adds the system `log-server` to the hosts that receive log messages:

```
log host add log-server
```

### Removing a Receiver Host

To remove a system from the list that receives Data Domain system log messages, use the command: `log host del host`.



For example, the following command removes the system *log-server* from the hosts that receive log messages: `log host del log-server`.

### Enabling Log Message Sending

To enable sending log messages to other systems, use the `log host enable` command.

```
log host enable
```

### Disabling Log Message Sending

To disable sending log messages to other systems, use the `log host disable` command.

```
log host disable
```

### Displaying the Log Message Sending Configuration

To display the list of systems that receive log messages and logging status (enabled or disabled), use the `log host show` command. The output is similar to the following:

```
# log host show
Remote logging is enabled.
Remote logging hosts
    log-server
```

## Managing Remote System Power with IPMI

When the DD System Manager host and a remote system both support Intelligent Platform Management Interface (IPMI), you can use DD System Manager to view the power status on a remote system and power up, power down, or power cycle the remote system. Similarly, you can use the DD OS CLI to manage remote system power with IPMI.

IPMI runs independently of DD OS and allows an IPMI user to manage system power as long as the remote system is connected to a power source and a network. An IP network connection is required between a management host and a remote system. When properly configured and connected, IPMI management eliminates the need to be physically present to power on or power off a remote system.

#### NOTICE

IPMI power removal is provided for emergency situations during which attempts to shut down power using DD OS commands fail. IPMI power removal simply removes power to the system, it does not perform an orderly shutdown of the DD OS file system. The proper way to remove and reapply power is to use the DD OS `system reboot` command. The proper way to remove system power is to use the DD OS `system poweroff` command and wait for the command to properly shut down the file system.

The following sections describe IPMI use:

- ◆ [Getting Started with IPMI on page 98](#)
- ◆ [Configuring IPMI for a Managed System on page 98](#)
- ◆ [Logging Into a Remote System for IPMI Power Management on page 102](#)
- ◆ [Managing Remote System Power After Login on page 102](#)

## Getting Started with IPMI

To use IPMI to manage power for a remote system, both the management system and the remote system to be managed must support IPMI, and the remote system must be configured to support IPMI user management.

---

### Note

If the ability to view the boot sequence of a remote Data Domain system is required using Serial Over LAN (SOL), see the ipmi chapter of the *EMC Data Domain Operating System Command Reference Guide*.

---

## IPMI Limitations

This topic lists the Data Domain systems that do not support IPMI.

IPMI is supported on all systems supported by this release except the following systems: DD140, DD610, and DD630.

## Terminology

The following terms are used to describe the status and responsibility of systems running IPMI:

- ◆ **Management host**—The management host is the system you use to manage power on a remote system. When you use DD System Manager, the management host is the system you select in the navigation panel. When you use the Data Domain OS CLI, the management host is the system on which you are using the CLI.
- ◆ **Remote system**—A remote system is the target system for which you want to manage power. DD System Manager supports management of remote systems that are and are not managed by DD System Manager.
- ◆ **Managed remote system**—A managed remote system is a system that appears in the DD System Manager navigation panel. You can use DD System Manager to configure IPMI on a managed system, and then you can use DD System Manager to manage power on that system.
- ◆ **Nonmanaged remote system**—A nonmanaged remote system is a system that does not appear in the DD System Manager navigation panel. You can use DD System Manager to manage power on a nonmanaged system, but the nonmanaged system must be configured to support IPMI users before you can manage it. If a remote system is managed by another DD System Manager host, you can use the current DD System Manager host to manage it, but you manage it as a nonmanaged remote system because it is not configured in the current DD System Manager host.

---

### Note

You can use the Data Domain OS CLI to enable IPMI and configure IPMI users on a nonmanaged remote host.

---

## Configuring IPMI for a Managed System

You can use DD System Manager to configure IPMI for any managed system that appears in the navigation panel. After you configure IPMI for a managed system, you can quickly begin remote system power management by selecting the managed system and entering the IPMI user information. Power management is easier for managed systems because you have control over the remote system configuration and you do not have to remember the host name or IP address for the remote system.

**Note**

You can configure IPMI for non-managed systems using the Data Domain OS CLI as described in the ipmi chapter of the *EMC Data Domain Operating System Command Reference Guide*.

The following topics describe how to configure IPMI for a managed system:

- ◆ [Viewing the IPMI Configuration for a Managed System on page 99](#)
- ◆ [Configuring an IPMI Port on a Managed System on page 100](#)
- ◆ [Managing IPMI Users for a Managed System on page 100](#)
- ◆ [Enabling and Disabling IPMI on a Managed System on page 102](#)

**Note**

If a system does not have the correct hardware or software to support IPMI, a message to that effect is generated when navigating to the configuration page.

## Viewing the IPMI Configuration for a Managed System

This topic describes how to view the IPMI configuration for a managed system.

**Procedure**

1. In the navigation panel, select the managed system to view.
2. Select **Maintenance** > **IPMI**.

The IPMI Configuration area shows the IPMI configuration for the managed system. The Network Ports table lists the ports on which IPMI can be enabled and configured. The IPMI Users table lists the IPMI users that can access the managed system.

**Table 18** Network Ports

Item	Description
Port	The logical name for a port that supports IPMI communications. See <a href="#">Configuring an IPMI Port on a Managed System on page 100</a> for additional information on port names and port configuration details.
Enabled	Whether the port is enabled for IPMI (Yes or No). See <a href="#">Enabling and Disabling IPMI on a Managed System on page 102</a> for details on how to change the status.
DHCP	Whether the port uses DHCP to set its IP address (Yes or No).
MAC Address	The hardware MAC address for the port.
IP Address	The port IP address.
Netmask	The subnet mask for the port.
Gateway	The gateway IP address for the port.

**Table 19** IPMI Users

Item	Description
User Name	The name of a user with authority to power manage the remote system. See <a href="#">Adding an IPMI User on page 101</a> for configuration details.

## Configuring an IPMI Port on a Managed System

When you configure a port for a managed system, you select the port from a network ports list and specify the IP configuration parameters for that port. The selection of ports displayed depends on the Data Domain system selected. Some systems support one or more dedicated ports, which can be used only for IPMI traffic. Other systems support ports that can be used for both IPMI traffic and all IP traffic supported by the physical interfaces in the **Hardware > Network > Interfaces** view. Shared ports are not provided on systems that provide dedicated IPMI ports.

The port names in the IPMI Network Ports list use the prefix `bmc`, which represents baseboard management controller. To determine if a port is a dedicated port or shared port, compare the rest of the port name with the ports in the network interface list. If the rest of the IPMI port name matches an interface in the network interface list, the port is a shared port. If the rest of the IPMI port name is different from the names in the network interface list, the port is a dedicated IPMI port.

### Note

DD4200, DD4500, and DD7200 systems are an exception to the naming ruled described earlier. On these systems, IPMI port, `bmc0a`, corresponds to shared port `ethMa` in the network interface list. EMC recommends that the shared port `ethMa` be reserved for IPMI traffic and system management traffic (using protocols such as HTTP, Telnet, and SSH). Backup data traffic should be directed to other ports.

When IPMI and nonIPMI IP traffic share an Ethernet port, EMC recommends that you do not use the link aggregation feature on the shared interface because link state changes can interfere with IPMI connectivity.

### Procedure

1. Select the managed system in the navigation panel.
2. Select **Maintenance > IPMI**.
3. In the **Network Ports** table, select a port to configure.
4. Above the **Network Ports** table, click **Configure**.  
The Configure Port dialog box appears.
5. Choose how network address information is assigned:
  - To collect the IP address, netmask, and gateway configuration from a DHCP server, select **Dynamic (DHCP)**.
  - To manually define the network configuration, select **Static (Manual)** and enter the IP address, netmask, and gateway address.
6. Click **Apply**.

## Managing IPMI Users for a Managed System

When you create an IPMI user with DD System Manager, you define a user name and password for users who can manage power for the selected managed system. After you

create an IPMI user, you can change the password for that user or delete the user. Each managed system has its own list of IPMI users. To give an IPMI user the authority to manage power on all managed systems, you must add that user to each of the managed systems.

---

### Note

The IPMI user list for each managed system is separate from the DD System Manager lists for administrator access and local users. Administrators and local users do not inherit any authorization for IPMI power management.

---

The following topics describe how to manage IPMI users for a managed system:

- ◆ [Adding an IPMI User on page 101](#)
- ◆ [Modifying an IPMI User Password on page 101](#)
- ◆ [Removing an IPMI User on page 101](#)

## Adding an IPMI User

This topic describes how to add an IPMI user to a managed system.

### Procedure

1. Expand **DD Network** in the navigation panel and select a managed system.
2. Select **Maintenance** › **IPMI**.
3. Above the IPMI Users table, click **Add**.
4. In the Add User dialog box, type the user name (16 or less characters) and password in the appropriate boxes (reenter the password in the **Verify Password** box).
5. Click **Create**.

The user entry appears in the **IPMI Users** table.

## Modifying an IPMI User Password

This topic describes how to modify the password for an IPMI user on a managed system.

### Procedure

1. Expand **DD Network** in the navigation panel and select a managed system.
2. Select **Maintenance** › **IPMI**.
3. In the IPMI Users table, select a user, and click **Change Password**.
4. In the Change Password dialog box, type the password in the appropriate text box and reenter the password in the **Verify Password** box.
5. Click **Update**.

## Removing an IPMI User

This topic describes how to delete an IPMI user from a managed system.

### Procedure

1. Expand **DD Network** in the navigation panel and select a managed system.
2. Select **Maintenance** › **IPMI**.
3. In the IPMI Users list, select a user and click **Delete**.
4. In the Delete User dialog box, click **OK** to verify user deletion.

## Enabling and Disabling IPMI on a Managed System

You can enable or disable IPMI on each IPMI network port. When you disable IPMI on an IPMI port that supports other Ethernet traffic, only IPMI traffic is disabled on that port. To change the status of a managed system's IPMI network port on a remote system:

### Procedure

- ◆ Enable a disabled IPMI network port by selecting the network port in the **Network Ports** table, and clicking **Enable**.
- ◆ Disable an enabled IPMI network port by selecting the network port in the **Network Ports** table, and clicking **Disable**.

## Logging Into a Remote System for IPMI Power Management

When you log into a remote system for IPMI power management, you specify the system to which you want to connect and the username and password for authentication.

### Procedure

1. Expand **DD Network** in the navigation panel and select a managed system that is different from the remote system you want to manage.

When you select a system in the navigation panel, it becomes the management system and you cannot use IPMI to manage the management system. You can use a management system to manage remote systems only.

2. Select **Maintenance > IPMI**.
3. In the IPMI Power Management area, click **Login to Remote System**.

The IPMI Power Management dialog box appears.

4. In the Target System area, select the type of remote system and identify the system:
  - **Managed System**—Select a system in the **Managed** list box.
  - **Nonmanaged System**—Enter the IP address or hostname of the remote system in the **Another System** box.

SSH access to remote systems is not supported. If no system appears in the Managed list box, configure managed systems as described in [Configuring IPMI for a Managed System on page 98](#). Nonmanaged systems must be configured to support power management as described in the ipmi chapter of the *EMC Data Domain Operating System Command Reference Guide*.

5. Enter the IPMI user name and password for the remote system.
6. Click **Connect**.

The IPMI Power Management dialog box appears.

## Managing Remote System Power After Login

After you log in to a remote system with DD System Manager, you can use IPMI to manage system power as described in the following topics:

- ◆ [Viewing the Power Status on page 103](#)
- ◆ [Changing the Power Status on page 103](#)

## Viewing the Power Status

After you log in to a remote system for IPMI power management, the IPMI Power Management dialog box appears and shows the target system identification and the current power status. The Status area always shows the current status.

---

### Note

The Refresh icon (the blue arrows) next to the status can be used to refresh the configuration status (for example, if the IPMI IP address or user configuration were changed within the last 15 minutes using the CLI commands).

---

Click **Done** to close the IPMI Power Management dialog box.

## Changing the Power Status

After you log in to a remote system for IPMI power management, the IPMI Power Management dialog box appears and displays buttons to power up, power down, and power cycle the remote system. The buttons that are active depend on the current power status of the remote system. The available buttons are:

- ◆ **Power Up**—Appears when the remote system is powered off. Click this button to power up the remote system.
- ◆ **Power Down**—Appears when the remote system is powered on. Click this button to power down the remote system.
- ◆ **Power Cycle**—Appears when the remote system is powered on. Click this button to power cycle the remote system.
- ◆ **Manage Another System**—Click this button to log into another remote system for IPMI power management.
- ◆ **Done**—Click to close the IPMI Power Management dialog box.

### NOTICE

The IPMI Power Down feature does not perform an orderly shutdown of the DD OS. This option can be used if the DD OS hangs and cannot be used to gracefully shutdown a system.

---





# CHAPTER 4

## Monitoring Data Domain Systems

This chapter includes:

- ◆ [About Monitoring Data Domain Systems](#)..... 106
- ◆ [Monitoring Using the DD Network Summary](#)..... 106
- ◆ [Monitoring a Single System](#)..... 108
- ◆ [About the Fibre Channel View](#)..... 109
- ◆ [Monitoring Chassis Status](#)..... 110
- ◆ [Working with Alerts](#)..... 113
- ◆ [Viewing Active Users](#)..... 119
- ◆ [Viewing System Statistics](#)..... 119
- ◆ [Working with Reports](#)..... 120
- ◆ [Viewing the Task Log](#)..... 126

## About Monitoring Data Domain Systems

DD System Manager provides a composite view of important statistics for a group of Data Domain systems, and detailed status for a single system and its components.

You can monitor Data Domain system operation with a variety of DD System Manager tools: reporting tools that automatically send emails containing status and alerts, log files that contain a record of important system events, and SNMP monitoring using third-party SNMP managers.

Automatic logging and reporting tools that provide system status to Customer Support and designated email recipients are important in monitoring system operation. Their setup and use are described in this chapter.

## Monitoring Using the DD Network Summary

The DD Network Summary presents key statistics about the health of managed Data Domain systems. The System Status, Space Usage, and Systems panes provide information to help you recognize problems immediately and determine which system has a problem.

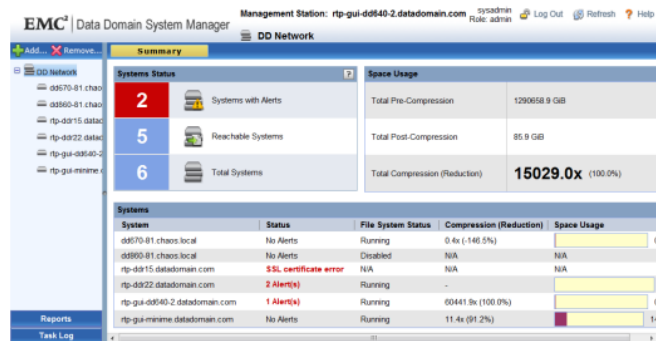
### Viewing the DD Network Status

#### Procedure

1. Select the DD Network icon in the navigation panel.
2. Click the Summary tab.

The DD Network Summary view appears. It presents a high-level view of important information for the systems in the network. The summary view displays summary statistics for all managed systems, space usage statistics, and status for individual systems.

Figure 3



### About the System Summary Statistics

The Systems Status list summarizes the following about the network.

Item	Description
Systems with Alerts	The number of systems with active alerts.

Item	Description
Reachable Systems	The total number of systems reporting to the DD System Manager. A system may not be reporting if: <ul style="list-style-type: none"> <li>the system is offline</li> <li>the network path to the system is down</li> <li>an SSL certificate error occurred</li> </ul>
Total Systems	The total number of managed systems configured on this DD System Manager system.

## About the Space Usage Statistics

The Space Usage list displays statistics for all managed systems.

Item	Description
Total Pre-Compression	The total amount of data sent to all managed systems by backup servers. This is the data written before compression.
Total Post-Compression	The total data amount for all systems after compression has been performed.
Average Compression (Reduction)	The average amount of compression as calculated for each individual system.

## About Individual-System Statistics

The Systems list summarizes important data for each of the managed systems.

Item	Description
System	The name of a managed system.
Status	If the system is reachable, No Alerts or <i>n</i> Alerts appears, where <i>n</i> is the number of active alerts. The status line changes to red with an active alert. If the system is not reachable, the status column displays one of the following: <ul style="list-style-type: none"> <li>Unknown</li> <li>Not reachable</li> <li>SSL certificate error</li> </ul>
File System Status	The status of the file system. Status can be: <ul style="list-style-type: none"> <li>Running</li> <li>Disabled</li> <li>N/A—The system is not reachable.</li> </ul>
Compression (Reduction)	The average amount of compression for the listed system or N/A if the system is not reachable.
Space Usage	A bar graph showing the disk space in use or N/A if the system is not reachable.

Select a system name in the Systems list to select the system in the navigation panel and display information on that system in the information panel.

## Monitoring a Single System

To help you proactively recognize trouble signs that keep a system from operating normally, DD System Manager presents system alerts, graphs, and logs. Procedures for working with these tools are provided in this section.

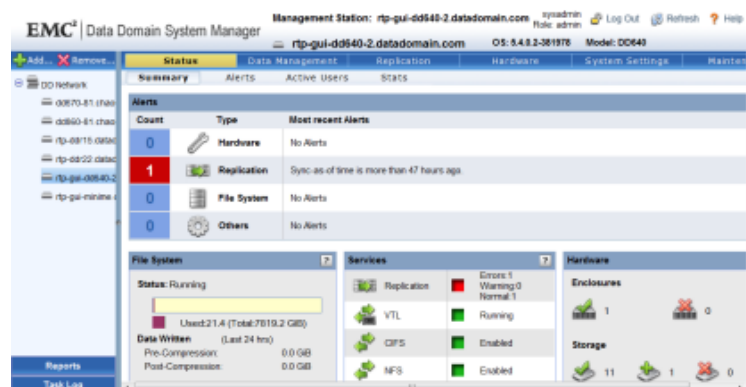
Sometimes, the system needs troubleshooting from Data Domain Customer Support. This section provides procedures for obtaining and sending system logs and reports.

## Viewing the System Status Summary

Expand **DD Network** in the navigation panel and select a system. If this is the first time you have selected this system during this DD System Manager session, the **Status > Summary** view appears. If a different view appears, select **Status > Summary**.

The Summary page shows important high-level information about the selected system. It displays alerts, file system statistics, service data, and hardware information. Click in one of the four display areas to display a DD System Manager view that provides more detailed information on the area topic.

Figure 4



### About the Alerts Summary

The Alerts summary shows the count, type, and the text of the most recent alerts in the system, for each subsystem (hardware, replication, file system, and others).

Click an alert to display the **Status > Alerts** view. See [Viewing Current Alerts on page 113](#).

### About the File System Summary

The File System summary shows file system statistics, including the operational status, compression factor, and data written statistics.

Click in the File System area to display the **Data Management > File System** view. See [Monitoring File System Usage on page 134](#).

### About the Services Summary

The Services summary presents the status of the system services, such as replication, VTL, CIFS, NFS, and DD Boost. The color-coded box shows the operational status (green for normal, yellow for warnings, or red for errors). The total numbers for warnings and errors are displayed as well.

Click a service to display additional information on that service:

- ◆ See [Checking Replication Status on page 322](#)
- ◆ See [Working with DD Virtual Tape Library on page 249](#)
- ◆ See [Monitoring CIFS Operation on page 220](#)
- ◆ See [View NFS Status on page 230](#)
- ◆ See [Monitoring DD Boost on page 245](#)

## About the Hardware Summary

The Hardware summary presents the status of the system hardware, such as disk drives and optional enclosures. The color-coded icons show the operational status (green for normal or red for degraded or failed). A count shows the number of enclosures, and the number of drives per condition (operational, spare, and failed).

---

### Note

Counts on the dashboards refer to the total number of errors, not the index number of the component exhibiting the error.

---

Click an icon to display the **Hardware > Storage** view for the hardware category; see [Managing System Storage on page 38](#)).

## Viewing System Details

### Procedure

1. Expand DD Network in the navigation panel and select a system.
2. Select **Maintenance > System**.

The System summary reports the model number of the system, the DD OS version, and the amount of time since the last reboot (System Uptime). On systems running DD OS 5.5.1 or DD OS 5.4.4 and later, the system serial number and chassis serial number are also displayed. On newer systems, such as DD4500 and DD7200, the system serial number is independent of the chassis serial number and remains the same during many types of maintenance events, including chassis replacements. On legacy systems, such as DD990 and earlier, the system serial number is set to the chassis serial number.

On systems running releases prior to DD OS 5.5.1 or DD OS 5.4.4, the chassis serial number appears.

## About the Fibre Channel View

The Fibre Channel view has two tabs: Physical Resources and Access Groups.

- ◆ [About the Physical Resources View on page 109](#)
- ◆ [About the Access Groups View on page 110](#)

## About the Physical Resources View

The **Hardware > Fibre Channel > Physical Resources** tab displays information about endpoints and initiators. For instructions on managing endpoints, see [About Endpoints on page 289](#). For instructions on managing initiators, see [About Initiators on page 294](#).

To display additional information on an endpoint, select a single endpoint in the Endpoints list, and the additional information appears in the Endpoint Details area. The following table describes the columns in the Endpoints list:

Item	Description
Endpoint	Name of the endpoint.
WWPN	World-Wide Port Name of the Fibre Channel port in the media server.
WWNN	World-Wide Node Name that is assigned to a node (an endpoint or device) in a Fibre Channel fabric.
Physical Port	The physical port number.
Enabled	The port operational state; either Enabled or Disabled.
Link Status	Either Online or Offline; that is, whether or not the port is up and capable of handling traffic.

The following table describes the columns in the Initiators list:

Item	Description
Name	Name of the initiator.
Service	Service support by the initiator, which is either VTL or DD Boost.
WWPN	World-Wide Port Name of the initiator.
WWNN	World-Wide Node Name of the initiator.
Vendor Name	Initiator's model.
Online Endpoints	Endpoints seen by this initiator. Displays <i>none</i> or <i>offline</i> if the initiator is not available.

## About the Access Groups View

The Access Group tab lists access groups by group name, service (VTL or DD Boost), endpoints, initiators, and number of devices.

Clicking a link to Configure DD Boost groups or VTL groups takes you to the DD Boost or VTL page where you can create an access group. See [Create Access Group on page 244](#) for DD Boost access groups. See [Working with Access Groups on page 283](#) for VTL access groups.

## Monitoring Chassis Status

The **Hardware > Chassis** view displays a block drawing of the chassis and its components—disks, fans, power supplies, NVRAM, CPUs, Memory, and so forth. The components that appear depend upon the system model.

The **Hardware > Chassis** view displays a block drawing of each enclosure in a system, including the chassis serial number and the enclosure status. On systems running DD OS 5.5.1 or DD OS 5.4.4 and later, the system serial number is also displayed. On newer systems, such as DD4500 and DD7200, the system serial number is independent of the chassis serial number and remains the same during many types of maintenance events,

including chassis replacements. On legacy systems, such as DD990 and earlier, the system serial number is set to the chassis serial number.

Within each block drawing are the enclosure components, such as disks, fans, power supplies, NVRAM, CPUs, and memory. The components that appear depend upon the system model.

### Procedure

1. Expand DD Network in the navigation panel and select a system.
2. Select **Hardware** > **Chassis**.

The Chassis view shows the system components. Enclosures for systems are shown below the chassis.

Components with problems show yellow (warning) or red (error); otherwise, the component displays OK.

3. Hover the cursor over a component to see detailed status.

The view includes information for:

- [Fans on page 111](#)
- [Temperature on page 111](#)
- [Power Supply on page 112](#)
- [PCI Slots on page 112](#)
- [NVRAM on page 112](#)

## Fans

Fans are numbered and correspond to their location in the chassis. The tooltip provides the following.

Item	Description
Description	The name of the fan.
Level	The current operating speed range (Low, Medium, High). The operating speed changes depending on the temperature inside the chassis.
Status	The health of the fan.

## Temperature

The tooltip shows temperature measurements for the CPUs, baseboard, midplane, and front panel of the chassis.

Item	Description
Description	The location within the chassis being measured. Components are dependent on the model. Some examples are: <ul style="list-style-type: none"> <li>• CPU 0 relative</li> <li>• CPU 1 relative</li> <li>• Baseboard</li> <li>• Mid-plane</li> </ul>

Item	Description
	<ul style="list-style-type: none"> <li>• Front panel</li> </ul>
C/F	The C/F column displays temperature in degrees Celsius and Fahrenheit. For CPUs (CPU <i>n</i> Relative), this column displays the number of degrees that each CPU is below the maximum allowable temperature and the actual temperature for the interior of the chassis (chassis ambient).
Status	Shows the temperature status: <ul style="list-style-type: none"> <li>• OK—The temperature is acceptable</li> <li>• Critical—The temperature is higher than the shutdown temperature.</li> <li>• Warning—The temperature is higher than the warning temperature (but lower than the shutdown temperature).</li> </ul>

Each Data Domain system features non-configurable temperature settings for warning messages and shutdown. If the overall temperature rises above the warning temperature setting, a warning message is generated. If the temperature reaches the shutdown temperature setting, the Data Domain system shuts down.

## Power Supply

The tooltips shows the status of the power supply (OK or DEGRADED if a power supply is absent or failed). You can also look at the back panel of the enclosure and check the LED for each power supply to identify those that need replacing.

## PCI Slots

The PCI Slots shown in the chassis view indicate the number of PCI slots and the numbers of each slot. There are no tooltips for the PCI slots.

## NVRAM

NVRAM shows information about the Non-Volatile RAM and the batteries.

Item	Description
Component	The component within the chassis being described: <ul style="list-style-type: none"> <li>• Memory Size</li> <li>• Battery number (The number of batteries depends on the system type.)</li> <li>• Current slot number for NVRAM</li> </ul>
Value	<ul style="list-style-type: none"> <li>• Memory Size—size in MBs</li> <li>• Battery number—Percent charged, status (Enabled/Disabled)</li> </ul>



## Working with Alerts

During normal operation, a managed system may produce warnings or encounter failures whereby administrators must be informed immediately. This communication is performed by means of an alert.

Alerts are sent out to designated individuals or groups so that appropriate actions can be taken promptly.

Alerts are sent as email (immediately via the notification settings or cumulatively as Daily Alert Summary email) and logged on the Current Alerts view. Alerts are also sent as SNMP traps. See the *MIB Quick Reference Guide* or the SNMP MIB for the full list of traps.

The Alerts views present lists of current and historical system alerts, and clicking on an alert shows its details. The Alerts view also allows you to configure alert notification settings and set when and to whom daily alert summaries are sent.

To view managed system alerts, do the following:

### Procedure

1. Expand **DD Network** in the navigation panel and select a managed system.
2. Select **Status > Alerts**.

### Results

The Alerts view appears and displays tabs for:

- ◆ [Working with the Current Alerts Tab on page 113](#)
- ◆ [Working with the Alerts History Tab on page 115](#)
- ◆ [Working with the Notification View on page 116](#)
- ◆ [Configuring the Daily Alert Summary Distribution List on page 118](#)

## Working with the Current Alerts Tab

The **Current Alerts** tab lists the alerts on the selected system that have not been corrected or manually cleared. A total of the current alerts appears below the list.

You can perform the following tasks on the **Current Alert** tab:

- ◆ [Viewing Current Alerts on page 113](#)
- ◆ [Filtering Current Alerts on page 114](#)
- ◆ [Clearing a Current Alert on page 114](#)

## Viewing Current Alerts

The Current Alerts tab lists the following information for each current alert:

Item	Description
ID	A unique numerical identifier for the alert.
Message	The alert message text.
Severity	The level of seriousness of the alert. For example, warning, critical, info, or emergency.
Date	The time and date the alert occurred.
Class	The subsystem where the alert occurred.

Item	Description
Object	The physical component where the alert is occurring.

Click an alert in the list to display additional information in the **Details** area, which displays the following:

Item	Description
Alert ID	A unique numerical identifier for the alert.
Name	A textual identifier for the alert.
Message	The alert message text.
Severity	The level of seriousness of the alert. For example, warning, critical, info, emergency.
Class	The subsystem and device where the alert occurred.
Date	The time and date the alert occurred.
Object ID	The physical component where the alert is occurring.
Event ID	An event identifier.
Description	More descriptive information about the alert.
Action	A suggestion to remedy the alert.
SNMP OID	SNMP object ID.

## Filtering Current Alerts

This topic describes how to limit the list of displayed alerts based on the alert severity and class.

### Procedure

1. Expand **DD Network** in the navigation panel and select a managed system.
2. In the Filter By area, select a **Severity** and **Class** to expose only alerts that pertain to those choices.
3. Click **Update**.

All alerts not matching the Severity and Class are removed from the list.

To remove filtering and return to the full listing of current alerts, click **Reset**.

## Clearing a Current Alert

An alert is automatically removed from the Current Alerts list when the underlying situation is corrected or when manually cleared. For example, an alert about a fan failure is removed when the fan is replaced with a working unit. This topic describes how to manually clear an alert and remove it from the Current Alerts list.

### Procedure

1. Expand **DD Network** in the navigation panel and select a managed system.
2. Select **Status > Alerts > Current Alerts**.
3. Select the checkbox of the alert in the list.
4. Click **Clear**.

The alert is moved to the Alerts History list.

## Working with the Alerts History Tab

The Alerts History tab lists cleared alert messages with the most recent alert listed first. This page can be used to see how healthy a managed system has been in the past and to track the actions that were taken to keep the system healthy. It is useful in spotting trends and avoiding problems.

You can perform the following tasks on the **Alerts History** tab:

- ◆ [View Alerts History on page 115](#)
- ◆ [Filter Alerts History on page 115](#)

### View Alerts History

The alert list displays the historical alerts and the following information for each alert:

Item	Description
ID	A unique numerical identifier for the alert.
Message	The alert message text.
Severity	The level of seriousness of the alert. For example, warning, critical, info, or emergency.
Date	The time and date the alert occurred.
Class	The subsystem where the alert occurred.
Object	The physical component where the alert is occurring.
Status	The current disposition of the alert (for example, Posted or Cleared).

Click an alert to display the following information in the Details area:

Item	Description
Alert ID	A unique numerical identifier for the alert.
Name	A textual identifier for the alert.
Message	The alert message text.
Severity	The level of seriousness of the alert. For example, warning, critical, info, emergency,
Class	The subsystem and device where the alert occurred.
Date	The time and date the alert occurred.
Object ID	The physical component where the alert is occurring.
Event ID	An event identifier.
Additional Information	More descriptive information about the alert.
Type	The type of alert.
Status	The status of the alert.
Clear By	The user name that cleared the alert.

### Filter Alerts History

The alert history list can be rearranged with the following options:

- ◆ Click any diamond in a column heading to sort the alert list according to the entries in that column. Click again to reverse the sort order.
- ◆ Use the Filter By options to filter using the Severity, Date, Class, and Status options, then Click **Update**.
- ◆ Use the Other option in the Date list to set a specific start and end date for when alerts were closed using the calendar exposed with the calendar icon.
- ◆ Click **Reset** to return to the default list display, where the latest alert is listed first.

## Working with the Notification View

The Notification view lists the group of email recipients who receive alert notifications and allows you to view and configure the notification groups.

The following topics describe tasks you can perform in the **Notifications** view:

- ◆ [Viewing the Group Notification List on page 116](#)
- ◆ [Filtering the Notifications List on page 116](#)
- ◆ [Creating a Notification Group on page 117](#)
- ◆ [Sending Test Email to a Notification Group on page 117](#)
- ◆ [Modifying a Notification Group on page 117](#)
- ◆ [Resetting the Notification Group Configuration on page 118](#)
- ◆ [Deleting a Notification Group on page 118](#)

### Viewing the Group Notification List

The group list on the Notification tab displays the following information:

Item	Description
Group Name	The name of the group receiving the notification.
Class	The number of classes being tracked.
Subscribers	The number of email subscribers in the group.

Click a group in the group list to display the following information in the Detailed Information area:

Item	Description
Class Attributes	The name of a class and the severity ranking that will trigger an alert.
Subscribers	The email addresses of subscribers in the group.

### Filtering the Notifications List

To filter (or search for an item) in the notifications group list, type a group name in the Group Name box or type a subscriber email in the Alert Email box, and click **Update**. The result is brought to the top of the notification list.

#### Note

Click **Reset** to return the group list to the default order.

## Creating a Notification Group

By default, all alerts are sent to the Alerts Summary email group, but groups that receive specific classes of alert notification are configurable. This topic describes how to create a notification group.

### Procedure

1. Expand **DD Network** in the navigation panel and select a managed system.
2. Select **Status > Alerts > Notification**.
3. Click **Add**.  
The Add Group dialog box appears.
4. Type the group name in the **Group Name** box.
5. Select the checkbox of one or more classes of which to be notified.
6. To change the default severity level (Warning) for a class, select another level in the associated list box.
7. Click **OK**.
8. Select the checkbox of a group in the Notifications group list, and click **Modify**.  
The Modify Group dialog box appears.
9. Click **Subscribers**.
10. To add subscribers to the group, click the **+** icon.  
The Email Address dialog box appears.
11. Enter the email address of a subscriber and click **OK**.
12. Repeat steps 10 and 11 for each subscriber that needs to be added to the group.
13. Click **Finish**.

## Sending Test Email to a Notification Group

This topic describes how to send a test email to the subscribers in a notification group.

### Procedure

1. Expand **DD Network** in the navigation panel and select a managed system.
2. Select **Status > Alerts > Notification**.
3. Select **More Tasks > Send Test Alert**.  
The Send Test Alert dialog box appears.
4. In the Notification Groups list, select the checkboxes of the groups to receive the test email and click **Next**.
5. If necessary, add or modify email addresses in the Additional Email Addresses list.
6. Click **Send Now**.

## Modifying a Notification Group

This topic describes how to modify the attribute classes in an existing group.

### Procedure

1. Expand **DD Network** in the navigation panel and select a system.
2. Select **Status > Alerts > Notification**.

3. Select the checkbox of the group to modify in the group list.
4. To modify the class attributes for a group, do the following:
  - a. Click **Configure** in the Class Attributes area.  
The Edit Group dialog box appears.
  - b. Select (or clear) the checkbox of one or more class attributes.
  - c. To change the severity level for a class attribute, select a level from the corresponding list box.
  - d. Click **OK**.
5. To modify the subscriber list for a group, do the following:
  - a. Click **Configure** in the Subscribers area.  
The Edit Subscribers dialog box appears.
  - b. To delete subscribers from the group list, select the checkboxes of subscribers to delete and click the Delete icon (X).
  - c. To add a subscriber, click the Add icon (+), type a subscriber email address, and click **OK**.
  - d. Click **OK**.
6. Click **OK**.

## Resetting the Notification Group Configuration

This topic describes how to remove all notification groups added to and any changes made to the Default group.

### Procedure

1. Expand **DD Network** in the navigation panel and select a managed system.
2. Select **Status > Alerts > Notification**.
3. Select **More Tasks > Reset Notification Groups**.
4. In the Reset Notification Groups dialog box, click **Yes** in the verification dialog box and **OK**.

## Deleting a Notification Group

This topic describes how to delete one or more existing notification groups.

### Procedure

1. Select one or more checkboxes of groups in the Notifications group list, and click **Delete**.  
The Delete Group dialog box appears.
2. Verify the deletion and click **OK**.

## Configuring the Daily Alert Summary Distribution List

Every morning, each managed system sends a Daily Alert Summary email to the subscribers configured for the alertsummary.list email group. The Daily Alert Summary email contains current and historical alerts showing messages about non-critical hardware situations and disk space usage numbers that should be addressed soon. An example would be a fan failure. A failed fan should be replaced as soon as is reasonably possible, but the system can continue operations. When Support receives the failure notification, they contact you to arrange a replacement component.

**Procedure**

1. Expand **DD Network** in the navigation panel and select a managed system.
2. Select **Status > Alerts > Daily Alert Summary**.
3. If the default deliver time of 8 AM is not acceptable, click **Schedule**.  
The Schedule Alert Summary dialog box appears.
4. Use the list boxes to select the hour, minute, and AM or PM for the summary report, and click **OK**.
5. To configure the daily alert subscriber list, click **Configure**.  
The Daily Alert Summary Mailing List dialog box appears.
6. Modify the daily alert subscriber list as follows:
  - To add a subscriber, click the **+** icon, type the email address, and click **OK**.
  - To modify an email address, select the checkbox for the subscriber, click the pencil icon, edit the email address, and click **OK**.
  - To delete an email address, select the checkbox for the subscriber and click **X**.
7. Click **Finish**.

## Viewing Active Users

**Procedure**

1. Expand **DD Network** in the navigation panel and select a system.
2. Select **Status > Active Users**.

The Active Users list appears and displays the following for each user:

Item	Description
Name	User name of the logged-in user.
Idle	Time since last activity of user.
Last Login From	System from which the user logged in.
Last Login Time	Datestamp of when user logged in.
TTY	Terminal notation for CLI login.
Session	Identifier of the user session.

**Note**

To manage local users, click **Go to Local Users** (see [Managing Local User Access to the System on page 68](#)).

## Viewing System Statistics

The **Status > Stats** view displays up to five graphs that show real-time subsystem performance statistics, such as CPU usage and disk traffic.

**Procedure**

1. Expand **DD Network** in the navigation panel and select a system.

2. Select **Status > Stats**.

The Performance Graphs area displays the currently selected graphs.

3. To change the selection of graphs to display, select and clear the checkboxes for graphs in the list box.
4. To view specific data-point information, hover over a graph point.
5. When a graph contains multiple data, you can use the checkboxes in the upper-right corner of the graph to select what to display. For example, if Read is not selected in the upper right of the disk activity graph, only write data is graphed.

### Results

Each graph shows usage over the last 200 seconds. Click **Pause** to temporarily stop the display. Click **Resume** to restart it and show points missed during the pause.

## About the Performance Statistics Graphs

The **Status > Stats** view can display the following graphs:

- ◆ [Disk on page 120](#)
- ◆ [File System Operations on page 120](#)
- ◆ [Network on page 120](#)
- ◆ [CPU: Recent CPU Usage on page 120](#)
- ◆ [Replication \(DD Replicator must be licensed\) on page 120](#)

### Disk

The amount of data in the appropriate unit of measurement based on the data received, such as KiB or MiB per second, going to and from all disks in the managed system.

### File System Operations

- ◆ NFS ops/s—The number of NFS operations per second.
- ◆ CIFS ops/s—The number of CIFS operations per second.

### Network

The amount of data in the appropriate unit of measurement based on the data received, such as KiB or MiB per second, that passes through each Ethernet connection. One line appears for each Ethernet port.

### CPU: Recent CPU Usage

The percentage of CPU usage at a given point in time.

### Replication (DD Replicator must be licensed)

The unit of measurement is the one most appropriate for the procedure.

- ◆ In: The total number of units of measurement, such as kilobytes per second, received by this side from the other side of the Replicator pair. For the destination, the value includes backup data, replication overhead, and network overhead. For the source, the value includes replication overhead and network overhead.
- ◆ Out: The total number of units of measurement, such as kilobytes per second, sent by this side to the other side of the Replicator pair. For the source, the value includes backup data, replication overhead, and network overhead. For the destination, the value includes replication and network overhead.

## Working with Reports

DD System Manager lets you generate reports to track space usage on a Data Domain system for up to the previous two years. You can also generate reports to help



understand replication progress. You can view reports on the file system daily and cumulatively, over a period of time.

The Reports view is divided into two sections. The upper section lets you create the various types of reports. The lower section lets you view and manage saved reports.

Reports display in a table format, and as charts, depending on the type of report. You can select a report for a specific Data Domain system and provide a specific time period.

The reports display historical data, not real-time data. After the report is generated, the charts remain static and do not update. Examples of the types of information you can get from the reports include:

- ◆ The amount of data that was backed up to the system and the amount of de-duplication that was achieved
- ◆ Estimates of when the Data Domain system will be full, based on weekly space usage trends
- ◆ Backup and compression utilization based on selected intervals
- ◆ Historical cleaning performance, including duration of cleaning cycle, amount of space that can be cleaned, and amount of space that was reclaimed
- ◆ Amount of WAN bandwidth used by replication, for source and destination, and if bandwidth is sufficient to meet replication requirements
- ◆ System performance and resource utilization

## Types of Reports

The types of reports that are available are:

- ◆ [File System Cumulative Space Usage Report on page 121](#)
- ◆ [File System Daily Space Usage Report on page 122](#)
- ◆ [Replication Status Report on page 123](#)
- ◆ [Replication Summary Report on page 124](#)

---

### Note

Replication reports can only be created if the system has a replication license and a valid replication context configured.

---

## File System Cumulative Space Usage Report

File System Cumulative Space Usage Reports include cumulative pre-compression, post-compression, and total compression factor data on the system during the specified duration. This report is used to analyze how much data is backed up, the amount of deduplication performed, and how much space is consumed.

The File System Cumulative Space Usage report lists the following information:

Item	Description
File System—Usage	
Data Written (GiB)	The amount of data written before compression. This is indicated by a purple shaded area on the report.
Time	The timeline for data that has been written. The time displayed on this report changes based upon the Duration selection when the chart was created.

<b>Item</b>	<b>Description</b>
Total Compression Factor	The total compression factor reports the compression ratio.
File System—Consumption	
Used (GiB)	The amount of space used after compression.
Time	The date the data was written. The time displayed on this report changes based upon the Duration selection when the chart was created.
Post Comp	The amount of storage used after compression.
Usage Trend	The dotted black line shows the storage usage trend. When the line reaches the red line at the top, the storage is almost full.
Size and Cleaning	Size is the Total Capacity on a Data Domain system. Cleaning is the Cleaning cycle (start and end time for each cleaning cycle). Administrators can use this information to decide when space cleaning should run and what throttle to set.
File System Weekly Cumulative Capacity	
Date (or Time for 24 hour report)	The last day of each week, based on the criteria set for the report. In reports, a 24-hour period ranges from noon-to-noon.
Data Written (Pre-Comp)	The cumulative data written before compression for the specified time period.
Used (Post-Comp)	The cumulative data written after compression for the specified time period.
Total Compression Factor	The total compression factor. This is indicated by a black line on the report.

## File System Daily Space Usage Report

File System Daily Space Usage Reports include daily pre-compression written, post-compression used, and total compression factor on the system during the specified duration. This report is used to analyze daily activities.

The File System Daily Space Usage report lists the following information.

<b>Item</b>	<b>Description</b>
File System Daily Space Usage	
Space Used (GiB)	The amount of space used. Post-comp is red shaded area. Pre-Comp is purple shaded area.
Time	The date the data was written.
Compression Factor	The total compression factor. This is indicated by a black square on the report.
File System Daily Capacity Utilization	
Date	The date the data was written.
Data Written (Pre-Comp)	The amount of data written pre-compression.
Used (Post-Comp)	The amount of storage used after compression.
Total Compression Factor	The total compression factor.

<b>Item</b>	<b>Description</b>
File System Weekly Cumulative Capacity	
Start Date	The first day of the week for this summary.
End Date	The last day of the week for this summary.
Available	Total amount of storage available.
Consumed	Total amount of storage used.
Data (Post -Comp)	The cumulative data written before compression for the specified time period.
Replication (Post-Comp)	The cumulative data written after compression for the specified time period.
Overhead	Extra space used for non-data storage.
Reclaimed by Cleaning	The total space reclaimed after cleaning.

## Replication Status Report

Replication Status reports include the status of the current replication job running on the system. This report is used to provide a snapshot of what is happening for all replication contexts to help understand the overall replication status on a Data Domain System.

The Replication Status report lists the following information:

<b>Item</b>	<b>Description</b>
Replication Context Summary	
ID	The Replication Context identification.
Source	Source system name.
Destination	Destination system name.
Type	Type of replication context: MTree, Directory, Collection, or Pool.
Status	Replication status types include: Error, Normal.
Sync as of Time	Time and date stamp of last sync.
Estimated Completion	The estimated time the replication should be complete.
Pre-Comp Remaining	The amount of pre-compressed data to be replicated. This only applies to Collection type.
Post-Comp Remaining	The amount of post-compressed data to be replicated. This only applies to Directory and Pool types.
Replication Context Error Status	
ID	The Replication Context identification.
Source	Source system name.
Destination	Destination system name.
Type	Replication context type: Directory or Pool.
Status	Replication status types include: Error, Normal, and Warning.
Description	Description of the error.

Item	Description
Replication Destination Space Availability	
Destination	Destination system name.
Space Availability (GiB)	Total amount of storage available.

## Replication Summary Report

Replication Summary reports provide performance information about a system's overall network in-and-out usage for replication, as well as per context levels over a specified duration. You select the contexts to be analyzed from a list.

The Replication Summary report lists the following information for the system and for the selected context:

**Table 20** Replication Summary Report Label Descriptions

Item	Description
Replication Summary (shown for system and context)	
Network In (MiB)	The amount of data entering the system. Network In is indicated by a thin green line.
Network Out (MiB)	The amount of data sent from the system. Network Out is indicated by a thick orange line.
Time	The date on which the data was written.
Pre-Comp Remaining (MiB)	The amount of pre-compressed data to be replicated. Pre-Comp Remaining is indicated by a blue line.

## Creating a Report

This topic describes how to create a report.

### Procedure

1. In the navigation panel, click the **Reports** button.  
The information panel displays a new report area and a list of saved reports.
2. Click a report type in the New Report area (see [Types of Reports on page 121](#) for descriptions of available reports).
3. In the System menu, select the system for which you want to create a report .
4. Select additional options for the report based on the type of report:
  - Duration— Last 4 Weeks, Last 7 Days, or Custom

### Note

In reports, the duration of a 24-hour day ranges from noon-to-noon.

- Contexts—Available contexts for working with the Replication Summary report
5. If you select Custom for the duration, enter Start and End Date and Time in the additional fields.

6. Click **Create**.

The report appears in a separate browser window and is added to the end of the Saved Reports list.

---

**Note**

If the report does not display, verify the option to block pop-up windows is enabled on your browser.

---

## Viewing Saved Reports

This topic describes how to view saved reports.

**Procedure**

1. In the navigation panel, click the **Reports** button.

The information panel displays a new report area and a list of saved reports.

2. Select the report you want to view in the Saved Reports area.

3. Click **View**.

The report appears in a new browser window.

---

**Note**

If the report does not appear, verify the option to block pop-up windows is enabled on your browser.

---

## Printing Saved Reports

This topic describes how to print saved reports.

**Procedure**

1. In the navigation panel, click **Reports**.

The information panel displays a new report area and a list of saved reports.

2. Select the report you want to view in the Saved Reports area.

3. Click **View**.

The report appears in a new browser window.

4. In the browser window, select **File > Print**.

## Deleting Saved Reports

**Procedure**

1. In the navigation panel, click the **Reports**.

The information panel displays a new report area and a list of saved reports.

2. Select the report you want to delete in the Saved Reports area. You can select multiple reports to delete. Click the box at the top to select all reports.

3. Click **Delete**.

A warning dialog box asks if you are sure you want to delete the selected reports.

4. Click **OK**, and click **Close**.

## Renaming Saved Reports

### Procedure

1. In the navigation panel, click **Reports**.  
The information panel displays a new report area and a list of saved reports.
2. Select the report you want to rename in the **Saved Reports** area.
3. Click **Rename**.  
The Rename Report dialog box appears.
4. Type a name for your report in the New Report Name box.  
It is a good idea to give the report a simple, descriptive name you can easily recognize.
5. Click **OK**.

## Viewing the Task Log

The task log displays a list of currently running jobs, such as, replication or system upgrades. DD System Manager can manage multiple systems and can initiate tasks on those systems. If a task is initiated on a remote system, the progress of that task is tracked in the management station task log, not in the remote system task log. This topic describes how to view that task log.

### Procedure

1. In the navigation panel, click **Task Log**.  
The Tasks view appears.
2. Select a filter by which to display the Task Log from the Filter By list box. You can select **All**, **In Progress**, **Failed**, or **Completed**.  
The Tasks view displays the status of all tasks based on the filter you select and refreshes every 60 seconds.
3. To manually refresh the Tasks list:
  - Click **Update** to update the task log.
  - Click **Reset** to display all tasks and remove any filters you have set.
4. To display detailed information about a task, select the task in the task list.

The following status information appears in the Detailed Information area:

Item	Description
System	The descriptive name of the managed system.
Task Description	A description of the type of task.
Start Time	The date and time the task started.
Status	The status of the task (completed, failed, or in progress).
End Time	The date and time the task ended.
Error Message	An applicable error message, if any.

5. To return to the managing system, select the system in the navigation panel.

# CHAPTER 5

## Working with the File System

This chapter includes:

- ◆ [About the File System](#)..... 128
- ◆ [Monitoring File System Usage](#)..... 134
- ◆ [Managing File System Operations](#)..... 140
- ◆ [Fast Copy Operations](#)..... 146

## About the File System

This section describes how to use the file system with the following topics:

- ◆ [How the File System Stores Data on page 128](#)
- ◆ [How the File System Reports Space Usage on page 128](#)
- ◆ [How the File System Uses Compression on page 129](#)
- ◆ [How the File System Implements Data Integrity on page 130](#)
- ◆ [How the File System Reclaims Storage Space with File System Cleaning on page 130](#)
- ◆ [Supported Interfaces on page 131](#)
- ◆ [Supported Backup Software on page 131](#)
- ◆ [Data Streams Sent to a Data Domain System on page 131](#)
- ◆ [File System Limitations on page 133](#)

### How the File System Stores Data

A Data Domain system is designed as a very reliable online system for backups and archive data. As new backups are added to the system, old backups are aged out. Such removals are normally done under the control of backup or archive software based on the configured retention period.

When backup software expires or deletes an old backup from a Data Domain system, the space on the Data Domain system becomes available only after the Data Domain system cleans the data of the expired backups from disk. A good way to manage space on a Data Domain system is to retain as many online backups as possible with some empty space (about 20% of total space available) to comfortably accommodate backups until the next scheduled cleaning run, which runs once a week by default.

Some storage capacity is used by Data Domain systems for internal indexes and other metadata. The amount of storage used over time for metadata depends on the type of data stored and the sizes of the stored files. With two otherwise identical systems, one system may, over time, reserve more space for metadata and have less space for actual backup data than the other if different data sets are sent to each system.

Space utilization on a Data Domain system is primarily affected by:

- ◆ The size and compressibility of the backup data.
- ◆ The retention period specified in the backup software.

High levels of compression result when backing up datasets with many duplicates and retaining them for long periods of time.

### How the File System Reports Space Usage

All DD System Manager windows and system commands display storage capacity using base 2 calculations. For example, a command that displays 1 GiB of disk space as used is reporting  $2^{30}$  bytes = 1,073,741,824 bytes.

- ◆ 1 KiB =  $2^{10}$  = 1024 bytes
- ◆ 1 MiB =  $2^{20}$  = 1,048,576 bytes
- ◆ 1 GiB =  $2^{30}$  = 1,073,741,824 bytes
- ◆ 1 TiB =  $2^{40}$  = 1,099,511,627,776 bytes



## How the File System Uses Compression

The file system uses compression to optimize available disk space when storing data, so disk space is calculated two ways: physical and logical. (See [Types of Compression on page 129](#).) Physical space is the actual disk space used on the Data Domain system. Logical space is the amount of uncompressed data written to the system.

The file system space reporting tools (DD System Manager graphs and `filesys show space` command, or the alias `df`) show both physical and logical space. These tools also report the size and amounts of used and available space.

When a Data Domain system is mounted, the usual tools for displaying a file system's physical use of space can be used.

The Data Domain system generates warning messages as the file system approaches its maximum capacity. The following information about data compression gives guidelines for disk use over time.

The amount of disk space used over time by a Data Domain system depends on:

- ◆ The size of the initial full backup.
- ◆ The number of additional backups (incremental and full) retained over time.
- ◆ The rate of growth of the backup dataset.
- ◆ The change rate of data.

For data sets with typical rates of change and growth, data compression generally matches the following guidelines:

- ◆ For the first full backup to a Data Domain system, the compression factor is generally 3:1.
- ◆ Each incremental backup to the initial full backup has a compression factor generally in the range of 6:1.
- ◆ The next full backup has a compression factor of about 60:1.

Over time, with a schedule of weekly full and daily incremental backups, the aggregate compression factor for all the data is about 20:1. The compression factor is lower for incremental-only data or for backups with less duplicate data. Compression is higher when all backups are full backups.

## Types of Compression

A Data Domain system compresses data at two levels: global and local. Global compression compares received data to data already stored on disks. Duplicate data does not need to be stored again, while data that is new is locally compressed before being written to disk.

### Local Compression

A Data Domain system uses a local compression algorithm developed specifically to maximize throughput as data is written to disk. The default algorithm (lz) allows shorter backup windows for backup jobs but uses more space. Local compression options provide a trade-off between slower performance and space usage. To change compression, see “Change Local Compression” on page 209.

After you change the compression, all new writes use the new compression type. Existing data is converted to the new compression type during cleaning. It takes several rounds of cleaning to recompress all of the data that existed before the compression change.

The initial cleaning after the compression change might take longer than usual. Whenever you change the compression type, carefully monitor the system for a week or two to verify that it is working properly.

## How the File System Implements Data Integrity

Multiple layers of data verification are performed by the DD OS file system on data received from backup applications to ensure that data is written correctly to the Data Domain system disks. This ensures the data can be retrieved without error.

The DD OS is purpose-built for data protection and it is architecturally designed for data invulnerability. There are four critical areas of focus, described in the following sections.

### End-to-End Verification

End-to-end checks protect all file system data and metadata. As data comes into the system, a strong checksum is computed. The data is deduplicated and stored in the file system. After all data is flushed to disk, it is read back, and re-checksummed. The checksums are compared to verify that both the data and the file system metadata are stored correctly.

### Fault Avoidance and Containment

New data never puts old data at risk. Data Domain uses a log-structured file system that never overwrites or updates existing data. New data is always written in new containers and appended to existing old containers. The old containers and references remain in place and are safe even in the face of software bugs or hardware faults that may occur when storing new backups.

### Continuous Fault Detection and Healing

Continuous fault detection and healing protects against storage system faults. The system periodically rechecks the integrity of the RAID stripes, and uses the redundancy of the RAID system to heal any faults. During a read, data integrity is re-verified and any errors are healed on the fly.

### File System Recoverability

Data is written in a self-describing format. The file system can be re-created, if necessary, by scanning the log and rebuilding it from the metadata stored with the data.

## How the File System Reclaims Storage Space with File System Cleaning

When your backup application (such as NetBackup or NetWorker) expires data, the data is marked by the Data Domain system for deletion. However, the data is not deleted immediately; it is removed during a cleaning operation.

- ◆ During the cleaning operation, the file system is available for all normal operations including backup (write) and restore (read).
- ◆ Although cleaning uses a significant amount of system resources, cleaning is self-throttling and gives up system resources in the presence of user traffic.
- ◆ Data Domain recommends running a cleaning operation after the first full backup to a Data Domain system. The initial local compression on a full backup is generally a factor of 1.5 to 2.5. An immediate cleaning operation gives additional compression by another factor of 1.15 to 1.2 and reclaims a corresponding amount of disk space.
- ◆ When the cleaning operation finishes, a message is sent to the system log giving the percentage of storage space that was reclaimed.

A default schedule runs the cleaning operation every Tuesday at 6 a.m. (tue 0600). You can change the schedule or you can run the operation manually (see “Modify a Cleaning Schedule” on page 208).

Data Domain recommends running the cleaning operation once a week.

---

#### Note

Any operation that disables the file system, or shuts down a Data Domain system during a cleaning operation (such as a system power-off or reboot) aborts the cleaning operation. The cleaning operation does not immediately restart when the system restarts. You can manually restart the cleaning operation or wait until the next scheduled cleaning operation.

With MTree replication, if a file is created and deleted while a snapshot is being replicated, then the next snapshot will not have any information about this file, and the system will not replicate any content associated with this file. Directory replication will replicate both the create and delete, even though they happen close to each other.

With the replication log that directory replication uses, operations like deletions, renaming, and so on, execute as a single stream. This can reduce the replication throughput. The use of snapshots by MTree replication avoids this problem.

---

## Supported Interfaces

The following interfaces are supported by the file system:

- ◆ NFS
- ◆ CIFS
- ◆ DD Boost
- ◆ VTL

## Supported Backup Software

Data Domain offers guidance on setting up backup software and backup servers for use with a Data Domain system. Because such information tends to change often, it is available on the Data Domain Support web site (<https://my.datadomain.com/>).

For more information about the backup applications that are supported and instructions for accessing the Data Domain Support web site compatibility matrices, see the section [Backup Software Requirements on page 21](#).

## Data Streams Sent to a Data Domain System

A data stream, in the context of Table 5, refers to a large byte-stream associated with sequential file access, such as a write stream to a backup file or a read stream from a restore image. A Repl Source or Destination stream refers to a directory replication operation or a DD Boost file replication stream associated with a file replication operation.

For optimal performance, Data Domain recommends the limits on simultaneous streams between Data Domain systems and your backup servers, as described in Table 5.

**Table 21** Data Streams Sent to a Data Domain System in DD OS 5.4

Model	RAM/ NVRAM	Backup Write Streams	Backup Read Streams	Repl <sup>1</sup> Source Streams	Repl <sup>1</sup> Dest Streams	Mixed
DD120, DD140, DD160, DD510, DD530, DD610	4 GB or 6 GB / 0.5 GB	16	4	15	20	w<= 16 ; r<= 4 ReplSrc<=15; ReplDest<=20; ReplDest+w<=16; w+r+ReplSrc <=16; Total<=20
DD565, DD620, DD630, DD640	8 GB / 0.5 GB or 1 GB	20	16	30	20	w<=20; r<=16; ReplSrc<=30; ReplDest<=20; ReplDest+w<=20; Total<=30
DD565	12 GB / 0.5 GB	45	20	45	45	w<=20; r<=16; ReplSrc<=45; ReplDest<=45; ReplDest+w<=45; Total<=45
DD580, DD580g	16 GB / 0.5 GB	45	30	60	45	w<=45; r<=30; ReplSrc<=60; ReplDest<=45; ReplDest+w<=45; Total<=60
DD640, DD660, DD670, DD690, DD690g	16 GB or 20 GB / 1 GB	90	30	60	90	w<=90; r<=30; ReplSrc<=60; ReplDest<=90; ReplDest+w<=90; Total<=90
DD690, DD690g	24 GB / 1GB	90	50	90	90	w<=90; r<=50; ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; Total<=90
DD670, DD860	36 GB / 1 GB	90	50	90	90	w<=90; r<=50; ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; Total<=140
DD880, DD880g	64 GB / 2 GB	180	50	90	180	w<=180; r<=50; ReplSrc<=90; ReplDest<=180; ReplDest+w<=180; Total<=180
DD860	72 GB <sup>2</sup> / 1 GB	90	50	90	90	w<=90; r<=50; ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; Total<=140
DD890	96 GB / 2 GB	180	50	90	180	w<=180; r<=50; ReplSrc<=90; ReplDest<=180; ReplDest+w<=180; Total<=180
DD990	128 or 256 GB <sup>2</sup> / 4 GB	540	150	270	540	w<=540; r<=150; ReplSrc<=270; ReplDest<=540; ReplDest+w<=540; Total<=540
DD2500	32 or 64 GB / 2 GB	180	50	90	180	w<=180; r<=50; ReplSrc<=90; ReplDest<=180; ReplDest+w<=180; Total<=180
DD4200	128 GB <sup>2</sup> / 4 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest+w<=270; Total<=270
DD4500	192 GB <sup>2</sup> / 4 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest+w<=270; Total<=270

**Table 21** Data Streams Sent to a Data Domain System in DD OS 5.4 (continued)

Model	RAM/ NVRAM	Backup Write Streams	Backup Read Streams	Repl <sup>1</sup> Source Streams	Repl <sup>1</sup> Dest Streams	Mixed
DD7200	128 or 256 GB <sup>2</sup> / 4GB	540	150	270	540	w<=540; r<=150; ReplSrc<=270; ReplDest<=540; ReplDest+w<=540; Total<=540

1: DirRepl, OptDup, MTreeRepl streams. 2: The Data Domain Extended Retention software option is available only for these devices with extended (maximum) memory.

## File System Limitations

There are some file system limitations to be aware of while working with the Data Domain system, as described in the following sections.

- ◆ [Limits on Number of Files in a Data Domain System on page 133](#)
- ◆ [Maximum Number of Supported Inodes on page 134](#)
- ◆ [Maximum Path Name Length on page 134](#)

### Limits on Number of Files in a Data Domain System

Data Domain recommends storing no more than 100 million files on a system. A larger number of files can adversely affect performance and the length of cleaning. Some processes, such as file system cleaning, may run much longer with a very large number of files. For example, the enumeration phase of cleaning takes about five minutes for one million files and over eight hours for 100 million files.

A system does not have a fixed limit on the number of files. Available disk space is used as needed to store data and the metadata that describes files and directories. In round numbers, each file or directory uses about 1000 bytes of metadata. A Data Domain system with five TB of space available could hold up to five billion empty files. The amount of space used by data in files directly reduces the amount of space available for metadata, and vice versa.

---

#### Note

The overall performance for the Data Domain system will fall to unacceptable levels if the system is required to support the maximum file amount, and the workload from the client machines is not carefully controlled.

As well, consider the overhead of about 500 post-comp bytes per empty file and about 1K post-comp bytes per non-empty file. Therefore, a five TB system could hold about 10 billion zero length files (if they were spread across multiple Mtrees) is less than five billion non-empty files, depending on the data compression factor. After that, the disk space occupied by user data will dominate the equation and the total number of files the Data Domain system can store will gradually decrease based upon the overall compression factor of the user data.

Many systems operate without problems with hundreds of millions of files. After a billion files, some processes or operations might be affected, for example:

- ◆ Lengthy cleaning operations which could be scheduled less frequently in a stable archive environment.
- ◆ AutoSupport operations (schedule to occur less frequently).

- ◆ Any process or command that needs to enumerate all the files.

If there are many small files, other considerations arise:

- ◆ Initial bulk migration of files may take a while. For example, at 70 MB/sec, the task will take at least four hours.
- ◆ The number of separate files that can be created per second, (even if the files are very small) may be more of a limitation than the number of MB/s that can be moved into a Data Domain system. When files are large, the file creation rate is not significant, but when files are small, the file creation rate dominates and may become a factor. The file creation rate is about 30 files per second per MTree per CIFS connection. This rate should be taken into account during system sizing when a bulk ingest of a large number of files is needed by a customer environment.
- ◆ File access latencies are affected by the number of files in a directory. To the extent possible, we recommend directory sizes of less than a thousand files. Larger directory sizes will experience slower responses to metadata operations such as listing the files in the directory and opening or creating a file.

## Maximum Number of Supported Inodes

An NFS or CIFS client request causes a Data Domain system to report a capacity of about two billion inodes (files and directories). A Data Domain system can exceed that number, but the reporting on the client may be incorrect.

## Maximum Path Name Length

The maximum length of a full path name (including the characters in `/data/coll/backup`) is 1023 bytes. The maximum length of a symbolic link is also 1023 bytes.

# Monitoring File System Usage

The File System view has tabs that show real-time data storage statistics, including current compression factors showing the space saved by using data deduplication, graphs of space usage amounts, consumption factors, and data written trends. There are also some options for managing file system cleaning, expansion, copying, and destruction.

## Accessing the File System View

### Procedure

1. Select a system in the Navigation panel.
2. Click the **Data Management** > **File System** tabs.

The File System view has a **File System** overview panel and six tabs which are described in detail in the following sections:

- [About the File System Overview Panel on page 135](#)
- [About the Summary View on page 135](#)
- [About the Configuration View on page 136](#)
- [About the Encryption View on page 137](#)
- [About the Space Usage View on page 138](#)
- [About the Consumption View on page 139](#)
- [About the Daily Written View on page 139](#)

## About the File System Overview Panel

The File System overview panel displays the file system State and the Clean Status.

### State

The State area contains an **Enable/Disable** button and shows the working state of the file system:

- ◆ Enabled and running—and the latest consecutive length of time the file system has been enabled and running.
- ◆ Disabled and shutdown.
- ◆ Enabling and disabling—in the process of becoming enabled or disabled.
- ◆ Destroying—if the file system is being deleted.
- ◆ Error—if there is an error condition, such as a problem initializing the file system.

### Clean Status

The **Clean Status** area contains a **Start/Stop Cleaning** button and shows the date the last cleaning operation occurred, or the current cleaning status if the cleaning operation is currently running. For example:

```
Cleaning finished at 2009/01/13 06:00:43
```

or, if the file system is disabled, shows:

```
Unavailable
```

## About the Summary View

Click the Summary tab to view important file system statistics as described in the following section.

### Space Usage

The first Space Usage pane shows the amount of disk space available and used by the file system components, based on the last cleaning.

- ◆ The `/data:post-comp` line shows amounts for compressed data in the `/data` directory.
- ◆ The `/ddvar` line shows amounts for log and core files. (Remove old logs and core files to free space in this area.)

For both of these, the following amounts are shown in real numbers and in the color-coded graph as described in [About the Space Usage View on page 138](#):

- ◆ Size—The amount of total physical disk space available for data.
- ◆ Used—the actual physical space used for compressed data. Warning messages go to the system log and an email alert is generated when the use reaches 90%, 95%, and 100%. At 100%, the Data Domain system accepts no more data from backup servers. If the Used amount is always high, check the cleaning schedule to see how often the cleaning operation runs automatically, then use the procedure [Modify a Cleaning Schedule on page 143](#) to run the operation more often. Also consider reducing the data retention period or splitting off a portion of the backup data to another Data Domain system.
- ◆ Available (GiB)—The total amount of space available for data storage. This figure can change because an internal index may expand as the Data Domain system fills with data. The index expansion takes space from the Avail GiB amount.

- ◆ Cleanable (GiB)—The amount of space that could be reclaimed if a cleaning were run.

The second Space Usage pane shows the compression factors:

- ◆ Currently Used—The amounts currently in use by the file system.
- ◆ Written in Last 24 Hours—The compression activity over the last day.

For both of these areas, the following is shown:

- ◆ Pre-Compression (GiB)—Data written before compression.
- ◆ Post-Compression (GiB)—Storage used after compression.
- ◆ Global-Comp Factor—Pre-Compression / (Size after global compression).
- ◆ Local-Comp Factor—(Size after global compression) / Post-Compression
- ◆ Total-Comp Factor—Pre-Comp / Post-Comp
- ◆ Reduction %—[(Pre-Comp - Post-Comp) / Pre-Comp] \* 100

## About the Archive Units View

The Archive Units view on the File System page is shown only when the optional DD Extended Retention license is activated. This view lists each archive unit and shows the unit's state (new, sealed, or target), its status (disabled or ready), and its size. If the unit has been sealed, meaning no more data can be added, the date that it was sealed is given.

Click the diamond symbol to the right of a column heading to sort the order of the values in reverse.

## About the Configuration View

To check the file system configuration settings, click the Configuration tab. The Configuration view shows the configurable options and the current cleaning schedule, along with Edit buttons to change those settings.

The Options settings and descriptions follow:

Options Settings	Description
Local Compression Type	The type of local compression in use. See: <ul style="list-style-type: none"> <li>• <a href="#">Types of Compression on page 129</a> for an overview.</li> <li>• <a href="#">Change Local Compression on page 144</a></li> </ul>
Report Relica as Writable	How applications see a replica. See: <ul style="list-style-type: none"> <li>• <a href="#">Change Read-only Settings on page 144</a></li> </ul>
Marker Type	Backup software markers (tape markers, tag headers, or other names are used) in data streams. See: <a href="#">Tape Marker Settings on page 146</a>
Staging Reserve	Manage disk staging. See: <ul style="list-style-type: none"> <li>• <a href="#">Working with Disk Staging on page 145</a></li> <li>• <a href="#">Configure Disk Staging on page 145</a></li> </ul>

Cleaning Schedule Settings	Description
----------------------------	-------------

Time	The date time cleaning operations run. See:
------	---



Cleaning Schedule Settings	Description
	<ul style="list-style-type: none"> <li>• <a href="#">Modify a Cleaning Schedule on page 143</a></li> </ul>
Throttle	<p>The system resources allocation. See:</p> <ul style="list-style-type: none"> <li>• <a href="#">Throttle the Cleaning Operation on page 143</a></li> </ul>

## About the Encryption View

The Encryption settings and descriptions follow:

Setting	Description
Encryption	
Status	<p>Status can be one of the following:</p> <ul style="list-style-type: none"> <li>• Not licensed—No other information provided.</li> <li>• Not configured—Encryption is licensed but not configured.</li> <li>• Enabled—Encryption is enabled and running.</li> <li>• Disabled—Encryption is disabled.</li> </ul>
Encryption Algorithm	<p>The algorithm used to encrypt the data:</p> <ul style="list-style-type: none"> <li>• AES 256-bit (CBC) (default)</li> <li>• AES 256-bit (GCM) (more secure but slower)</li> <li>• AES 128-bit (CBC) (not as secure as 256-bit)</li> <li>• AES 128-bit (GCM) (not as secure as 256-bit)</li> </ul> <p>See <a href="#">Changing the Encryption Algorithm on page 159</a> for details.</p>
Encryption Passphrase	<p>When configured, shows as “*****.”To change the passphrase, see <a href="#">Managing the Encryption Passphrase on page 160</a></p>
File System Lock	
Status	<p>The File System Lock status is either:</p> <ul style="list-style-type: none"> <li>• Unlocked—The feature is not enabled.</li> <li>• Locked—The feature is enabled.</li> </ul>
Key Management	
Key Manager	<p>Either the internal Data Domain Embedded Key Manager or the optional RSA Data Protection Manager (DPM) Key Manager. Click <b>Configure</b> to switch between key managers (if both are configured), or to modify Key Manager options.</p>
Server	<p>The name of the RSA Key Manager Server.</p>
Server Status	<p>Online or offline, or the error messages returned by the RSA Key Manager Server.</p>
Key Class	<p>A specialized type of security class used by the optional RSA Data Protection Manager (DPM) Key Manager that groups cryptographic keys with similar characteristics. The Data Domain system retrieves a</p>

Setting	Description
	key from the RSA server by key class. A key class to be set up to either return the current key, or to generate a new key each time.
	<p><b>Note</b></p> <p>The Data Domain system supports only key classes configured to return the current key.</p>
Port	The port number of the RSA server.
FIPS mode	Whether or not the imported host certificate is FIPS compliant. The default mode is enabled.
Encryption Keys	<p>Lists keys by ID numbers. Shows when a key was created, how long it is valid, its type (RSA DPM Key Manager or the Data Domain internal key), state (see <a href="#">Table 22 on page 151</a>), and its post-compression size. Selected keys in the list can be:</p> <ul style="list-style-type: none"> <li>• Synchronized so the list shows new keys added to the RSA server (but are not usable until the file system is restarted).</li> <li>• Deleted.</li> <li>• Destroyed.</li> </ul>

## About the Space Usage View

The Space Usage view contains a graph that displays a visual representation of data usage for the MTree. Click the **Data Management > MTree > Space Usage** tabs.

- ◆ Click a point on a graph line to display a box with data at that point.
- ◆ Click **Print** (at the bottom on the graph) to open the standard Print dialog box.
- ◆ Click **Show in new window** to display the graph in a new browser window.

The lines of the graph denote measurement for:

- ◆ **Pre-comp Written**—The total amount of data sent to the MTree by backup servers. Pre-compressed data on an MTree is what a backup server sees as the total uncompressed data held by an MTree-as-storage-unit, shown with the Space Used (left) vertical axis of the graph.
- ◆ **Post-comp Used**—The total amount of disk storage in use on the MTree, shown with the Space Used (left) vertical axis of the graph.
- ◆ **Comp Factor**—The amount of compression the Data Domain system has performed with the data it received (compression ratio), shown with the Compression Factor (right) vertical axis of the graph.

### Checking Historical Space Usage

On the Space Usage graph, clicking an interval (that is, 7d, 30d, 60d, or 120d) on the Duration line above the graph allows you to change the number of days of data shown on the graph, from 7 to 120 days.

To see space usage for intervals over 120 days, issue the following command:

```
# filesys show compression [summary | daily | daily-detailed] {[last n
{hours | days | weeks | months}] | [start date [end date]]}
```

## About the Consumption View

The Consumption view shows the space used over time, shown in relation to total system capacity.

- ◆ Click a point on a graph line to display a box with data at that point.
- ◆ Click **Print** (at the bottom on the graph) to open the standard Print dialog box.
- ◆ Click **Show in new window** to display the graph in a new browser window.

The lines of the graph denote measurement for:

- ◆ **Capacity**—The total amount of disk storage available for data on the Data Domain system. The amount is shown with the Space Used (left) vertical axis of the graph. Clicking the Capacity checkbox toggles this line on and off.
- ◆ **Post-comp**—The total amount of disk storage in use on the Data Domain system. Shown with the Space Used (left) vertical axis of the graph.
- ◆ **Comp Factor**—The amount of compression the Data Domain system has performed with the data it received (compression ratio). Shown with the Compression Factor (right) vertical axis of the graph.
- ◆ **Cleaning**—A grey diamond is displayed on the chart each time a file system cleaning operation was started.
- ◆ **Data Movement**—The amount of disk space moved to the archiving storage area (if the Archive license is enabled).

### Checking Historical Consumption Usage

On the Consumption graph, clicking an interval (that is, 7d, 30d, 60d, 120d) on the Duration line above the graph allows you to change the number of days of data shown on the graph, from 7 to 120 days.

## About the Daily Written View

The Data Written pane contains a graph that displays a visual representation of data that is written daily to the MTree over a period of time, selectable from 7 to 120 days. The data amounts are shown over time for pre- and post-compression amounts.

It also provides totals for global and local compression amounts, and pre-compression and post-compression amounts.

- ◆ Click a point on a graph line to display a box with data at that point.
- ◆ Click **Print** (at the bottom on the graph) to open the standard Print dialog box.
- ◆ Click **Show** in new window to display the graph in a new browser window.

The lines on the graph denote measurements for:

- ◆ **Pre-Comp**—The total amount of data written to the MTree by backup servers. Pre-compressed data on an MTree is what a backup server sees as the total uncompressed data held by an MTree -as-storage-unit.
- ◆ **Post-Comp**—The total amount of data written to the MTree after compression has been performed, as shown in GiBs.
- ◆ **Total Comp**—The total amount of compression the Data Domain system has performed with the data it received (compression ratio). Shown with the Total Compression Factor (right) vertical axis of the graph.

### Checking Historical Written Data

On the Daily Written graph, clicking an interval (for example, 7d, 30d, 60d, or 120d) on the Duration line above the graph allows you to change the number of days of data shown on the graph, from 7 to 120 days.

Below the Daily Written graph, the following totals display for the current duration value:

- ◆ Pre-comp
- ◆ Post-comp
- ◆ Global-comp factor
- ◆ Local-comp factor
- ◆ Total-comp factor

## When the File System Is Full or Nearly Full

A Data Domain system has three progressive levels of being full. As each level is reached, progressively more operations are disallowed. At each level, deleting data and performing a file system cleaning operation makes disk space available for continued operation. Deleting files and removing snapshots do not immediately reclaim disk space, but allow the next cleaning operation to reclaim the space.

- ◆ Level 1—At the first level of fullness, no more new data can be written to the file system. An informative out of space alert is generated.  
Remedy—Delete unneeded datasets, reduce the retention period, delete snapshots, and perform a file system cleaning operation.
- ◆ Level 2—At the second level of fullness, files cannot be deleted. This is because deleting files also require free space but the system has so little free space available that it cannot even delete files.  
Remedy—Expire snapshots and perform a file system cleaning operation.
- ◆ Level 3—At the third and final level of fullness, attempts to expire snapshots, delete files, or write new data fail.  
Remedy—Perform a file system cleaning operation to free enough space to at least delete some files or expire some snapshots and then rerun cleaning.

## Monitor the Space Usage with Email Alerts

Alerts are generated when the file system is at 90%, 95%, and 100% full. To receive these alerts, add the user to the alert emailing list. To join the alert email list, see [Working with the Notification View on page 116](#).

# Managing File System Operations

The following file system operations are described in this section:

- ◆ [Performing Basic Operations on page 140](#)
- ◆ [Performing Cleaning on page 143](#)
- ◆ [Modifying Basic Settings on page 144](#)

## Performing Basic Operations

Basic file system operations include enabling and disabling the file system, and in the rare occasion, destroying a file system.

## Creating the File System

There are three reasons to create a file system:

- ◆ For a new Data Domain system.
- ◆ When a system is started after a clean installation.

- ◆ After a file system has been destroyed.

To create the file system:

#### Procedure

1. Verify that storage has been installed and configured. If the system does not meet this prerequisite, a warning message is displayed. Install and configure the storage before attempting to create the file system.
2. Select a system in the Navigational pane.
3. Click the **Data Management** › **File System** tabs.
4. From the **More Tasks** menu, select **Create File System**.

The File System Create dialog box shows the approximate size of the file system. Check **Enable file system** after creation to start using this file system as soon as it is created. Click **Next**.

5. A summary displays the file system size and whether the file system is to be automatically enabled. Click **Back** if you want to change the enable the file system option. Clicking **Finish** starts the file system creation.
6. A progress bar measures the file system creation's progress. A check mark indicates that a step of the procedure has completed. When a check mark Completed is displayed, click **OK**.

## Enable or Disable the File System

The option to enable or disable the file system is dependent on the current state of the file system—if its enabled, you can disable it and vice versa.

- ◆ Enabling the file system allows Data Domain system operations to begin. This ability is available to administrative users only.
- ◆ Disabling the file system halts all Data Domain system operations, including cleaning. This ability is available to administrative users only.

#### CAUTION

Disabling the file system when a backup application is sending data to the system can cause the backup process to fail. Some backup software applications are able to recover by restarting where they left off when they are able to successfully resume copying files; others might fail, leaving the user with an incomplete backup.

#### Procedure

1. Select a system in the Navigational pane.
2. Click the **Data Management** › **File System** tabs.
3. In the Overview pane, click **Enable** or **Disable** in the State area.
4. Click **OK** and **Close**.

## Expanding the File System

You might need to expand the size of a file system if the suggestions given in [When the File System Is Full or Nearly Full on page 140](#) do not clear enough space for normal operations.

A file system may not be expandable, however, for these reasons:

- ◆ There are no unused disks or enclosures in the Active or Retention tiers.
- ◆ An expanded storage license is not installed.

- ◆ There are not enough capacity licenses installed.

To expand the file system:

#### Procedure

1. Select a system in the Navigation panel.
2. Click the **Data Management** > **File System** tabs.
3. From the **More Tasks** menu, select **Expand Capacity**.

The Expand File System Capacity window shows the current size of the file system and notes how much additional storage space is available for expansion.

- If enough capacity is available for expansion requirements, continue to step 7.
- If capacity needs to be added, continue with the next step.

4. Click **Configure** to allocate existing storage to the file system.

The Configure Storage dialog box is displayed.

5. In the Available Storage area, click the checkboxes of the storage devices to use and click **Add to Tier**.

System storage must be moved from the Available Storage area to the Active Tier storage area before the file system can be created.

6. Click **OK** and **Close** in the progress dialog box.
7. Click **Finish** to expand the file system into the available storage.

## Destroy the File System

Destroying the file system should be done only under the direction of Customer Support. This action deletes all data in the Data Domain file system, including virtual tapes. Deleted data is not recoverable. This operation also removes Replication configuration settings.

This operation is used when it is necessary to cleaning out existing data, to create a new collection replication destination, or to replace a collection source, or for security reasons because the system is being removed from operation.



**The optional Write zeros to disk operation writes zeros to all file system disks, effectively removing all traces of data. If the Data Domain system contains a large amount of data, this operation can take many hours, or a day, to complete.**

---

#### Note

As this is a destructive procedure, this operation is available to administrative users only.

---

#### Procedure

1. From the **More Tasks** menu, select **Destroy**.
2. In the Destroy File System dialog box, enter the sysadmin password. (It is the only accepted password.)
3. Optionally, click the checkbox for **Write zeros to disk** to completely remove data.
4. Click **OK**.

## Performing Cleaning

To start or stop cleaning, or to modify the default cleaning schedule (every Tuesday at 6 a.m. with 50% throttle), use one of the following procedures:

### Manually Start Cleaning

To immediately start a cleaning operation:

#### Procedure

1. In the Overview pane, click **Start Cleaning** in the Clean Status area.  
The **Start File System Clean** dialog box is displayed.
2. In the Throttle Percentage text box, enter a system throttle amount. This is the percentage of CPU usage dedicated to cleaning. The default is 50 percent.
3. Click **OK**.
4. The Start File System Clean dialog box allows you to monitor the cleaning operation progress. Click **Close** to exit the progress dialog box.

### Manually Stop Cleaning

To immediately stop a cleaning operation (stopping the process means that all work done so far is lost):

#### Procedure

1. In the Overview pane, click **Stop Cleaning** in the **Clean Status** area.  
The Stop File System Clean dialog box is displayed.
2. Click **OK**.

### Modify a Cleaning Schedule

Use the following procedure to change the cleaning schedule.

#### Procedure

1. Click the **Data Management** > **File System** > **Configuration** tabs.
2. In the Clean Schedule area, click **Edit**.  
The Modify Schedule dialog box is displayed.
3. Select the cleaning frequency, such as daily, bi-weekly, or monthly. Weekly is the default.
4. Enter the start day of the week, time, and throttle percentage.

---

#### Note

Cleaning leverages all resources when they are idle.

5. Click **OK**.

### Throttle the Cleaning Operation

If necessary, you can modify the throttle settings to change the amount of system resources used by the cleaning process using the [Modify a Cleaning Schedule](#) on page 143 process.

## Modifying Basic Settings

The Modify Settings option allows you to change the type of compression used, marker types, Replica write status, and Staging Reserve percentage, as described in the following sections:

- ◆ [Change Local Compression on page 144](#)
- ◆ [Change Read-only Settings on page 144](#)
- ◆ [Configure Disk Staging on page 145](#)

### Change Local Compression

---

#### Note

Do not change the type of local compression unless it is necessary.

---

To change the type of local compression in use:

#### Procedure

1. Click the **Data Management** > **File System** > **Configuration** tabs.
2. In the Options area, click **Edit**.  
The Modify Settings dialog box is displayed.
3. In the Local Compression Type area, click the drop-down list and select a new compression type.

Option	Description
none	Do not compress data.
lz	The default algorithm that gives the best throughput. Data Domain recommends the lz option.
gzfast	A zip-style compression that uses less space for compressed data, but more CPU cycles (twice as much as lz). Gzfast is the recommended alternative for sites that want more compression at the cost of lower performance.
gz	A zip-style compression that uses the least amount of space for data storage (10% to 20% less than lz on average; however, some datasets get much higher compression). This also uses the most CPU cycles (up to five times as much as lz). The gz compression type is commonly used for nearline storage applications in which performance requirements are low.

4. Click **OK** at the confirmation dialog box.
5. Click **Close** to exit the status dialog box.

### Change Read-only Settings

Some backup applications must see the replica as writable to do a restore or vault operation from the replica. To change the replica to writable:

#### Procedure

1. Click the **Data Management** > **File System** > **Configuration** tabs.
2. In the Options area, click **Edit**.

The Modify Settings dialog box is displayed.



3. In the Report Replica as Writable pane, click the **Enable** checkbox.
4. Click **OK**.
5. Click **Close** to exit the status dialog box.

## Working with Disk Staging

Disk staging enables a Data Domain system to serve as a staging device, where the system is viewed as a basic disk via a CIFS share or NFS mount point. Disk staging can be used in conjunction with your backup software, such as Symantec's NetBackup (NBU) and EMC's NetWorker.

The Data Domain disk staging feature does not require a license and is disabled by default.

---

### Note

The VTL feature is not required or supported when the Data Domain system is used as a Disk Staging device.

---

The reason that some backup applications use disk staging devices is to enable tape drives to stream continuously. After the data is copied to tape, it is retained on disk for as long as space is available. Should a restore be needed from a recent backup, more than likely the data is still on disk and can be restored from it more conveniently than from tape. When the disk fills up, old backups can be deleted to make space. This delete-on-demand policy maximizes the use of the disk.

In normal operation, the Data Domain System does not reclaim space from deleted files until a cleaning operation is done. This is not compatible with backup software that operates in a staging mode, which expects space to be reclaimed when files are deleted. When you configure disk staging, you reserve a percentage of the total space—typically 20 to 30 percent—in order to allow the system to simulate the immediate freeing of space.

The amount of available space is reduced by the amount of the staging reserve. When the amount of data stored uses all of the available space, the system is full. However, whenever a file is deleted, the system estimates the amount of space that will be recovered by cleaning and borrows from the staging reserve to increase the available space by that amount. When a cleaning operation runs, the space is actually recovered and the reserve restored to its initial size. Since the amount of space made available by deleting files is only an estimate, the actual space reclaimed by cleaning may not match the estimate. The goal of disk staging is to configure enough reserve so that you do not run out before cleaning is scheduled to run.

## Configure Disk Staging

To enable disk staging and specify the staging reserve percentage:

### Procedure

1. Click the **Data Management** > **File System** > **Configuration** tabs.
2. In the Options area, click **Edit**.  
The Modify Settings dialog box is displayed.
3. In the Staging Reserve pane, click the **Enable** checkbox.
4. Enter a value in the % of **Total Space** text box.

This value represents the percentage of the total disk space to be reserved for disk staging, typically 20 to 30 percent.

5. Click **OK**.

## Tape Marker Settings

Backup software from some vendors insert markers (tape markers, tag headers, or other names are used) in all data streams (both file system and VTL backups) sent to a Data Domain system. Markers can significantly degrade data compression on a Data Domain system. As such, the default marker type auto is set and cannot be changed by the user. If this setting is not compatible with your backup software, contact your contracted support provider.

---

### Note

For information about how applications work in a Data Domain environment, see [Application Compatibility Matrices and Integration Guides on page 21](#). You can use these matrices and integration guides to troubleshoot vendor-related issues.

---

## Fast Copy Operations

A fast copy operation clones files and directory trees of a source directory to a target directory on a Data Domain system. The `force` option allows the destination directory to be overwritten if it exists. Executing the fast copy operation displays a progress status dialog box.

---

### Note

A fast copy operation makes the destination equal to the source, but not at a specific time. There are no guarantees that the two are or were ever equal if you change either folder during this operation.

---

## Perform a Fast Copy Operation

### Procedure

1. Using the Data Domain System Manager, click the **Data Management > File System** tabs and select **Fast Copy** from **More Tasks**.

The Fast Copy dialog box is displayed.

2. In the **Source** text box, enter the pathname of the directory where the data to be copied resides. For example, `/data/col1/backup/.snapshot/snapshot-name/dir1`.

---

### Note

`col1` uses an ell followed by the number 1.

3. In the **Destination** text box, enter the pathname of the directory where the data will be copied to. For example, `/data/col1/backup/dir2`. This destination directory must be empty, or the operation fails.
  - If the Destination directory exists, click the checkbox **Overwrite existing destination if it exists**.
4. Click **OK**.
5. In the progress dialog box that appears, click **Close** to exit.

# CHAPTER 6

## Managing Encryption of Data at Rest

This chapter includes:

- ◆ [How Encryption of Data at Rest Works](#)..... 148
- ◆ [Configuring Encryption](#)..... 148
- ◆ [About Key Management](#)..... 149
- ◆ [Key Manager Setup](#)..... 154
- ◆ [Changing Key Managers after Setup](#)..... 157
- ◆ [Checking Settings for Encryption of Data at Rest](#)..... 157
- ◆ [Enabling and Disabling Encryption of Data at Rest](#)..... 157
- ◆ [Locking and Unlocking the File System](#)..... 158
- ◆ [Managing the Encryption Passphrase](#)..... 160

## How Encryption of Data at Rest Works

Data encryption protects user data if the Data Domain system is stolen or if the physical storage media is lost during transit, and eliminates accidental exposure of a failed drive if it is replaced. If an intruder circumvents network security controls and gains access to encrypted data, the data is unreadable and unusable without the proper cryptographic keys.

When data enters the Data Domain system using any of the supported protocols (NFS, CIFS, VTL, DD Boost, and NDMP Tape Server), the stream is segmented, fingerprinted, de-duplicated (global compression), then grouped into multi-segment compression regions, locally compressed, and then encrypted before stored to disk.

Once enabled, the Encryption at Rest feature encrypts all data entering the Data Domain system. You cannot enable encryption at a more granular level.

### CAUTION

**Data that has been stored before the encryption feature is enabled does not automatically get encrypted. To protect all of the data on the system, when you configure encryption, be sure to enable the option to encrypt existing data.**

---

#### Additional Notes:

The `filesys encryption apply-changes` command applies any encryption configuration changes to all data present in the file system during the next cleaning cycle. For more information about this command, see the *EMC DD OS Command Reference Guide*.

You can use all of the currently supported backup applications described in the Backup Application Matrix on the Support portal with Encryption of Data at Rest.

Data Domain Replicator software can be used with the encryption option, enabling encrypted data to be replicated using collection, directory, MTree, or application-specific managed file replication and with the various topologies. Each replication form works uniquely with encryption and offers the same level of security. For more information, see the section [Using Encryption of Data at Rest with Replication on page 304](#).

Files locked using the Data Domain Retention Lock software options can be stored, encrypted, and replicated.

The autosupport feature includes information about the state of encryption on the Data Domain system:

- ◆ Whether or not encryption is enabled
- ◆ The Key Manager in effect and which keys are used
- ◆ The encryption algorithm that is configured
- ◆ The state of the file system

## Configuring Encryption

If the Encryption Status on the **Data Management > File System > Encryption** tab shows Not Configured, click **Configure** to set up encryption on the Data Domain system.

This procedure includes configuring a key manager.

Complete the following information:

- ◆ Passphrase

In the text fields, provide the user name and password of a Security Officer account (an authorized user in the Security User group on that Data Domain system). Enter a passphrase. See [Managing the Encryption Passphrase on page 160](#).

- ◆ Algorithm
  - Select an encryption algorithm from the drop-down list or accept the default AES 256-bit (CBC).  
The AES 256-bit Galois/Counter Mode (GCM) is the most secure algorithm but it is significantly slower than the Cipher Block Chaining (CBC) mode.
  - Determine what data is to be encrypted: existing and new or only new. Existing data will be encrypted during the first cleaning cycle after the file system is restarted. Encryption of existing data can take longer than a standard file system cleaning operation.
- ◆ Key Manager (select one of the two)
  - Embedded Key Manager  
By default, the Data Domain Embedded Key Manager is in effect after you restart the file system unless you configure the RSA DPM Key Manager.  
  
You can enable or disable key rotation. If enabled, enter a rotation interval between one-to-12 months.
  - RSA DPM Key Manager

---

#### Note

The RSA DPM Key Manager requires setup on both an RSA DPM server and on the Data Domain system. Follow the instructions in [RSA DPM Key Manager Encryption Setup on page 154](#) before selecting the RSA DPM Key Manager in the Data Domain interface. You can enable encryption using the Embedded Key Manager before configuring the RSA DPM Key Manager. You can then switch to the RSA DPM Key Manager after performing the RSA DPM Key Manager Encryption Setup and following the procedure described in [Changing Key Managers after Setup on page 157](#).

---

- ◆ Enter the name or the IP address of the Key Manager server.
- ◆ Choose the key class that the Data Domain system will use to generate the key from the menu.
- ◆ Enter the port number (443 is the default).
- ◆ Select whether the imported host certificate is FIPS compliant. The default mode is enabled.

The Summary shows your selected configuration values. Review them for correctness. To change a value, click Back to navigate to the page where it was entered and modify it.

A system restart is necessary to enable encryption. To apply the new configuration, select the option to restart the file system.

---

#### Note

Applications may experience an interruption while the file system is restarted.

---

## About Key Management

Encryption keys determine the output of the cryptographic algorithm. They are protected by a passphrase, which encrypts the encryption key before it is stored in multiple locations on disk. The passphrase is user generated and requires both an administrator and a security officer to change it.

A key manager controls the generation, distribution, and lifecycle management of multiple encryption keys. A Data Domain system can use either the Embedded Key Manager or the RSA Data Protection Manager (DPM) Key Manager. Only one can be in effect at a time. When encryption is enabled on a Data Domain system, the Embedded Key Manager is in effect by default. If you configure the RSA DPM Key Manager, it replaces the Embedded Key Manager and remains in effect until you disable it. A file system restart is required for a new key manager to be operational.

Both key managers provide multiple keys, although only one encryption key is active on a Data Domain system. If the RSA DPM Key Manager is configured and enabled, the Data Domain system uses keys provided by the RSA DPM Key Manager Server. If the same DPM Key Manager manages multiple Data Domain systems, all will have the same active key—if they are synced and the Data Domain file system has been restarted. The Embedded Key Manager generates its keys internally.

Both key managers rotate keys and support a maximum of 254 keys. The Embedded Key Manager allows you to specify how many months a key is in effect before being replaced (after the file system is restarted). The RSA DPM Key Manager rotates keys on a regular basis, depending on the key class. The Embedded Key Manager key rotation is managed on the Data Domain system. The DPM Key Manager key rotation is managed on the RSA DPM Key Manager server.

The section covers the following major topics:

- ◆ [Rectifying Lost or Corrupted Keys on page 150](#)
- ◆ [Key Manager Support on page 150](#)
- ◆ [Working with the RSA DPM Key Manager on page 151](#)
- ◆ [How the Cleaning Operation Works on page 153](#)
- ◆ [Working with the Embedded Key Manager on page 153](#)

## Rectifying Lost or Corrupted Keys

You can create a file that contains all of your system's current encryption keys. Your support provider can use this file to import keys back to your system should they become lost or corrupted. It is recommended that you create an export file on a regular basis.

You are prompted for the Security Officer's credential to export the keys. For additional key file protection, you can use a passphrase that differs from the one used in Data Domain system. After exporting, it is recommended that you save the key file in a secure file server not accessible only by authorized users. You must remember the passphrase used for the key file. If the passphrase is lost or forgotten, the Data Domain system cannot import and restore the keys. Enter:

```
# filesys encryption keys export
```

## Key Manager Support

Both Key Managers support all DD OS file system protocols.

### DD Extended Retention

Data Domain systems with DD Extended Retention Software do not support encryption of data at rest. Therefore, DD Extended Retention software cannot be added to those Data Domain systems that are encryption enabled or either have encrypted data on them.

### Replication

When configuring Data Domain systems for directory MTree replication, configure each Data Domain system separately. The two systems can use either the same or a different key class, and the same or different key managers.

For collection replication configuration, the Data Domain system must be configured on the source. After a replication break the original replica Data Domain system has to be

configured for the Key Manager. If not, the Data Domain system continues to use the latest known key.

## Working with the RSA DPM Key Manager

### Encryption Key States

One Activated-RW key is always in effect. If the active key is compromised, the RSA DPM Key Manager provides a new key. When the Data Domain system detects the new key, it issues an alert for the administrator to restart the file system.

Expired keys become read only for the existing data on the Data Domain system, and a new active key is applied to all new data that is ingested. When a key is compromised, the existing data is re-encrypted using the new encryption key after a file system clean is run. If the maximum number of keys is reached, unused keys must be deleted to make room for new keys.

To view information about the encryption keys that are on Data Domain system, open the Data Domain System Manager and go to the **Data Management > File System > Encryption System** tab. Keys are listed by ID number in the **Encryption Key** section of the Encryption page. The following information is given for each key: when a key was created, how long it is valid, its type (RSA DPM or Data Domain), its state, such as Activated-RW or Deactivated, and its post-compression size.

**Table 22** DPM Encryption Key States Supported by Data Domain

State	Definition
Pending-Activated	The key has just been created. After a file system restarts, the key becomes Activated-RW.
Activated-RW and Activated-RO	Both Activated-RW and Activated-RO read the data encrypted with their keys, respectively. Activated-RW is the latest activated key.
De-Activated	A key becomes deactivated when the current time exceeds the validity period. The key is used for reading.
Compromised	The key can only decrypt. After all of the data encrypted with the compromised key is re-encrypted, the state changes to Destroyed Compromised. The keys are re-encrypted when a file system clean is run. You can delete a Destroyed Compromised key, if necessary.
Marked-For-Destroy	You have marked the key as destroyed for the data to be re-encrypted.
Destroyed	After re-encrypting all data encrypted with this key, the DD OS changes it from Marked-For-Destroy to Destroyed. Also, when the key that is destroyed is compromised, it becomes Compromised-Destroyed. You can delete keys that are Destroyed and Compromised-Destroyed.  <b>Note</b> A key is not destroyed in the Data Domain system until a cleaning operation is run and completed.

## How Keys are Kept in Sync with the RSA DPM Key Manager

An automatic key sync is performed every day at midnight. A manual key sync is required only if you cannot wait for the scheduled sync. Whenever new keys are synced on the Data Domain system, an alert is generated. This alert is cleared after the Data Domain file system is restarted.

After the RSA DPM Key Manager Server generates new keys, click the **Sync** button to have them display in the Encryption Key list on the Data Domain System Manager's Encryption tab.

---

**Note**

A file system restart is necessary if keys have changed since the last sync.

**Procedure**

1. Using the Data Domain System Manager, select the Data Domain system you are working with in the Navigational pane.

---

**Note**

Always perform Data Domain System Manager functions on the system you have selected in the Navigational pane.

2. Click **Data Management > File System** and then select the Encryption tab.
3. In the Encryption Keys section, select the **RSA DPM** key, and click **Sync**.

## Destroying a Key (RSA DPM Key Manager)

Destroy a key if you do not want any data to be encrypted with it. Follow this procedure only when you are running out of keys (no more keys can be added). The maximum number of keys is 254. This procedure requires security officer credentials. For information about the security officer, see [Creating Local Users on page 71](#) and [Enabling Security Authorization on page 74](#).

To change an RSA DPM key to a state in which it can be deleted:

**Procedure**

1. Deactivate the key on the RSA DPM Server.
2. Restart the file system for the key to be deactivated on the Data Domain system.
3. Using the DD System Manager, navigate to the **Data Management > File System > Encryption** tabs.
4. In the Encryption Keys section, select the key in the list to be destroyed.
5. Click **Destroy....**
6. Enter your security officer user name and password.
7. Confirm that you want to destroy the key by clicking **OK**.
8. After the file system clean has run, the key state changes to Destroyed.
9. Delete the key. See [How to Delete a Key on page 152](#).

## How to Delete a Key

You need to delete a key only when the number of keys has exceeded the maximum 254 limit.

This procedure requires security officer credentials.

You can delete Key Manager keys that are in the Destroyed or Compromised-Destroyed states.

**Procedure**

1. Using the Data Domain System Manager, navigate to the **Data Management > File System > Encryption** tabs.



2. In the Encryption Keys section, select the key or keys in the list to be deleted.
3. Click **Delete Keys**.
4. Enter your security officer user name and password.
5. Confirm that you want to delete the key or keys by clicking **OK**.

## How the Cleaning Operation Works

Encryption affects the performance of cleaning operations during which all data encrypted with the Compromised or Marked-For-Destroyed keys is re-keyed using the Activated-RW key. At the end of the cleaning operation, there will be no data that is encrypted with the Compromised or Marked-For-Destroyed keys. Also, any data written by the cleaning operation is encrypted with the Activated-RW key.

## Working with the Embedded Key Manager

The Embedded Key Manager does not support a compromised key state. If a key is compromised, destroy the key as described below and issue the file system clean command to re-key all affected data to the latest Activated-RW key.

After the key rotation policy is configured, a new key is automatically created at the next rotation. An alert informs you of the creation of a new key. You need to perform a file system restart to put the new key in effect after which the old key is deactivated. You can disable the key rotation policy by clicking the disable button associated with the Embedded Key Manager Key's rotation status.

---

### Note

If the Data Domain system has already used 254 keys on the system, deleting a key is not sufficient. You must force a new key to be created by issuing the command `filesys encryption embedded-key-manager keys create`. Otherwise, a new key is not created until the next key-rotation date.

---

## Destroying a Key

To destroy a key:

### Procedure

1. Using the Data Domain System Manager, navigate to the **Data Management > File System > Encryption** tabs.
2. In the Encryption Keys section, select the key in the list to be destroyed.
3. Click **Destroy Key**.
4. Enter your security officer user name and password.
5. Confirm that you want to destroy the key by clicking **OK**.
6. After the file system clean has run, the key state changes to Destroyed.
7. Delete the key.

### After you finish

To reconfigure the embedded-key-manager key rotation policy on the destination Data Domain system after a collection replication break, enter:

```
# filesys encryption embedded-key-manager keys set key-rotation-policy {months | none }
```

To use a new key, enter:

```
# fileys encryption embedded-key-manager keys create  
# fileys restart
```

## Deleting a Destroyed Key

To delete a key that is in a Destroyed state:

### Procedure

1. Issue the `delete` command after a `fileys clean` is completed at least one time on the destination Data Domain system.
2. Issue the following commands to delete a key in Destroyed state:

```
# fileys clean start
```

3. After the clean completes, enter:

```
# fileys encryptions keys delete keyid
```

## Key Manager Setup

Follow the instructions for the type of key manager you are using:

- ◆ Data Domain Embedded Key Manager, [Configuring Encryption on page 148](#)
- ◆ [RSA DPM Key Manager Encryption Setup on page 154](#)

## RSA DPM Key Manager Encryption Setup

DPM Key Manager must be set up on both the RSA DPM Server and on the Data Domain system.

### Performing this Setup on the RSA DPM Server

The main steps for setting up the RSA DPM Server (using its graphical user interface) are as follows:

---

#### Note

See the latest version of the *RSA Data Protection Manager Server Administrator's Guide* for more information about each step of this procedure.

Algorithm and cipher mode settings set on the RSA DPM Key Manager Server are ignored by the Data Domain system. Configure these settings on the Data Domain system.

---

#### Procedure

1. Create an identify for the Data Domain system using the X509 certificate. A secure channel is created based on this certificate.
2. Create a key class with the proper attributes:
  - Key length: 256 bits.
  - Duration: For example, six months or whatever matches your policy.
  - Auto-key generation: Select to have keys automatically generated.

---

**Note**

Multiple Data Domain systems can share the same key class. For more information about key classes, see [About RSA DPM Key Classes on page 155](#).

---

3. Create an identity using the Data Domain system's host certificate as its identity certificate. The identity and the key class have to be in the same identity group.
4. Import the certificates. See [Importing the Certificates on page 155](#).

## About RSA DPM Key Classes

The Data Domain system retrieves a key from RSA DPM Key Manager by key class. A key class is a specialized type of security class used by the RSA DPM Key Manager that groups cryptographic keys with similar characteristics.

The RSA DPM Key Manager Server allows a key class to be set up to either return the current key, or to generate a new key each time. The Data Domain system supports only the key classes configured to return the current key. *Do not use a key class that is configured to generate a new key each time.*

---

**Note**

If the key length is not 256 bits, the DPM configuration will fail.

---

## Importing the Certificates

DD OS does not support the RSA DPM Key Manager Server's Auto Registration Certificate capability, which uploads an auto registered certificate directly, or imports multiple certificates. This means that you must import the CA and Host certificates for a Data Domain system.

---

**Note**

You must obtain CA and Host certificates that are compatible with the RSA DPM Key Manager. You can request these certificates from third-party certificates authorities, or create them using appropriate SSL utility tools.

DD OS supports certificates without any extension and certificates with server and client extensions for use with both the Data DD Manager and RSA DPM Key Manager. Certificates with client extensions are supported only by RSA DPM Key Manager, and certificates with server extensions are supported only by the DD System Manager.

Alerts:

---

- ◆ If HTTPS fails to restart due to corrupted imported certificates, self-signed certificates are used. If this occurs, a managed alert, UnusableHostCertificate, is issued. To clear the alert, delete the corrupted certificates and re-import new certificates.
- ◆ If imported certificates are removed, for example during a system headswap and the imported certificates fail to copy over, a managed alert, MissingHostCertificate, is issued. Re-import the certificates to clear the alert.

After obtaining the certificates, import them to the Data Domain system as follows:

---

**Note**

The following prerequisites must be met:

- ◆ The Host certificate should be in PKCS12 format.
  - ◆ The CA certificate should be in PEM format.
  - ◆ If the system passphrase is not set, you cannot import the host certificate. The passphrase is set when you enable encryption. To change it, see [Managing the Encryption Passphrase on page 160](#).
- 

**Procedure**

1. Configure the RSA DPM Key Manager Server to use the CA and Host certificates. For instructions, see the *RSA DPM Key Manager Server Administration Guide*.
2. Import the certificates by redirecting the certificate files using `ssh` command syntax. See the *EMC DD OS Command Reference Guide* for details.

```
ssh sysadmin@<Data-Domain-system> adminaccess certificate import {host
password password |ca } < path_to_the_certificate
```

For example, to import the host certificate `host.p12` from your personal computer's desktop over to the Data Domain system `DD1` using `ssh`, enter:

```
# ssh sysadmin@DD1 adminaccess certificate import host password
abc123 < C:\host.p12
```

3. Import the CA certificate, for example, `ca.pem`, from your desktop to `DD1` via SSH by entering:

```
# ssh sysadmin@DD1 adminaccess certificate import ca < C:\ca.pem
```

## Performing this Setup on the Data Domain System

Using the DD System Manager for the Data Domain system setup, follow these steps:

**Procedure**

1. Complete the DPM Key Manager setup on the RSA DPM Server.
2. The Data Domain system must be able to resolve its own IP address using its hostname. If this mapping has not been added to the DNS server, use this command line to add the entry to the `/etc/hosts` file:

```
# net hosts add ipaddr host-list
```

where `ipaddr` is the IP address of Data Domain system and `host-list` is the hostname of the Data Domain system.

---

**Note**

By default, the `fips-mode` is enabled. If the PKCS #12 client credential is not encrypted with the FIPS 140-2 approved algorithm, such as RC2, then you must disable `fips-mode`. See the *Data Domain Operating System Command Reference Guide* for information about disabling `fips-mode`.

---

3. Log into the DD System Manager and select the Data Domain system you are working with in the Navigation panel.

**Note**

Always perform DD System Manager functions on the system you have selected in the Navigation panel.

4. Click the **Data Management** › **File System** › **Encryption** tab.
5. Follow the instructions in [Configuring Encryption on page 148](#) and select the **DPM Key Manager**. If encryption has already been set up, following the instructions in [Changing Key Managers after Setup on page 157](#).

## Changing Key Managers after Setup

**Procedure**

1. Using the DD System Manager, select the Data Domain system you are working with in the Navigation panel.
2. Click **Data Management** › **File System** › **Encryption** tab.
3. Under Key Management, click **Configure**.
4. Enter your security officer username and password.
5. Select which Key Manager to use.
  - **Embedded Key Manager:** Select to enable or disable key rotation. If enabled, enter a rotation interval between 1-to-12 months. Select **Restart the file system now**, and click OK.
  - **RSA DPM Key Manager:** Enter the server name, key class, port (the default is 443), and whether the imported host certificate is FIPS compliant. The default mode is enabled. Select **Restart the file system now**, and click OK.

## Checking Settings for Encryption of Data at Rest

To check the settings for the Encryption feature, click the **Data Management** › **File System** › **Encryption** tabs. The currently used Key Manager is shown as Enabled. For a description of the Encryption settings, see [About the Encryption View on page 137](#).

## Enabling and Disabling Encryption of Data at Rest

After configuring Encryption, the status is enabled and the Disabled button is active. When Encryption is disabled, the Enabled button is active.

### Enable Encryption of Data at Rest

Use the Data Domain System Manager to enable the Encryption feature:

**Procedure**

1. Using the Data Domain System Manager, select the Data Domain system you are working with in the Navigational pane.
2. In the Encryption view, click the **Enable** button.
3. In the Enable Encryption dialog box, select **Restart** the file system and click **OK**.
4. The Configure Encryption Status shows the implementation status. Click **Close** when the process is complete or **OK** to exit.

---

**Note**

Applications may experience an interruption while the file system is restarted.

---

**After you finish**

Encryption will be enabled once the file system is restarted.

## Disable Encryption of Data at Rest

Use the Data Domain System Manager to disable the Encryption feature:

**Procedure**

1. Using the Data Domain System Manager, select the Data Domain system you are working with in the Navigational pane.
2. In the Encryption view, click the **Disable** button.  
The Disable Encryption dialog box is displayed.
3. In the Security Officer Credentials area, enter the user name and password of a security officer.
4. Select **Restart the file system now** and click **OK**.

The data will be unencrypted during the next cleaning cycle.

## Locking and Unlocking the File System

Use this procedure when an encryption-enabled Data Domain system (and its external storage devices) are being transported, or if you want to lock a disk that is being replaced. The procedure requires two accounts: Security Officer and System Administration roles.

**Procedure**

1. Using the DD System Manager, select the Data Domain system you are working with in the Navigation panel.
2. Go the File System Lock area of the **Data Management > File System > Encryption** view.  
The Status shows whether the file system is Locked or Unlocked.
3. Disable the file system by clicking the **Disabled** button in the File System status area.
4. Use either the procedure [Locking the File System on page 158](#) or [Unlock the File System on page 159](#).

## Locking the File System

To lock the file system, Encryption must be enabled and the file system must be disabled.

**Procedure**

1. Using the DD System Manager, select the Data Domain system you are working with in the Navigation panel.
2. Navigate to the File System Lock area of the **Data Management > File System > Encryption** view. Click **Lock File System**.
3. In the text fields of the Lock File System dialog box, provide:
  - The user name and password of a Security Officer account (an authorized user in the Security User group on that Data Domain system).

- The current and a new passphrase.
4. Click **OK**.

This procedure creates a new passphrase and destroys the cached copy of the current passphrase. Therefore, anyone who does not possess the new passphrase will not be able to decrypt the data.

---

#### Note

Changing the passphrase requires two-user authentication to protect against the possibility of a rogue employee's shredding the data.

---

#### **CAUTION**

**Be sure to take care of the passphrase. If the passphrase is lost, you will never be able to unlock the file system and access the data. The data will be irrevocably lost.**

5. Shut down the system:

---

#### **CAUTION**

**Do not use the chassis power switch to power off the system. Enter the following command on the command line instead.**

```
# system poweroff The 'system poweroff' command shuts down
the system and turns off the power. Continue? (yes|no|?)
[no]:
```

6. Transport the system or remove the disk being replaced.
7. Power on the system and continue with [Unlock the File System on page 159](#).

## Unlock the File System

This procedure prepares an encrypted file system for use after it has arrived at its destination.

### Procedure

1. Using the Data Domain System Manager, select the Data Domain system you are working with in the Navigational pane.
2. Navigate to the **File System Lock** area of the **Data Management > File System > Encryption** view. Click **Unlock File System**.

The Unlock File System dialog box is displayed.

3. In the text fields, provide the current passphrase.
4. Click **OK**.

The Enable File System Status dialog box displays.

5. Click **Close** to exit.

If the passphrase is incorrect, the file system does not start and the system reports the error. Re-enter the correct passphrase, as directed in the previous step.

## Changing the Encryption Algorithm

To change the encryption algorithm:

### Procedure

1. Using the Data Domain System Manager, select the Data Domain system you are working with in the Navigational pane.
2. Go to the **Data Management > File System** and select the **Encryption** tab.
3. To change the Encryption Algorithm used to encrypt the Data Domain system, click **Change Algorithm**.

The Change Algorithm dialog box is displayed.

4. Select an encryption algorithm from the drop-down list or accept the default AES 256-bit (CBC).

The AES 256-bit Galois/Counter Mode (GCM) is the most secure algorithm but it is significantly slower than the Cipher Block Chaining (CBC) mode.

---

#### Note

To reset the algorithm to the default AES 256-bit (CBC) click Reset to default.

5. Determine what data will be encrypted:
  - To encrypt existing and new data on the system, select **Apply to Existing data, Restart file system now**, and click **OK**. Existing data will be encrypted during the first cleaning cycle after the file system is restarted.

---

#### Note

Encryption of existing data can take longer than a standard file system clean operation

- To encrypt only new data, select **Restart file system now** and click **OK**.
6. The status is displayed. Click **Close** when the process is complete.

---

#### Note

Applications may experience an interruption while the file system is restarted.

## Managing the Encryption Passphrase

The encryption passphrase is a human-readable (understandable) key (like a smart card) which is used to generate a machine usable AES 256 encryption key.

The administrator can change the passphrase without having to manipulate the actual encryption keys. Changing the passphrase indirectly changes the encryption of the keys, but does not affect user data or the underlying encryption key. To change the passphrase, see [Changing the Encryption Passphrase on page 161](#).

The passphrase allows a Data Domain system to be transported with encryption keys on the system, but without the passphrase being stored on it. If the system is stolen in transit, an attacker cannot easily recover the data; at most, they can recover the encrypted user data and the encrypted keys.

The passphrase is stored internally on a hidden part the Data Domain storage subsystem. This allows the Data Domain system to boot and continue servicing data access without any administrator intervention.

Changing the passphrase requires two-user authentication to protect against shredding the data.



## Changing the Encryption Passphrase

The passphrase is set when encryption is enabled, but it can be changed with the following procedure:

### Procedure

1. Using the DD System Manager, click the **Data Management** > **File System** > **Encryption** tabs.
2. To change the Encryption Passphrase, click **Change Passphrase**.  
The Change Passphrase dialog box is displayed.
3. In the text fields, provide:
  - The user name and password of a Security Officer account (an authorized user in the Security User group on that Data Domain system).
  - The current and a new passphrase.
4. To enable the encryption feature, click the checkbox for **Enable the file system now**.
5. Click **OK**.

### CAUTION

**Be sure to take care of the passphrase. If the passphrase is lost, you will never be able to unlock the file system and access the data. The data will be irrevocably lost.**

---



# CHAPTER 7

## Working with DD Retention Lock

This chapter includes:

- ◆ [About DD Retention Lock Software](#)..... 164
- ◆ [Supported Data Access Protocols](#)..... 166
- ◆ [Enabling DD Retention Lock on an MTree](#)..... 167
- ◆ [Client-Side Retention Lock File Control](#)..... 170
- ◆ [System Behavior with DD Retention Lock](#)..... 175

## About DD Retention Lock Software

When data is locked on an MTree that is enabled with the DD Retention Lock software option, the DD Retention Lock software helps ensure that data integrity is maintained where any data that is locked cannot be overwritten, modified, or deleted for a user-defined retention period of up to 70 years.

---

### Note

DD Retention Lock is not supported on GDA systems.

---

There are two DD Retention Lock editions:

- ◆ *EMC Data Domain Retention Lock Governance Edition* retains the functionality of the Data Domain Retention Lock product prior to DD OS 5.2. You can use Data Domain Retention Lock Governance software to define retention policies on data that is to be retained for a specific period of time to meet internal IT governance policies implemented by the system administrator.
- ◆ *EMC Data Domain Retention Lock Compliance Edition* enables you to meet the strictest data permanence requirements of regulatory standards, such as those of SEC 17a-4(f). The full list of regulatory standards includes:
  - CFTC Rule 1.31b
  - FDA 21 CFR Part 11
  - Sarbanes-Oxley Act
  - IRS 98025 and 97-22
  - ISO Standard 15489-1
  - MoREQ2010

For certification information, see *Compliance Assessment - Summary and Conclusions: EMC Data Domain Retention Lock Compliance Edition* at:

[http://powerlink.emc.com/km/live1/en\\_US/Communications/Testimonial\\_Profile/cohasset-associates\\_data-domain-retention-lock-compliance-ar.pdf](http://powerlink.emc.com/km/live1/en_US/Communications/Testimonial_Profile/cohasset-associates_data-domain-retention-lock-compliance-ar.pdf)

(Login is required.)

Compliance with these standards ensures that files locked on the Data Domain system using the Data Domain Retention Lock Compliance software cannot be altered or destroyed before the retention period expires. Data Domain Retention Lock Compliance requires a security officer for implementation of policies. An audit log file is accessible by the administrator or security officer.

Each edition requires a separate, add-on license, and either or both can be used on a single Data Domain system.

The retention locking protocol is the same for DD Retention Lock Governance and Compliance. The differences in use stem from the system behavior for DD Retention Lock Compliance, since it places strict restrictions to meet compliance requirements. For specifics, refer to “System Behavior with Retention Lock” on page 269. For an overview, see the EMC Data Domain *Retention Lock Software (White Paper)* available at:

[http://powerlink.emc.com/km/live1/en\\_US/Offering\\_Technical/White\\_Paper/h10666-data-domain-retention-lock-wp.pdf](http://powerlink.emc.com/km/live1/en_US/Offering_Technical/White_Paper/h10666-data-domain-retention-lock-wp.pdf)

(Login is required.)

The DD Retention Lock Governance edition does not require a security officer and provides a higher degree of flexibility for archive data retention on DD systems.

For archive compliance storage requirements, SEC rules require that a separate copy of retention-locked data must be stored with the same retention requirements as the original. Retention-locked files can be replicated via DD Replicator software to another Data Domain system. If a retention-locked file is replicated, it remains retention locked on the destination system, with the same level of protection as the source file. For more information see page 269 for using replication with DD Retention Lock Governance, and page 272 for using replication with DD Retention Lock Compliance.

## DD Retention Lock Protocol

Only files that are explicitly committed to be retention-locked files are retention locked on the Data Domain system. Files are committed to be retention-locked files through client-side file commands issued while DD Retention Lock Governance or Compliance is enabled on the MTree containing the files (see [Client-Side Retention Lock File Control on page 170](#)).

---

### Note

Linux, Unix, and Windows client environments are supported.

Files that are written to shares or exports that are not committed to be retained (even if DD Retention Lock Governance or Compliance is enabled on the MTree containing the files) can be modified or deleted at any time.

Retention locking prevents any modification or deletion of files under retention from occurring directly from CIFS shares or NFS exports during the retention period specified by a client-side *atime* update command (see [Setting Retention Locking on a File on page 171](#)). Some archive applications and backup applications can issue this command when appropriately configured (see [Supported Data Access Protocols on page 166](#)). Applications or utilities that do not issue this command cannot lock files using DD Retention Lock.

Retention-locked files are always protected from modification and premature deletion, even if retention locking is subsequently disabled or if the retention-lock license is no longer valid.

You cannot rename or delete non-empty folders or directories within an MTree that is retention-lock enabled. However, you can rename or delete empty folders or directories and create new ones.

The retention period of a retention-locked file can be extended (but not reduced) by updating the file's *atime* (see [Extending Retention Locking on a File on page 173](#)).

For both DD Retention Lock Governance and Compliance, once the retention period for a file expires, the file can be deleted using a client-side command, script, or application. However, the file cannot be modified even after the retention period for the file expires. The Data Domain system never automatically deletes a file when its retention period expires.

## DD Retention Lock Flow

The general flow of activities with DD Retention Lock is as follows:

1. Enable MTrees for DD Retention Lock Governance or Compliance retention locking using the DD System Manager or DD OS commands issued from the system console. (See [Extending Retention Locking on a File on page 173](#).)
2. Commit files to be retention locked on the Data Domain system using client-side commands issued by an appropriately configured archiving or backup application, manually, or via scripts. (See [Client-Side Retention Lock File Control on page 170](#).)

---

**Note**

Windows clients may need to download utility programs for DD OS compatibility.

---

3. Optionally, extend file retention times using client-side commands. (See [Extending Retention Locking on a File on page 173.](#))
4. Optionally, delete files with expired retention periods using client-side commands. (See [Deleting or Expiring a File on page 174.](#))

## Supported Data Access Protocols

DD Retention Lock is compatible with industry-standard, NAS-based Write-Once-Read-Many (WORM) protocols, and integration is qualified with archive applications such as Symantec Enterprise Vault, EMC SourceOne, EMC Cloud Tiering Appliance, EMC DiskXtender, and so on. Customers using backup applications such as CommVault can also develop custom scripts to use the EMC Data Domain Retention Lock software option.

To check whether an archive application is tested and certified for DD Retention Lock, refer to the *EMC Data Domain Archive Product Compatibility Matrix*.

The protocol support of DD Retention Lock software is as follows:

- ◆ NFS is supported with both DD Retention Lock Governance and Compliance.
  - ◆ CIFS is supported with both DD Retention Lock Governance and Compliance.
  - ◆ VTL is supported with DD Retention Lock Governance, but not with DD Retention Lock Compliance.
- 

**Note**

Virtual tapes, here referred to as *tapes*, are represented as files on the Data Domain file system.

- When you create a storage pool, a collection of tapes that map to a directory on the file system, you are creating an MTree, unless you specifically select to create the older style directory pool (for backward compatibility). You can also convert storage pools created prior to DD OS 5.3 to MTrees. These MTrees can be retention locked and replicated.
  - You can retention-lock one or more tapes using the `vtl tape modify` command, described in the *EMC Data Domain Operating System Command Reference Guide*.  
The `mtree retention-lock revert path` command can be used to revert the retention-locked state of tapes locked with the `vtl tape modify` command. After the tape is unlocked, updates can be made to it. The unlocked state won't be visible via the DD System Manager or CLI until the VTL service is disabled then enabled; however, updates will be applied to the unlocked tape. This capability is only for the DD Retention Lock Governance Edition.
  - The retention time for tapes can be displayed using the `vtl tape show` command with the `time-display retention` argument.
  - You can retention-lock an individual tape using the DD System Manager. For instructions, see [Working with DD Virtual Tape Library on page 249.](#)
-

- ◆ DD Boost is not supported with either DD Retention Lock Governance or DD Retention Lock Compliance.  
Currently, EMC DD Retention Lock software is not integrated with the DD Boost protocol. If client-side scripts are used to retention-lock backup files or backup images, and if a backup application (Symantec NetBackup, for example) is also used on the system via DD Boost, be aware that the backup application does not share the context of the client-side scripts. Thus, when a backup application attempts to expire or delete files that were retention locked via the client-side scripts, space is not released on the EMC Data Domain system.

Data Domain recommends that administrators change their retention period policy to align with the retention lock time. This applies to all the backup applications that are integrated with DD Boost: Avamar, Symantec NetBackup, Symantec Backup Exec, EMC NetWorker, and so on.

## Enabling DD Retention Lock on an MTree

Only files within DD Retention Lock Governance or Compliance enabled MTrees can be retention-locked.

MTrees enabled for DD Retention Lock Compliance cannot be converted to DD Retention Lock Governance MTrees and vice versa.

The procedures that follow show how to enable MTrees for either DD Retention Lock Governance or DD Retention Lock Compliance.

## Enabling DD Retention Lock Governance on an MTree

Complete the steps below to add a DD Retention Lock Governance license to a system, and then enable DD Retention Lock Governance on one or more MTrees.

### Procedure

1. Log in to the DD System Manager.  
The DD System Manager window appears with `DD Network` in the Navigation Panel.
2. Select a Data Domain system.  
In the Navigation Panel, expand `DD Network` and select a system.
3. Add the DD Retention Lock Governance license, if it is not present.  
Click the `System Settings > Licenses` tabs. If `RETENTION-LOCK-GOVERNANCE` is not listed, add it as follows:
  - a. In the Licensed Features Panel, click **Add**. The Add License Key dialog box appears.
  - b. In the License Key text box, enter the license key.

---

### Note

License keys are case-insensitive. Include the hyphens when entering keys.

---

- c. Click **OK**. The added license appears in the license list.
4. Select an MTree for retention locking.  
Click the `Data Management > MTree` tab, then the checkbox for the MTree you want to use for retention locking. (You can also create an empty MTree and add files to it later.)
  5. Display information for the MTree you selected.  
Click the MTree Summary tab.

6. Bring up the MTree's Modify Retention Lock dialog box.

Scroll down to Retention Lock and click **Edit** in the Retention Lock area. The dialog box appears.

7. Enable DD Retention Lock Governance on the MTree and change the default minimum and maximum retention lock periods for the MTree, if necessary.

Perform the following actions in the Modify Retention Lock dialog box:

- a. Click the **Enable** checkbox to enable DD Retention Lock Governance on the MTree.
- b. To change the minimum or maximum retention period for the MTree, go to the Retention Period Panel and modify the minimum or maximum time period:

Type a number for the interval in the text box (for example, 5 or 14).

From the drop-down list, select an interval (minutes, hours, days, years).

---

#### Note

Specifying a minimum retention period of less than 12 hours, or a maximum retention period longer than 70 years, results in an error.

---

- c. Click **OK** to save your settings.

After you close the Modify Retention Lock dialog box, updated MTree information appears.

8. Check retention lock information for the MTree.

Note the following retention lock fields:

- Top:
  - The Status field indicates the Read/Write access for the MTree, the type of retention locking on the MTree, and whether retention locking is enabled or disabled.
- Bottom:
  - The Status field indicates whether retention locking is enabled for the MTree.
  - The Retention Period field indicates minimum and maximum retention periods for the MTree. The retention period specified for a file in the MTree must be equal to or greater than the minimum retention period and equal to or less than the maximum retention period.
  - The UUID field is a unique identification number generated for the MTree.

Repeat steps 4–8 to enable additional MTrees.

---

#### Note

To check retention lock configuration settings for any MTree, select the MTree in the Navigation Panel, then click the Summary tab.

---

9. When you are finished, exit the DD System Manager.

#### After you finish

Go to [Client-Side Retention Lock File Control on page 170](#) to retention-lock files in a retention-lock-enabled MTree.

## Enabling DD Retention Lock Compliance on an MTree

Complete the steps below to add a DD Retention Lock Compliance license to a system, set up a system administrator and one or more security officers, configure and enable the



system to use DD Retention Lock Compliance software, and then enable DD Retention Lock Compliance on one or more MTrees.

---

**Note**

For DD OS 5.4, the DD System Manager does not support DD Retention Lock Compliance.

---

**Procedure**

1. Add the DD Retention Lock Compliance license on the system, if it is not present.

- a. First, check whether the license is already installed.

```
license show
```

- b. If the RETENTION-LOCK-COMPLIANCE feature is not displayed, install the license.

```
license add license-key
```

---

**Note**

License keys are case-insensitive. Include the hyphens when entering keys.

---

2. Set up and one or more security officer user accounts according to Role-Based Access Control (RBAC) rules (see the *EMC Data Domain Operating System Command Reference Guide*).

- a. In the sysadmin role, add a security officer account.

```
user add user role security
```

- b. Enable the security officer authorization.

```
authorization policy set security-officer enabled
```

3. Configure and enable the system to use DD Retention Lock Compliance.
- 

**Note**

Enabling DD Retention Lock compliance enforces many restrictions on low-level access to system functions used during troubleshooting. Once enabled, the only way to disable DD Retention Lock Compliance is to completely initialize and reload the system, which will result in destroying all data on the system.

---

- a. Configure the system to use DD Retention Lock Compliance.

```
system retention-lock compliance configure
```

The system automatically reboots.

- b. After the reboot process is complete, enable DD Retention Lock Compliance on the system.

```
system retention-lock compliance enable
```

4. Enable compliance on an MTree that will contain retention-locked files.

```
mtree retention-lock enable mode compliance mtree mtree-path
```

---

**Note**

Compliance cannot be enabled on /backup, DD Boost storage unit MTrees, or pool MTrees.

---

5. To change the default minimum and maximum retention lock periods for a compliance-enabled MTree, enter the following commands with security officer authorization.

```
mtree retention-lock set min-retention-period period mtree
mtree-path
mtree retention-lock set max-retention-period period mtree
mtree-path
```

---

**Note**

Specifying a minimum retention period of less than 12 hours, or a maximum retention period longer than 70 years, results in an error.

---

Repeat steps 4 and 5 to enable additional MTrees.

**After you finish**

Go to [Client-Side Retention Lock File Control on page 170](#) to retention lock files in a retention-lock-enabled MTree.

## Client-Side Retention Lock File Control

This section describes the DD Retention Lock client command interface for locking files stored on EMC Data Domain systems. Client commands are the same for DD Retention Lock Governance and Compliance. Linux, Unix, and Windows client environments are supported; however, Windows clients may need to download utility programs with commands to lock files.

---

**Note**

If your application already supports industry-standard WORM, writing a WORM file to a DD Retention Lock Governance or Compliance enabled MTree will lock the file on the Data Domain system. The retention time in the application should agree with the DD Retention Lock settings. You do not need to use the commands described in this section. To check whether an application is tested and certified for the DD Retention Lock, refer to the *EMC Data Domain Archive Product Compatibility Matrix*.

---

**Note**

Some client machines using NFS, but running a legacy OS, cannot set retention time later than 2038. The NFS protocol doesn't impose the 2038 limit and allows to specifying times until 2106. Further, DD OS doesn't impose the 2038 limit.

---

Client-side commands are used to manage the retention locking of individual files. These commands apply to all retention-lock-capable Data Domain systems and must be issued in addition to the setup and configuration of the DD Retention Lock software on the Data Domain system (see [Enabling DD Retention Lock on an MTree on page 167](#)).

**Required Tools for Windows Clients**

You need the `touch.exe` command to perform retention-locking from a Windows-based client.

To obtain this command, download and install utilities for Linux/Unix-based applications according to your Windows version. These utilities are best recommendations from EMC and should be used per customer environment.

---

**Note**

The `touch` command for Windows may have a different format than the Linux examples in this chapter.

- ◆ For Windows 8, Windows 7, Windows Server 2008, Windows Vista, Windows Server 2003, and Windows XP:  
<http://sourceforge.net/projects/unxutils/files/latest>
  - ◆ For Windows Server 2008, Windows Vista Enterprise, Windows Vista Enterprise 64-bit edition, Windows Vista SP1, Windows Vista Ultimate, and Windows Vista Ultimate 64-bit edition:  
<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=23754>
  - ◆ For Windows Server 2003 SP1 and Windows Server 2003 R2:  
<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=20983>
- 

Follow the installation instructions provided and set the search path as needed on the client machine.

**Client Access to Data Domain System Files**

After an MTree is enabled for DD Retention Lock Governance or Compliance, you can:

- ◆ Create a CIFS share based on the MTree. This CIFS share can be used on a client machine.
  - ◆ Create an NFS mount for the MTree and access its files from the NFS mount point on a client machine.
- 

**Note**

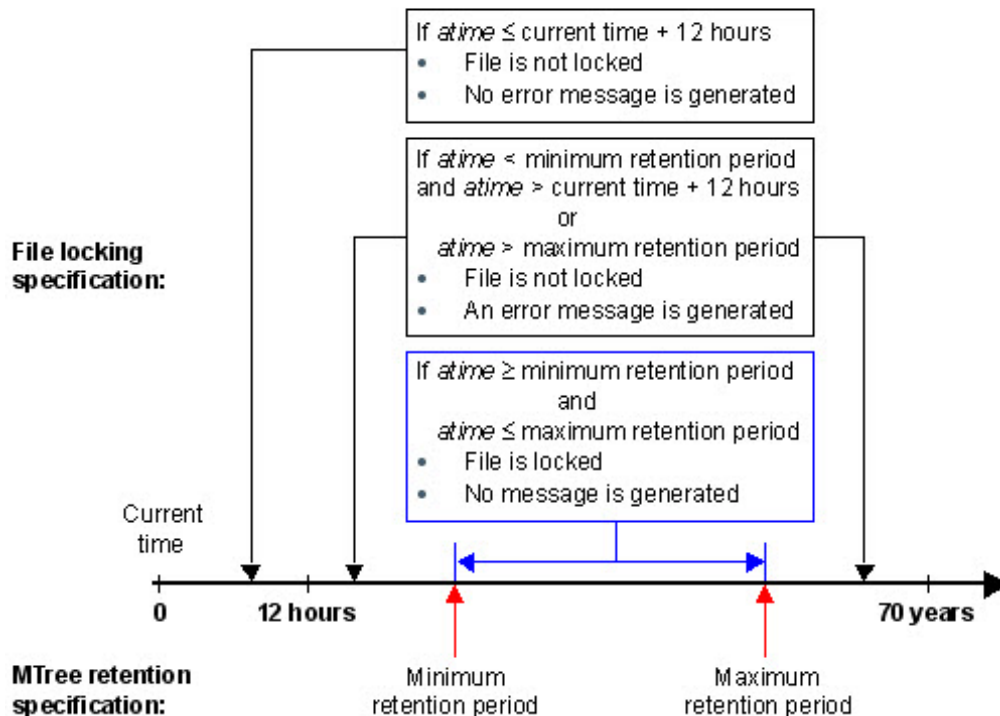
- ◆ The commands listed in this section are to be used only on the client. They cannot be issued through the DD System Manager or CLI.
  - ◆ Command syntax may vary slightly, depending on the utility you are using.
- 

## Setting Retention Locking on a File

To retention lock a file, change the last access time (*atime*) of the file to the desired retention time of the file, that is, the time at which the file can be deleted. This is usually performed by the archive application, and all the archive applications that are qualified on Data Domain systems today (per the *EMC Data Domain Archive Product Compatibility Matrix*) follow the basic locking protocol outlined here.

The future *atime* you specify must respect the minimum and maximum retention periods of the file's MTree (as offsets from the current time), as shown in the next figure.

**Figure 5** Valid and Invalid *atime*s for Retention Locking Files  
**For DD Retention Lock Governance and Compliance**



**Note**

Some client machines using NFS, but running a legacy OS, cannot set retention time later than 2038. The NFS protocol doesn't impose the 2038 limit and allows to specifying times until 2106. Further, DD OS doesn't impose the 2038 limit.

Errors are permission-denied errors (referred to as EACCESS, a standard POSIX error). These are returned to the script or archive application setting the *atime*.

**Note**

A file must be completely written to the Data Domain system before it is committed to be a retention-locked file.

The following command can be used on clients to set the *atime*:

```
touch -a -t [atime] [filename]
```

The format of *atime* is:

```
[ [YY]YY ] MMDDhhmm[.ss]
```

For example, suppose the current date and time is 1 p.m. on January 18, 2012 (that is, 201201181300), and the minimum retention period is 12 hours. Adding the minimum retention period of 12 hours to that date and time results in a value of 201201190100. Therefore, if the *atime* for a file is set to a value greater than 201201190100, that file becomes retention locked.

The following command:

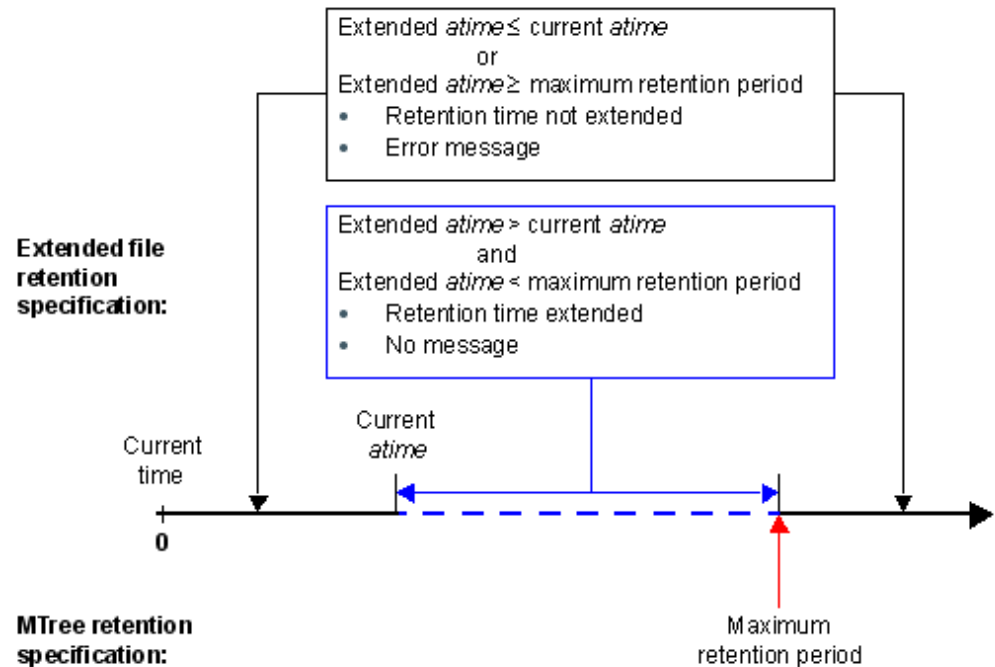
```
ClientOS# touch -a -t 201412312230 SavedData.dat
```

will lock file `SavedData.dat` until 10:30 p.m. December 31, 2014.

## Extending Retention Locking on a File

To extend the retention time of a retention-locked file, set the file's *atime* to a value greater than the file's current *atime* but less than the maximum retention period of the file's MTree (as an offset from the current time), as shown in the next figure.

**Figure 6** Valid and Invalid *atimes* for Extending Retention Locking on Files



For example, changing the *atime* from 201412312230 to 202012121230 using the following command:

```
ClientOS# touch -a -t 202012121230 SavedData.dat
```

will cause the file to be locked until 12:30 p.m. December 12, 2020.

### Note

Some client machines using NFS, but running a very old OS, cannot set retention time later than 2038. The NFS protocol doesn't impose the 2038 limit and allows to specifying times until 2106. Further, DD OS doesn't impose the 2038 limit.

Errors are permission-denied errors (referred to as EACCESS, a standard POSIX error). These are returned to the script or archive application setting the *atime*.

## Identifying a Retention-Locked File

The *atime* value for a retention-locked file is its retention time. To determine whether a file is retention locked, try to set the *atime* of the file to a value earlier than its current *atime*. This action will fail with a permission-denied error if and only if the file is a retention-locked file.

First, list the current *atime* value, and then execute the `touch` command with an earlier *atime* using these commands:

```
ls -l --time=atime [filename]
```

```
touch -a -t [atime] [filename]
```

The following example shows the command sequence:

```
ClientOS# ls -l --time=atime SavedData.dat
202012121230
ClientOS# touch -a -t 202012111230 SavedData.dat
```

If the *atime* of `SavedData.dat` is 202012121230 (12:30 p.m. December 12, 2020) and the `touch` command specifies an earlier *atime*, 202012111230 (12:30 p.m. December 11, 2020), the `touch` command fails, indicating that `SavedData.dat` is retention-locked.

---

#### Note

The `--time=atime` option is not supported in all versions of Unix.

---

## Specifying a Directory and Touching Only Those Files

In this routine, *root directory to start from* contains the files on which you want to change access times using this client system command:

```
find [root directory to start from] -exec touch -a -t
[expiration time] {} \;
```

For example:

```
ClientOS# find [/backup/data1/] -exec touch -a -t 202012121230 {} \;
```

## Reading a List of Files and Touching Only Those Files

In this routine, *name of file list* is the name of a text file that contains the names of the files on which you want to change access times. Each line contains the name of one file. Here is the client system command syntax:

```
touch -a -t [expiration time] `cat [name of file list]`
```

For example:

```
ClientOS# touch -a -t 202012121230 `cat /backup/data1/filelist.txt`
```

## Deleting or Expiring a File

You can delete or expire a file with an expired retention lock using a client application, or delete a file using a standard file-delete command.

Expiring a file using an application makes the file inaccessible to the application. The file may or may not actually be removed from the Data Domain system by the expiration operation. If it is not removed, the application often provides a separate delete operation.

---

#### Note

If the retention period of the retention-locked file has not expired, the delete operation results in a permission-denied error.

---

You must have the appropriate access rights to delete the file, independent of the DD Retention Lock software.

## Using `ctime` or `mtime` on Retention-Locked Files

**ctime**

*ctime* is the last-metadata-change time of a file. It gets set to the current time when any of the follow events occur:

- ◆ A non-retention-locked file is retention locked.
- ◆ The retention time of a retention-locked file is extended.
- ◆ A retention-locked file is reverted.

**mtime**

*mtime* is the last-modified time of a file. It changes only when the contents of the file change. So, the *mtime* of a retention-locked file cannot change.

## System Behavior with DD Retention Lock

System behavior topics are discussed separately for DD Retention Lock Governance and DD Retention Lock Compliance in the sections that follow.

### DD Retention Lock Governance

Certain DD OS commands behave differently when using DD Retention Lock Governance. The following sections describe the differences for each.

#### Replication

---

**Note**

For information on using DD Replicator software, see [Working with DD Replicator on page 297](#).

---

Collection replication, MTree replication, and directory replication replicate the locked or unlocked state of files. That is, files that are governance retention locked on the source are governance retention locked on the destination and have the same level of protection. For replication, both the source and destination systems must have a DD Retention Lock Governance license installed.

Replication is supported between systems that are:

- ◆ Running the same major DD OS version (for example, both systems are running DD OS 5.4.x.x).
  - ◆ Running DD OS versions within the next two consecutive higher or lower major releases (for example, 5.2.x.x to 5.4.x.x or 5.4.x.x to 5.2.x.x). Cross-release replication is supported only for directory and MTree replication.
- 

**Note**

MTree replication is not supported for DD OS 5.0 and earlier.

---

Be aware that:

- ◆ Collection replication and MTree replication replicate the minimum and maximum retention periods configured on MTrees to the destination system.
- ◆ Directory replication does not replicate the minimum and maximum retention periods to the destination system.

The procedure for configuring and using collection, MTree, and directory replication is the same as for Data Domain systems that do not have a DD Retention Lock Governance license.

### Replication Resync

The `replication resync destination` command tries to bring the destination into sync with the source when the MTree or directory replication context is broken between destination and source systems. This command cannot be used with collection replication. Note that:

- ◆ If the destination MTree or directory contains retention-locked files that do not exist on the source, then resync will fail.
- ◆ If the destination directory has retention lock enabled, but the source directory does not have retention lock enabled, then a resync of a directory replication will fail.
- ◆ With MTree replication, resync will succeed if the source MTree does not have retention lock enabled while the destination MTree has retention lock enabled or vice versa, as long as the destination MTree does not contain retention-locked files not present on the source.

### Fastcopy

When `filesystem fastcopy source src destination dest` is run on a system with a DD Retention Lock Governance enabled MTree, it does not copy the locked or unlocked state of files. Files that are retention locked on the source are not retention locked on the destination.

If you try to fastcopy to a destination that has retention-locked files, the fastcopy operation aborts when it encounters retention-locked files on the destination.

### Filesys Destroy

When `filesystem destroy` is run on a system with a DD Retention Lock Governance enabled MTree:

- ◆ All data is destroyed, including retention-locked data.
- ◆ All `filesystem` options are returned to their defaults. This means that retention locking is disabled and the minimum and maximum retention periods are set back to their default values on the newly created file system.

---

#### Note

This command is not allowed if DD Retention Lock Compliance is enabled on the system.

---

### MTree Delete

When the `mtree delete mtree-path` command attempts to delete a DD Retention Lock Governance enabled (or previously enabled) MTree that currently contains data, the command returns an error.

---

#### Note

The behavior of `mtree delete` is a similar to a command to delete a directory—an MTree with retention lock enabled (or previously enabled) can be deleted only if the MTree is empty.

---

## DD Retention Lock Compliance

Certain DD OS commands behave differently when using DD Retention Lock Compliance. The following sections describe the differences for each.



## Replication

---

### Note

For information on using DD Replicator software, see [Working with DD Replicator on page 297](#).

---

In DD OS 5.4, an MTree enabled with Retention Lock Compliance can be replicated via MTree and collection replication only. Directory replication is not supported.

MTree and collection replication replicate the locked or unlocked state of files. Files that are compliance retention locked on the source are compliance retention locked on the destination and have the same level of protection. Minimum and maximum retention periods configured on MTrees are replicated to the destination system.

To perform collection replication, the same security officer user must be present on both the source and destination systems before starting replication to the destination system and afterward for the lifetime of the source/replica pair.

### Replication Resync

The `replication resync destination` command can be used with MTree replication, but not with collection replication.

- ◆ If the destination MTree contains retention-locked files that do not exist on the source, then resync will fail.
- ◆ Both source and destination MTrees must be enabled for Retention Lock Compliance, or resync will fail.

## Replication Procedures

Below are MTree and collection replication procedures supported for Retention Lock Compliance.

- ◆ MTree replication:
    - [Replicating an MTree: One-to-One Topology on page 177](#)
    - [Replicating an MTree: One-to-Many Topology on page 178](#)
    - [Adding Retention Lock Compliance Protection to an Existing MTree Replication Pair on page 179](#)
  - ◆ Upgrading to MTree replication from DD OS 5.2 collection replication:
    - [Converting a Collection Replication Pair to MTree Replication Pairs on page 180](#)
  - ◆ Collection replication:
    - [Performing Collection Replication on page 181](#)
    - [Adding Retention Lock Compliance Protection to an Existing Collection Replication Pair on page 182](#)
- 

### Note

For full descriptions of the commands used in this section, see the *EMC Data Domain Operating System Command Reference Guide*.

---

## Replicating an MTree: One-to-One Topology

This procedure describes how to replicate a Retention Lock Compliance enabled MTree from a source system to a destination system.

Instructions for creating a compliance-enabled MTree and retention locking files in the MTree are provided in:

- ◆ [Enabling DD Retention Lock on an MTree on page 167](#)
- ◆ [Client-Side Retention Lock File Control on page 170](#)

Complete these activities on the source system before starting.

### Procedure

1. On the destination system only:

- a. Complete steps 1–3 in [Enabling DD Retention Lock Compliance on an MTree on page 168](#).
- b. Create a replication context.

```
replication add source mtree://source-system-name/data/
coll/mtree-name destination mtree://destination-system-
name/data/coll/mtree-name
```

2. On the source system only:

- a. Create a replication context.

```
replication add source mtree://source-system-name/data/
coll/mtree-name destination mtree://destination-system-
name/data/coll/mtree-name
```

- b. Initialize the replication context.

```
replication initialize mtree://destination-system-name/
data/coll/mtree-name
```

- c. Confirm that replication is complete.

```
replication status mtree://destination-system-name/data/
coll/mtree-name detailed
```

This command reports 0 pre-compressed bytes remaining when replication is finished.

### Replicating an MTree: One-to-Many Topology

This procedure describes how to replicate a Retention Lock Compliance enabled MTree from a source system to multiple destination systems.

Instructions for creating a compliance-enabled MTree and retention locking files in the MTree are provided in:

- ◆ [Enabling DD Retention Lock Compliance on an MTree on page 168](#)
- ◆ [Client-Side Retention Lock File Control on page 170](#)

Complete these activities on the source system before starting the step below.

### Procedure

1. On the each destination system:

- a. Complete steps 1–3 in [Enabling DD Retention Lock Compliance on an MTree on page 168](#).
- b. Create a replication context.

```
replication add source mtree://source-system-name/data/
coll/mtree-name destination mtree://destination-system-
name/data/coll/mtree-name
```

## 2. On the source system only:

- a. Create a replication context for each destination system.

```
replication add source mtree://source-system-name/data/
coll/mtree-name destination mtree://destination-system-
name/data/coll/mtree-name
```

- b. Initialize the replication context for each destination system MTree.

```
replication initialize mtree://destination-system-name/
data/coll/mtree-name
```

- c. Confirm that replication is complete for each destination system.

```
replication status mtree://destination-system-name/data/
coll/mtree-name detailed
```

This command reports 0 pre-compressed bytes remaining when replication is finished.

## Adding Retention Lock Compliance Protection to an Existing MTree Replication Pair

Replication of MTrees enabled for Retention Lock Compliance was introduced with DD OS 5.3. This procedure describes how to add DD Retention Lock Compliance protection to an existing MTree replication pair that is not enabled for retention locking.

### Procedure

## 1. On both the source and destination systems:

- Complete steps 1–3 in [Enabling DD Retention Lock Governance on an MTree on page 167](#).
- Break the current MTree context on the replication pair.

```
replication break mtree://destination-system-name/data/
coll/mtree-name
```

- c. Create the new replication context.

```
replication add source mtree://source-system-name/data/
coll/mtree-name destination mtree://destination-system-
name/data/coll/mtree-name
```

## 2. On the source system only:

- Continue with steps 4-5 in [Enabling DD Retention Lock Governance on an MTree on page 167](#). (Security officer authorization is required.)
- Lock files in the compliance-enabled MTree using the instructions in [Client-Side Retention Lock File Control on page 170](#).

- c. Ensure that both source and destination (replica) MTrees are the same.

```
replication resync mtree://destination-system-name/data/
coll/mtree-name
```

- d. Check the progress of resync.

```
replication watch mtree://destination-system-name/data/
coll/mtree-name
```

- e. Confirm that replication is complete.

```
replication status mtree://destination-system-name/data/
coll/mtree-name detailed
```

This command reports 0 pre-compressed bytes remaining when replication is finished.

## Converting a Collection Replication Pair to MTree Replication Pairs

This procedure is for customers who used collection replication under DD Retention Lock Compliance in DD OS 5.2 and want to upgrade compliance-enabled MTrees in the collection replication pair to MTree replication pairs.

### Procedure

1. On the source system only:

- a. Create a snapshot for each Retention Lock Compliance enabled MTree.

```
snapshot create snapshot-name /data/coll/mtree-name
```

- b. Synchronize the collection replication pair.

```
replication sync col://destination-system-name
```

- c. Confirm that replication is complete.

```
replication status col://destination-system-name detailed
```

This command reports 0 pre-compressed bytes remaining when replication is finished.

- d. View snapshot information for each Retention Lock Compliance enabled MTree.

```
snapshot list mtree /data/coll/mtree-name
```

Note the snapshot names for use later.

2. On the destination system only:

- a. Confirm that the replication is complete.

```
replication status mtree://destination-system-name/data/coll/mtree-name detailed
```

This command reports 0 pre-compressed bytes remaining when replication is finished.

- b. View each MTree snapshot replicated to the destination system.

```
snapshot list mtree /data/coll/mtree-name
```

- c. Ensure that all Retention Lock Compliance MTree snapshots have been replicated by comparing the snapshot names generated here with those generated on the source system.

```
snapshot list mtree /data/coll/mtree-name
```

3. On the both the source and destinations systems:

- a. Disable the file system.

```
filesystem disable
```

- b. Break the collection replication context.

```
replication break col://destination-system-name
```

- c. Enable the file system. (Security officer authorization may be required.)

```
filesystem enable
```

- d. Add a replication context for each Retention Lock Compliance enabled MTree.

```
replication add source mtree://source-system-name/data/coll/mtree-name destination mtree://destination-system-name/data/coll/mtree-name
```

---

**Note**

Source and destination MTree names must be the same.

---

## 4. On the source system only:

- a. Ensure that both source and destination MTrees are the same.

```
replication resync mtree://destination-system-name
```

- b. Check the progress of resync.

```
replication watch destination
```

- c. Confirm that replication is complete.

```
replication status mtree://destination-system-name/data/col1/mtree-name detailed
```

This command reports 0 pre-compressed bytes remaining when replication is finished.

## Performing Collection Replication

This procedure describes how to replicate /data/col1 from a compliance-enabled source system to a compliance-enabled destination system.

### Procedure

## 1. On the source system only:

- a. Go to
- [Enabling DD Retention Lock Governance on an MTree on page 167](#)
- and complete steps 1–3.
- 
- b. Create the replication context.

```
replication add source col://source-system-name
destination col://destination-system-name
```

## 2. On the destination system only:

- a. Destroy the file system.

```
filesystem destroy
```

- b. Go to
- [Enabling DD Retention Lock Governance on an MTree on page 167](#)
- and complete steps 1–2.

---

**Note**

For collection replication the same security officer account must be used on both the source and destination systems.

---

- c. Create a file system, but do not enable it.

```
filesystem create
```

- d. Create the replication context.

```
replication add source col://source-system-name
destination col://destination-system-name
```

- e. Configure and enable the system to use DD Retention Lock Compliance.

```
system retention-lock compliance configure
```

(The system automatically reboots and executes the `system retention-lock compliance enable` command.)

## 3. On the source system only:

## a. Initialize the replication context.

```
replication initialize source col://source-system-name
destination col://destination-system-name
```

## b. Confirm that replication is complete.

```
replication status col://destination-system-name detailed
```

This command reports 0 pre-compressed bytes remaining when replication is finished.

## Adding Retention Lock Compliance Protection to an Existing Collection Replication Pair

This procedure describes how to add DD Retention Lock Compliance protection to a collection replication pair that was created without DD Retention Lock Compliance enabled on the source and destination systems.

### Procedure

## 1. On both source and destination systems:

## a. Disable the replication.

```
replication disable col://destination-system-name
```

b. Go to [Enabling DD Retention Lock Governance on an MTree on page 167](#) and complete steps 1–2.

## 2. On the source system first, and then on the destination system:

## a. Configure and enable the system to use DD Retention Lock Compliance.

```
system retention-lock compliance configure
```

(The system automatically reboots by executing the `system retention-lock compliance enable` command.)

## b. Enable the replication context.

```
replication enable col://destination-system-name
```

## Fastcopy

When the `fileSYS fastcopy source src destination dest` command is run on a system with a DD Retention Lock Compliance enabled MTree, it does not copy the locked or unlocked state of files. Files that are retention locked on the source are not retention locked on the destination.

If you try to fastcopy to a destination that has retention-locked files, the fastcopy operation aborts when it encounters retention-locked files on the destination.

## CLI Usage

A Data Domain system with DD Retention Lock Compliance has the following key considerations:

- ◆ Commands that break compliance cannot be run. The following commands are disallowed:

- `fileSYS archive unit del archive-unit`
- `fileSYS destroy`
- `mtree delete mtree-path`

- `mtree retention-lock reset {min-retention-period period | max-retention-period period} mtree mtree-path`
- `mtree retention-lock disable mtree mtree-path`
- `mtree retention-lock revert`
- `user reset`
- ◆ The following command requires security officer authorization if the license being deleted is for DD Retention Lock Compliance:
  - `license del license-feature [license-feature ...] | license-code [license-code ...]`
- ◆ The following commands require security officer authorization if DD Retention Lock Compliance is enabled on an MTree specified in the command:
  - `mtree retention-lock set {min-retention-period period | max-retention-period period} mtree mtree-path`
  - `mtree rename mtree-path new-mtree-path`
- ◆ The following commands require security officer authorization if DD Retention Lock Compliance is enabled on the system:
  - `alerts notify-list reset`
  - `config set timezone zonename`
  - `config reset timezone`
  - `cifs set authentication active-directory realm { [dc1 [dc2 ...]]`
  - `license reset`
  - `ntp add timeserver time server list`
  - `ntp del timeserver time server list`
  - `ntp disable`
  - `ntp enable`
  - `ntp reset`
  - `ntp reset timeservers`
  - `replication break {destination | all}`
  - `replication disable {destination | all}`
  - `system set date MMDDhhmm[[CC]YY]`

## System Clock

DD Retention Lock Compliance implements an internal security clock to prevent malicious tampering with the system clock. The security clock closely monitors and records the system clock. If there is an accumulated two-week skew within a year between the security clock and the system clock, the Data Domain file system (DDFS) is disabled and can be resumed only by a security officer.

### Finding the System Clock Skew

You can run the DD OS command `system retention-lock compliance status` (security officer authorization required) to get system and security clock information, including the last recorded security clock value, and the accumulated system clock variance. This value is updated every 10 minutes.

## Removing the System Clock Skew

Clock skew is updated every time the security clock records a new value for the system clock. After 1 year, it is reset to 0. At any time, you can run the DD OS command `system set date MMDDhhmm[[CC]YY]` to set the time of the system clock (security officer authorization required). If the clock skew becomes larger than the preset value (2 weeks), the file system is disabled. Complete these steps to restart the file system and remove the skew between security and system clocks.

### Procedure

1. At the system console, enable the file system.

```
fileSYS enable
```

2. At the prompt, confirm that you want to quit the `fileSYS enable` command and check whether the system date is right.

3. Display the system date.

```
system show date
```

4. If the system date is not correct, set the correct date (security officer authorization is required) and confirm it.

```
system set date MMDDhhmm[ [CC]YY]  
system show date
```

5. Enable the file system again.

```
fileSYS enable
```

6. At the prompt, continue to the enabling procedure.

7. A security officer prompt appears. Complete the security officer authorization to start the file system. The security clock will automatically be updated to the current system date.



# CHAPTER 8

## Working with MTrees

This chapter includes:

- ◆ [About MTrees](#)..... 186
- ◆ [Monitoring MTree Usage](#)..... 191
- ◆ [Managing MTree Operations](#)..... 192

## About MTrees

An MTree is a logical partition of the Data Domain file system for use in the following ways: DD Boost storage units, VTL pools, or an NFS/CIFS share. MTrees allow granular management of snapshots, quotas, and Retention Lock. For systems that have DD Extended Retention and granular management of data migration policies from Active Tier to Retention Tier, MTree operations can be performed on a specific MTree as opposed to the entire file system.

---

### Note

Although a Data Domain system supports a maximum of 100 MTrees, system performance might degrade rapidly if more MTrees are actively engaged in read or write streams.

- ◆ DD990, DD890, and DD880 platforms running DD OS 5.3, and later, support 32 currently accessible MTrees.
- ◆ DD4200, DD4500, and DD7200 platforms running DD OS 5.4, and later, support 32 currently accessible MTrees.
- ◆ All other platforms running DD OS 5.2 and DD OS 5.3, and later, support 14 currently accessible MTrees.

The degree of degradation depends on overall I/O intensity and other file system loads. For optimum performance, constrain the number of simultaneously active MTrees to a maximum of 14 or 32 as described above. Also, whenever possible, aggregate operations on the same MTree into a single operation.

There can be up to 99 MTrees designated for MTree replication contexts. For example, seven Data Domain systems, each with 14 active MTrees can replicate into a single Data Domain system.

---

## Quotas

MTree quotas apply only to the logical data written to the MTree. An administrator can set the storage space restriction for an MTree, Storage Unit, or VTL pool to prevent it from consuming excess space. There are two kinds of quota limits: hard limits and soft limits. You can set either a soft or hard limit or both a soft and hard limit. Both values must be integers, and the soft value must be less than the hard value.

When a soft limit is set, an alert is sent when the MTree size exceeds the limit, but data can still be written to it. When a hard limit is set, data cannot be written to the MTree when the hard limit is reached. Therefore, all write operations fail until data is deleted from the MTree.

See [Configure and Enable/Disable MTree Quotas on page 193](#).

## About the MTree Overview Panel

By default, the MTree overview lists all the active MTrees on the system and shows real-time data storage statistics. Information in the overview area is helpful in visualizing space usage trends. Click the **Data Management** > **MTree** tabs.

- ◆ Click a checkbox of an MTree in the list to display details and perform configuration in the Summary view.
- ◆ Enter text (wildcards are supported) in the Filter By MTree Name field and click **Update** to list specific MTree names in the list.

- ◆ Delete filter text and click **Reset** to return to the default list.

The MTree overview information includes:

Item	Description
MTree Name	The pathname of the MTree.
Quota Hard Limit	Percentage of hard limit quota used.
Quota Soft Limit	Percentage of hard limit quota used.
Last 24 hr Pre-Comp (pre-compression)	Amount of raw data from the backup application that has been written in the last 24 hours.
Last 24 hr Post-Comp (post-compression)	Amount of storage used after compression in the last 24 hours.
Last 24 hr Comp Ratio	The compression ratio for the last 24 hours.
Weekly Avg Post-Comp	Average amount of compressed storage used in the last five weeks.
Last Week Post-Comp	Average amount of compressed storage used in the last seven days.
Weekly Avg Comp Ratio	The average compression ratio for the last five weeks.
Last Week Comp Ratio	The average compression ratio for the last seven days.

## About the Summary View

Click the Summary tab to view important file system statistics as described in the following section.

### View Detail Information

Selecting an MTree in the overview list displays additional details in this area.

The detailed information for a selected MTree includes:

Item	Description
Full Path	The pathname of the MTree.
Pre-Comp Size	The current amount of raw data from the backup application that has been written to the MTree.
Status	<p>The current status of the MTree (combinations are supported). Status can be:</p> <ul style="list-style-type: none"> <li>• D: Deleted</li> <li>• RO: Read-only</li> <li>• RW: Read/write</li> <li>• RD: Replication destination</li> <li>• RLCE: Retention Lock Compliance enabled</li> <li>• RLCD: Retention Lock Compliance disabled</li> <li>• RLGE: Retention Lock Governance enabled</li> <li>• RLGD: Retention Lock Governance disabled</li> </ul>

Item	Description
Quota	
Quota Enforcement	Enabled or Disabled.
Pre-Comp Soft Limit	Current value. Click Configure to revise the quota limits.
Pre-Comp Hard Limit	Current value. Click Configure to revise the quota limits.
Quota Summary	Percentage of Hard Limit used.
Protocols	
CIFS Shared	<p>The CIFS share status. Status can be:</p> <ul style="list-style-type: none"> <li>• Yes—The MTree or its parent directory is shared.</li> <li>• Partial—The subdirectory under this MTree is shared.</li> <li>• No—This MTree and its parent or child directories are not shared.</li> </ul> <p>Click the CIFS link to go to the CIFS view.</p>
NFS Exported	<p>The NFS export status. Status can be:</p> <ul style="list-style-type: none"> <li>• Yes—The MTree or its parent directory is exported.</li> <li>• Partial—The subdirectory under this MTree is exported.</li> <li>• No—This MTree and its parent or child directories are not exported.</li> </ul> <p>Click the NFS link to go to the NFS view.</p>
DD Boost Storage Unit	<p>The DD Boost export status. Status can be:</p> <ul style="list-style-type: none"> <li>• Yes—The MTree is exported.</li> <li>• No—This MTree is not exported.</li> <li>• Unknown—There is no information.</li> </ul> <p>Click the DD Boost link to go to the DD Boost view.</p>
VTL Pool	If applicable, the name of the VTL pool that was converted to an MTree.

## View MTree Replication Information

If the selected MTree is configured for replication, summary information about the configuration displays in this area. Otherwise, this area displays `No Record Found`.

- ◆ Click the Replication link to go to the Replication page for configuration and to see additional details.

The Replication information includes:

Item	Description
Source	The source MTree pathname.
Destination	The destination MTree pathname.
Status	The status of the MTree replication pair. Status can be Normal, Error, or Warning.
Synced As Of Time	The last day and time the replication pair was synchronized.

## View MTree Snapshot Information

If the selected MTree is configured for snapshots, summary information about the snapshot configuration displays in this area.

- ◆ Click the **Snapshots** link to go to the Snapshots page to perform configuration or to see additional details.
- ◆ Click **Assign Snapshot Schedules** to assign a snapshot schedule to the selected MTree. Select the schedule's checkbox, and then click **OK** and **Close**. To create a snapshot schedule, click **Create Snapshot Schedule** (see [Create a Snapshot Schedule on page 202](#) for instructions).

The snapshot information includes:

Item	Description
Total Snapshots	The total number of snapshots created for this MTree. A total of 750 snapshots can be created for each MTree.
Expired	The number of snapshots in this MTree that have been marked for deletion, but have not been removed with the clean operation as yet.
Unexpired	The number of snapshots in this MTree that are marked for keeping.
Oldest Snapshot	The date of the oldest snapshot for this MTree.
Newest Snapshot	The date of the newest snapshot for this MTree.
Next Scheduled	The date of the next scheduled snapshot.
Assigned Snapshot Schedules	The name of the snapshot schedule assigned to this MTree.

## View MTree Retention Lock Information

If the selected MTree is configured for one of the Retention Lock software options, summary information about the Retention Lock configuration displays in this area.

### Note

For information on how to manage Retention Lock for an MTree, see [Working with DD Retention Lock on page 163](#).

The Retention Lock information includes:

Item	Description
Status	Indicates whether Retention Lock is enabled or disabled.

Item	Description
Retention Period	Indicates the minimum and maximum Retention Lock time periods.
UUID	Shows either: <ul style="list-style-type: none"> <li>the unique identification number generated for an MTree when the MTree is enabled for Retention Lock</li> <li>that the Retention Lock on a file in the MTree has been reverted</li> </ul>

## Enabling and Managing DD Retention Lock Settings

### Procedure

1. Go to the Data Management › MTree › Summary tab.
2. In the Retention Lock area, click **Edit**.
3. In the Modify Retention Lock dialog box, select **Enable** to enable Retention Lock on the Data Domain system.
4. Modify the minimum or maximum retention period (the feature must be enabled first), in the Retention Period pane.
5. Select an interval (minutes, hours, days, years). Click **Default** to show the default values.
6. Click **OK**.

### Results

After you close the Modify Retention Lock dialog box, updated MTree information is displayed in the Retention Lock summary area.

## About the Space Usage View

The Space Usage view contains a graph that displays a visual representation of data usage for the MTree. Click the **Data Management › MTree › Space Usage** tabs.

- ◆ Click a point on a graph line to display a box with data at that point.
- ◆ Click **Print** (at the bottom on the graph) to open the standard Print dialog box.
- ◆ Click **Show in new window** to display the graph in a new browser window.

The lines of the graph denote measurement for:

- ◆ **Pre-comp Written**—The total amount of data sent to the MTree by backup servers. Pre-compressed data on an MTree is what a backup server sees as the total uncompressed data held by an MTree-as-storage-unit, shown with the Space Used (left) vertical axis of the graph.
- ◆ **Post-comp Used**—The total amount of disk storage in use on the MTree, shown with the Space Used (left) vertical axis of the graph.
- ◆ **Comp Factor**—The amount of compression the Data Domain system has performed with the data it received (compression ratio), shown with the Compression Factor (right) vertical axis of the graph.

### Checking Historical Space Usage

On the Space Usage graph, clicking an interval (that is, 7d, 30d, 60d, or 120d) on the Duration line above the graph allows you to change the number of days of data shown on the graph, from 7 to 120 days.

To see space usage for intervals over 120 days, issue the following command:

```
# fileysys show compression [summary | daily | daily-detailed] {[last n
{hours | days | weeks | months}] | [start date [end date]]}
```

## About the Daily Written View

The Data Written pane contains a graph that displays a visual representation of data that is written daily to the MTree over a period of time, selectable from 7 to 120 days. The data amounts are shown over time for pre- and post-compression amounts.

It also provides totals for global and local compression amounts, and pre-compression and post-compression amounts.

- ◆ Click a point on a graph line to display a box with data at that point.
- ◆ Click **Print** (at the bottom on the graph) to open the standard Print dialog box.
- ◆ Click **Show** in new window to display the graph in a new browser window.

The lines on the graph denote measurements for:

- ◆ **Pre-Comp**—The total amount of data written to the MTree by backup servers. Pre-compressed data on an MTree is what a backup server sees as the total uncompressed data held by an MTree -as-storage-unit.
- ◆ **Post-Comp**—The total amount of data written to the MTree after compression has been performed, as shown in GiBs.
- ◆ **Total Comp**—The total amount of compression the Data Domain system has performed with the data it received (compression ratio). Shown with the Total Compression Factor (right) vertical axis of the graph.

### Checking Historical Written Data

On the Daily Written graph, clicking an interval (for example, 7d, 30d, 60d, or 120d) on the Duration line above the graph allows you to change the number of days of data shown on the graph, from 7 to 120 days.

Below the Daily Written graph, the following totals display for the current duration value:

- ◆ Pre-comp
- ◆ Post-comp
- ◆ Global-comp factor
- ◆ Local-comp factor
- ◆ Total-comp factor

## Monitoring MTree Usage

### Procedure

1. Select a system in the Navigation panel.
2. Click the **Data Management** > **MTree** tabs.

The MTree view shows a list of configured MTrees, and when selected in the list, details of the MTree are shown in the Summary tab. The Space Usage and Daily Written tabs show graphs that visually display space usage amounts and data written trends for a selected MTree. The view also contains options that allow MTree configuration for CIFS, NFS, and DD Boost, as well as sections for managing snapshots and Retention Lock for an MTree.

The MTree view has an MTree overview panel and three tabs which are described in detail in the following sections:

- [About the MTree Overview Panel on page 186](#)
- [About the Summary View on page 135](#)
- [About the Space Usage View on page 138](#)
- [About the Daily Written View on page 139](#)

## Managing MTree Operations

The following MTree operations are described in this section:

- ◆ [Create an MTree on page 192](#)
- ◆ [Configure and Enable/Disable MTree Quotas on page 193](#)
- ◆ [Delete an MTree on page 194](#)
- ◆ [Undelete an MTree on page 195](#)
- ◆ [Renaming an MTree on page 195](#)
- ◆ [Replicating a System with Quotas to One Without on page 195](#)

### Create an MTree

MTrees are created in the area `/data/col1/mtree_name`.

To create an MTree:

#### Procedure

1. Select a system in the Navigational pane.
2. Click the **Data Management** > **MTree** tabs.
3. In the MTree overview area, click **Create**.

The Create MTree dialog box is displayed.

4. Enter the name of the MTree in the MTree Name text box. MTree names can be up to 50 characters. The following characters are acceptable:
  - Upper- and lower-case alphabetical characters: A-Z, a-z
  - Numbers: 0-9
  - Embedded space
  - comma (,)
  - period (.), as long as it does not precede the name.
  - explanation mark (!)
  - number sign (#)
  - dollar sign (\$)
  - per cent sign (%)
  - plus sign (+)
  - at sign @)
  - equal sign (=)
  - ampersand (&)
  - semi-colon (;)
  - parenthesis [(and)]



- square brackets ([and])
  - curly brackets ({and})
  - caret (^)
  - tilde (~)
  - apostrophe (unslanted single quotation mark)
  - single slanted quotation mark (')
5. Setting storage space restrictions for the MTree to prevent an MTree from consuming excess space, enter either a soft or hard limit quota setting, or both a hard and soft limit. With a soft limit an alert is sent when the MTree size exceeds the limit, but data can still be written to the MTree. Data cannot be written to the MTree when the hard limit is reached.

---

#### Note

The quota limits are pre-compressed values.  
To set quota limits for the MTree, select **Set to Specific value** and enter the value.  
Select the unit of measurement: MiB, GiB, TiB, or PiB.

---



---

#### Note

When setting both soft and hard limits, a quota's soft limit cannot exceed the quota's hard limit.

---

6. Click **OK**.  
The new MTree displays in the MTree table.
- 

#### Note

You may need to expand the width of the MTree Name column to see the entire pathname.

---

## Configure and Enable/Disable MTree Quotas

An administrator can set the storage space restriction for an MTree, Storage Unit, or VTL pool to prevent it from consuming excess space. The Data Management > Quota page shows the administrator how many MTree have no soft or hard quotas set, and for MTrees with quotas set, the percentage of pre-compressed soft and hard limits used.

---

#### Note

MTree quotas apply to ingest operations. These quotas can be applied to data on systems that have the DD Extended Retention software, regardless of which tier it resides on; as well as VTL, DD Boost, CIFS, and NFS.

Snapshots are not counted.

Quota cannot be set on the `/data/coll/backup` directory.

The maximum quota value allowed is 4096 PiB.

---

### Configure MTree Quotas

You can use the MTree tab or the Quota tab to configure MTree quotas.

#### Configuring MTree Quotas (MTree Tab)

To configure MTree quotas using the MTree tabs:

**Procedure**

1. Select a system in the Navigation panel.
2. Click the **Data Management** > **MTree** tab.
3. Click the **Summary** tab. In the Quota area, click the **Configure** button.
4. In the Configure Quota for MTrees dialog box, enter values for hard and soft quotas and select the unit of measurement: MiB, GiB, TiB, or PiB.
5. Click **OK**.

**Configuring MTree Quotas (Quota Tab)****Procedure**

1. Select a system in the Navigation panel.
2. Click the **Data Management** > **Quota** tabs.
3. Select an MTree.

**Note**


---

Quotas cannot be set on the `/data/coll/backup` directory.

---

4. Click the **Configure Quota** button.
5. In the Configure Quota for MTrees dialog box, enter values for hard and soft quotas and select the unit of measurement: MiB, GiB, TiB, or PiB.
6. Click **OK**.

**Enable/Disable MTree Quotas**

To enable or disable MTree quotas:

**Procedure**

1. Select a system in the Navigational pane.
2. Click the **Data Management** > **Quota** tab.
3. In Quota Settings, Quota Enforcement, click the **Disable** button if quota enforcement is Enabled. Click **Enable** if it is disabled.

**Delete an MTree**

Deleting an MTree removes the MTree from the MTree table and removes all data in that MTree at the next file system cleaning.

---

**Note**

Because the MTree and its associated data are not removed until file cleaning is run, you cannot create a new MTree with the same name as a deleted MTree until the deleted MTree is completely removed from the file system via cleaning.

---

**Procedure**

1. Select a system in the Navigational pane.
2. Click the **Data Management** > **MTree** tabs.
3. In the MTree overview area, click **Delete**.
4. Click **OK** at the Warning dialog box.
5. Click **Close** in the Delete MTree Status dialog box after viewing the progress.

## Undelete an MTree

An undelete of an MTree retrieves a deleted MTree and its data and places it back in the MTree table. This undelete is possible only if file cleaning has not been run after the MTree was marked for deletion.

### Procedure

1. Select a system in the Navigational pane.
2. Click the **Data Management** › **MTree** tabs.
3. From the More Tasks menu, select **Undelete**.  
The Undelete MTree dialog box is displayed.
4. Click the checkboxes of the MTrees you wish to bring back and click **OK**.
5. Click **Close** in the Undelete MTree Status dialog box after viewing the progress.  
The recovered MTree displays in the MTree table.

## Renaming an MTree

### Procedure

1. Select a system in the Navigation panel.
2. Click the **Data Management** › **MTree** tabs.
3. Select an MTree in the MTree table.
4. Select the Summary tab.
5. In the Detailed Information overview area, click **Rename**.  
The Rename MTree dialog box displays.
6. Enter the name of the MTree in the New MTree Name text box.  
See [Create an MTree on page 192](#) for a list of allowed characters.
7. Click **OK**.  
The renamed MTree displays in the MTree table.

## Replicating a System with Quotas to One Without

### Note

Quotas were introduced as of DD OS 5.2.

When replicating a Data Domain system with a DD OS that supports quotas to a system with a DD OS that does not have quotas, do one of the following:

- ◆ A reverse resync, which takes the data from the system without quotas and puts it back in an MTree on the system that has quotas enabled (and which continues to have quotas enabled).
- ◆ A reverse initialization from the system without quotas, which takes its data and creates a new MTree on the system that supports quotas but does not have quotas enabled because it was created from data on a system without quotas.



# CHAPTER 9

## Working with Snapshots

This chapter includes:

- ◆ [About Snapshots](#)..... 198
- ◆ [Monitoring Snapshots and Their Schedules](#).....199
- ◆ [Managing Snapshots](#)..... 200
- ◆ [Managing Snapshot Schedules](#)..... 202
- ◆ [Recover Data from a Snapshot](#).....204

## About Snapshots

This chapter describes how to use the snapshot feature with MTree.

A snapshot saves a read-only copy (called a *snapshot*) of a designated MTree at a specific time. You can use a snapshot as a restore point. You can manage MTree snapshots and schedules and display information about the status of existing snapshots. For more information about MTree, see [About MTree on page 186](#).

---

### Note

Snapshots created on the source Data Domain system are replicated to the destination with collection and MTree replication. It is not possible to create snapshots on a Data Domain system that is a replica for collection replication. It is also not possible to create a snapshot on the destination MTree of MTree replication. Directory replication does not replicate the snapshots, and it requires you to create snapshots separately on the destination system.

---

Snapshots for the MTree named `backup` are created in the system directory `/data/coll/backup/.snapshot`. Each directory under `/data/coll/backup` also has a `.snapshot` directory with the name of each snapshot that includes the directory. Each MTree has the same type of structure, so an MTree named `SantaClara` would have a system directory `/data/coll/SantaClara/.snapshot`, and each subdirectory in `/data/coll/SantaClara` would have a `.snapshot` directory as well.

---

### Note

The `.snapshot` directory is not visible if only `/data` is mounted. When the MTree itself is mounted, the `.snapshot` directory is visible.

---

An expired snapshot remains available until the next file system cleaning operation.

The maximum number of snapshots allowed per MTree is 750. Warnings are sent when the number of snapshots per MTree reaches 90% of the maximum allowed number (from 675 to 749 snapshots), and an alert is generated when the maximum number is reached. To clear the warning, expire snapshots and then run the file system cleaning operation.

---

### Note

To identify an MTree that is nearing the maximum number of snapshots, check the Snapshots panel of the MTree page (see [View MTree Snapshot Information on page 189](#)).

---

Snapshot retention for an MTree does not take any extra space, but if a snapshot exists and the original file is no longer there, the space cannot be reclaimed.

---

### Note

Snapshots and CIFS Protocol: As of DD OS 5.0, the `.snapshot` directory is no longer visible in the directory listing in Windows Explorer or DOS CMD shell. You can access the `.snapshot` directory by entering its name in the Windows Explorer address bar or the DOS CMD shell. For example, `\\dd\backup\.snapshot` or `Z:\.snapshot` when `Z:` is mapped as `\\dd\backup`).

---

## Monitoring Snapshots and Their Schedules

The Snapshots view provides detailed and summary information about the status of snapshots and snapshot schedules.

### About the Snapshots View

The Snapshot view contains the following components:

- ◆ [Snapshots Overview Pane on page 199](#)
- ◆ [Snapshots View on page 199](#)
- ◆ [Schedules View on page 199](#)

### Snapshots Overview Pane

The Snapshots overview pane displays following snapshot information.

Field	Description
Total Snapshots (Across all MTrees)	The total number of snapshots, active and expired, on all MTrees in the system.
Expired	The number of snapshots that have been marked for deletion, but have not been removed with the clean operation as yet.
Unexpired	The number of snapshots that are marked for keeping.
Next file system clean scheduled	The date the next scheduled file system clean operation will be performed.

### Snapshots View

The Snapshots tab displays a list of snapshots and lists the following information.

Field	Description
Selected Mtree	A drop-down list that selects the MTree the snapshot operates on.
Filter By	Items to search for in the list of snapshots that display. Options are: <ul style="list-style-type: none"> <li>• <b>Name</b>—Name of the snapshot (wildcards are accepted).</li> <li>• <b>Year</b>—Drop-down list to select the year.</li> <li>• <b>Status</b>—Drop-down list to select the status (Expired or Unexpired).</li> </ul>
Name	The name of the snapshot image.
Creation Time	The date the snapshot was created.
Expires On	The date the snapshot expires.
Status	The status of the snapshot, which can be Expired or blank if the snapshot is active.

### Schedules View

The Schedules tab displays a list of snapshot schedules and lists the following information.

Field	Description
Name	The name of the snapshot schedule
Days	The days the snapshots will be taken.
Times	The time of day the snapshots will be taken.
Retention Period	The amount of time the snapshot will be retained.
Snapshot Name Pattern	A string of characters and variables that translate into a snapshot name (for example, <code>scheduled-%Y-%m-%d-%H-%M</code> , which translates to “scheduled-2010-04-12-17-33”).

## Managing Snapshots

Managing snapshots includes the following topics:

- ◆ [Create a Snapshot on page 200](#)
- ◆ [Modify a Snapshot Expiration Date on page 201](#)
- ◆ [Rename a Snapshot on page 201](#)
- ◆ [Expiring a Snapshot on page 201](#)

### Create a Snapshot

A snapshot can be created manually, when an unscheduled snapshot is required.

To manually create a snapshot, use the following procedure.

#### Procedure

1. Click the **Data Management** > **Snapshots** tabs to open the Snapshots view.  
The Snapshots tab is active by default.
2. In the Snapshots view, click **Create**.  
The Create dialog box is displayed.
3. In the Name text field, enter the name of the snapshot.
4. In the MTree(s) area, select a checkbox of one or more MTrees in the Available MTrees pane and click **Add**.  
The MTrees move to the Selected MTrees pane.
5. In the Expiration area, select one of these expiration options:
  - a. **Never Expire**
  - b. Enter a number for the In text field, and select **Days**, **Weeks**, **Month**, or **Years** from the drop-down list. The snapshot will be retained until the same time of day as when it is created.
  - c. Enter a date (using the format *mm/dd/yyyy*) in the On text field, or click **Calendar** and click a date. The snapshot will be retained until midnight (00:00, the first minute of the day) of the given date.
6. Click **OK** and **Close**.  
The snapshot is added to the list.



## Modify a Snapshot Expiration Date

The administrator may wish to modify snapshot expiration dates, for example, when snapshots take up too much disk space, if snapshots were created too frequently, or if a date needs to be extended for auditing or compliance.

To modify a scheduled snapshot expiration date:

### Procedure

1. Click the Data Management › Snapshots › tabs to open the Snapshots view.  
The Snapshots tab is active by default.
2. Click the checkbox of the snapshot entry in the list and click **Modify Expiration Date**.  
The Modify dialog box is displayed.

---

### Note

More than one snapshot can be selected by clicking additional checkboxes.

3. In the Expiration area, select one of the following for the expiration date:
  - a. **Never Expire**
  - b. In the In text field, enter a number and select **Days, Weeks, Month, or Years** from the drop-down list. The snapshot will be retained until the same time of day as when it is created.
  - c. In the On text field, enter a date (using the format *mm/dd/yyyy*) or click **Calendar** and click a date. The snapshot will be retained until midnight (00:00, the first minute of the day) of the given date.
4. Click **OK**.

## Rename a Snapshot

To rename a snapshot:

### Procedure

1. Click the Data Management › Snapshots tabs to open the Snapshots view.  
The Snapshots tab is active by default.
2. Click the checkbox of the snapshot entry in the list and click **Rename**.  
The Rename dialog box is displayed.
3. In the Name text field, enter a new name.
4. Click **OK**.

## Expiring a Snapshot

Snapshots cannot be deleted. To free up disk space, you can expire snapshots manually and they will be deleted in the next cleaning operation after the expiry date. This operation can be used to remove snapshots that are no longer needed, but their scheduled expiration date has not occurred, or that have no expiration date.

To expire a scheduled snapshot:

### Procedure

1. Click the **Data Management › Snapshots** tabs to open the Snapshots view.  
The Snapshots tab is active by default.

2. Click the checkbox next to snapshot entry in the list and click **Expire**.

---

**Note**

More than one snapshot can be selected by selecting additional checkboxes. The snapshot is marked as Expired in the Status column and will be deleted at the next cleaning operation.

---

## Managing Snapshot Schedules

This section describes how to set up and manage a series of snapshots that are automatically taken at regular intervals in the future. Such a series of snapshots is called a “snapshot schedule,” or “schedule” for short.

Multiple snapshot schedules can be active at the same time.

---

**Note**

If multiple snapshots are scheduled to occur at the same time, only one is retained. Which one is retained is indeterminate, thus only one snapshot should be scheduled for a given time.

---

### Create a Snapshot Schedule

To add a snapshot schedule:

**Procedure**

1. Click the **Data Management > Snapshots > Schedules** tabs to open the Schedules view.
2. Click **Create**.  
The Create dialog appears.
3. In the **Name** text field, enter the name of the schedule.
4. In the **Snapshot Name Pattern** text box, enter a name pattern.  
Use alphabetic characters, numbers, `_`, `-`, and variables, such as `%Y-%m-%d-%H-%M` that translate into current values.
5. Click **Validate Pattern & Update Sample**. The name displays in the Live Sample field.
6. Click **Next**.
7. Select the date when the schedule is to be executed:
  - a. **Weekly**—Click checkboxes next to the days of the week or select **Every Day**.
  - b. **Monthly**—Click the **Selected Days** option and click the dates on the calendar, or select the **Last Day of the Month** option.
  - c. Click **Next**.
8. Select the time of day when the schedule is to be executed:
  - a. **At Specific Times**—Click **Add** and in the Time dialog that appears, enter the time in the format `hh:mm`, and click **OK**.
  - b. **In Intervals**—Click the drop-down arrows to select the start and end time `hh:mm` and AM or PM. Click the **Interval** drop-down arrows to select a number and then the hours or minutes of the interval.
  - c. Click **Next**.

9. In the Retention Period text entry field, enter a number and click the drop-down arrow to select days, months, or years, and click **Next**.

Schedules must explicitly specify a retention time.

10. Review the parameters in the schedule summary and click **Finish** to complete the schedule or **Back** to change any entries.
11. If an MTree is not associated with the schedule, a warning dialog box asks if you would like to add an MTree to the schedule. Click **OK** to continue (or **Cancel** to exit).
12. To assign an MTree to the schedule, in the MTree area, click the checkbox of one or more MTrees in the Available MTrees pane, click **Add** and **OK**.

The MTrees move to the Selected MTrees pane.

## Naming Conventions for Snapshots Created by a Schedule

The naming convention for scheduled snapshots is the word `scheduled` followed by the date when the snapshot is to occur, in the format `scheduled-yyyy-mm-dd-hh-mm`. For example, `scheduled-2009-04-27-13-30`.

The name “`mon_thurs`” is the name of a snapshot schedule. Snapshots generated by that schedule might have the names `scheduled-2008-03-24-20-00`, `scheduled-2008-03-25-20-00`, etc.

## Modify a Snapshot Schedule

To modify a snapshot schedule:

### Procedure

1. In the schedule list, select the schedule and click **Modify**.  
The Modify Schedule dialog appears.
2. In the Name text field, enter the name of the schedule and click **Next**.  
Use alphanumeric characters, and the `_` and `-`.
3. Select the date when the schedule is to be executed:
  - a. **Weekly**—Click checkboxes next to the days of the week or select **Every Day**.
  - b. **Monthly**—Click the **Selected Days** option and click the dates on the calendar, or select the **Last Day of the Month** option.
  - c. Click **Next**.
4. Select the time of day when the schedule is to be executed:
  - a. **At Specific Times**—Click the checkbox of the scheduled time in the Times list and click **Edit**. In the Times dialog that appears, enter a new time in the format `hh:mm`, and click **OK**. Or click **Delete** to remove the scheduled time.
  - b. **In Intervals**—Click the drop-down arrows to select the start and end time `hh:mm` and AM or PM. Click the Interval drop-down arrows to select a number and then the hours or minutes of the interval.
  - c. Click **Next**.
5. In the Retention Period text entry field, enter a number and click the drop-down arrow to select days, months, or years, and click **Next**.
6. Review the parameters in the schedule summary and click **Finish** to complete the schedule or **Back** to change any entries.

## Delete a Snapshot Schedule

To delete a snapshot schedule,

### Procedure

1. In the schedule list, click the checkbox to select the schedule and click **Delete**.
2. In the verification dialog box, click **OK** and then **Close**.

## Recover Data from a Snapshot

The fastcopy operation can be used to retrieve data stored in a snapshot. See [Fast Copy Operations on page 146](#) for details.

# CHAPTER 10

## Working with CIFS

This chapter includes:

◆ CIFS Overview .....	206
◆ Performing CIFS Setup .....	206
◆ Working with Shares .....	210
◆ Managing Access Control .....	215
◆ Monitoring CIFS Operation .....	220
◆ Performing CIFS Troubleshooting .....	222

## CIFS Overview

The Common Internet File System (CIFS) clients can have access to the system directories on the Data Domain system. The `/data/coll/backup` directory is the destination directory for compressed backup server data. The `/ddvar` directory contains Data Domain system core and log files.

Clients, such as backup servers that perform backup and restore operations with a Data Domain System, at the least, need access to the `/data/coll/backup` directory. Clients that have administrative access need to be able to access the `/ddvar` directory to retrieve core and log files.

As part of the initial Data Domain system configuration, CIFS clients were configured to access these directories. This chapter describes how to modify these settings and how to manage data access using the Data Domain System Manager and `cifs` command.

---

### Note

- ◆ The Data Domain System Manager Data Management > CIFS page allows you to perform major CIFS operations (such as enabling and disabling CIFS, setting authentication, managing shares, and viewing configuration and share information).
  - ◆ The `cifs` command contains all the options to manage CIFS backup and restores between Windows clients and Data Domain systems, and display CIFS statistics and status. For complete information about the `cifs` command, see the *EMC DD OS Command Reference Guide*.
  - ◆ For information about the initial system configuration, see the *EMC DD OS Initial Configuration Guide*.
  - ◆ For information about setting up clients to use the Data Domain system as a server, see the related tuning guide, such as the *CIFS Tuning Guide*, which is available from the Data Domain support web site. From the Documentation > Integration Documentation page, select the vendor from the list and click **OK**. Select the tuning guide from the list.
- 

## Performing CIFS Setup

CIFS Setup topics.

### Prepare Clients for Access to Data Domain Systems

#### Procedure

1. Log into the Data Domain Support web site.
2. In the Systems pane, click **Documentation**.
3. On the Documentation page, click **Integration Documentation**.
4. Select the vendor for the client system's operating system, such as Microsoft, and click **OK**.
5. Select the appropriate tuning document, such as the *CIFS Tuning Guide*.
6. Follow the instructions given in the tuning document.

## Enabling CIFS Services

After configuring client access, as described in [Prepare Clients for Access to Data Domain Systems on page 206](#), enable CIFS services, which allows the client to access the system using the CIFS protocol.

### Procedure

1. For the Data Domain system that is selected in the DD System Manager Navigation tree, click **Data Management** > **CIFS**.
2. In the CIFS Status area, click **Enable**.

## Naming the CIFS Server

The hostname for the Data Domain system that serves as the CIFS server is set during the system's initial configuration. To change a CIFS server name, see the procedures in [Setting Authentication Parameters on page 207](#).

A Data Domain system's hostname should match the name assigned to its IP address, or addresses, in the DNS table. Otherwise authentication, as well as attempts to join a domain, can fail. If you need to change the Data Domain system's hostname, use the `net set hostname` command, and also modify the system's entry in the DNS table.

When the Data Domain system acts as a CIFS server, it takes the hostname of the system. For compatibility purposes, it also creates a NetBIOS name. The NetBIOS name is the first component of the hostname in all uppercase letters. For example, the hostname `jp9.oasis.local` is truncated to the NetBIOS name `JP9`. The CIFS server responds to both names.

You can have the CIFS server respond to different names at the NetBIOS levels by changing the NetBIOS hostname.

## Changing the NetBIOS Hostname

### Procedure

1. Display the current NetBIOS name by entering:

```
# cifs show config
```

2. Use the `cifs set nb-hostname nb-hostname` command.

## Setting Authentication Parameters

The DD System Manager Configure Authentication dialog box allows you to set the authentication parameters that the Data Domain system uses for working with CIFS.

The Data Domain system can join the active directory (AD) domain or the NT4 domain, or be part of a workgroup (the default). If you did not use the Data Domain System Manager's Configuration Wizard to set the join mode, use the procedures in this section to choose or change a mode.

The authentication configuration procedures are:

- ◆ [Configuring Authentication for Active Directory on page 208](#)
- ◆ [Configure Authentication for Workgroups on page 209](#)
- ◆ [Resetting the Authentication Mode to the Default \(Workgroup\) on page 209](#)

## Configuring Authentication for Active Directory

The Data Domain system must meet all active-directory requirements, such as a clock time that differs no more than five minutes from that of the domain controller. (See [Managing Access Control on page 215](#) for information about synchronizing clock time with a domain controller.)

---

### Note

When a Data Domain system that is already joined to an AD domain is joined to a different domain, the DD OS does not remove the system's account in the previously joined domain. In some AD configurations, these two accounts might cause authentication failure for domain users. To prevent this problem, manually delete the Data Domain system's account in the previously joined domain.

---

To set Active Directory authentication parameters:

### Procedure

1. On the CIFS page, click **Configure Authentication**.
2. Select **Configure Authentication**.  
The Configure Authentication dialog appears.
3. From the Mode drop-down list, select **Active Directory**.  
The active-directory mode joins a Data Domain System to an active-directory domain.
4. In the **Realm Name** text box, enter the full realm name for the system, such as **domain1.local**.
5. In the Domain Joining Credential area, enter a user name and password.  
Enter either a user on your company's domain, or a user in a domain that is a trusted domain of your company. The user name and password must be compatible with Microsoft requirements for the Active Directory domain being joined. This user must have permission to create accounts in the domain.
6. Click the **Advanced** tab to set additional information.
7. Optionally, to modify a CIFS server name, in the CIFS Server Name area, change the name of the CIFS server (for information about the CIFS server name, see [Naming the CIFS Server on page 207](#)):
  - Click the checkbox to use the default CIFS server name.
  - Deselect the checkbox and enter the CIFS server name in the text box.
8. In the Domain Controller area, determine how domain controllers are assigned:
  - For automatic assignment, select **Automatically Assign Domain Controllers**, which is the default and recommended.
  - To add specific domain controllers, select **Manually Assign Domain Controllers** and enter a controller name in the text box. You can add up to three controller names.  
You can enter fully qualified domain names, hostnames, or IP addresses.
9. Optionally, to join a specific Organizational Unit in the active directory, in the Organizational Unit area, set the name of the Organizational Unit:
  - Click the checkbox to use the default Organizational Unit.
  - Deselect the checkbox and enter the Organizational Unit name in the text box.



10. Optionally, to use DDNS, click the **Enable** checkbox.

11. Click **OK**.

## Configure Authentication for Workgroups

The workgroup mode means that the Data Domain system authenticates CIFS clients using local user accounts that are defined on the Data Domain system.

To set Workgroup authentication parameters:

### Procedure

1. On the CIFS page, click **Configure Authentication**.

The Configure Authentication dialog appears.

2. From the Mode drop-down list, select **Workgroup**.

3. Enter the name of the workgroup in the Workgroup Name text box.

4. Click the Advanced tab to configure additional settings.

5. Optionally, to modify a CIFS server name, in the CIFS Server Name area, change the name of the CIFS server:

- Click the checkbox to use the default CIFS server name.
- Deselect the checkbox and enter the CIFS server name in the text box.

## Resetting the Authentication Mode to the Default (Workgroup)

### Procedure

1. On the CIFS page, click **Configure Authentication**.

The Configure Authentication dialog appears.

2. From the Mode drop-down list, select **Workgroup** (default).

3. Click **OK**.

## Specify a WINS Server

From the CLI, the WINS server can be set when the Data Domain system needs to join a NT4 domain. This option does not need to be set for active directory domain or workgroup authentication.

### Procedure

◆ Use this command:

```
cifs set wins-server ipaddr
```

.

### After you finish

---

#### Note

If CIFS clients are using NetBIOS, a WINS server may be needed to resolve NetBIOS names to IP addresses.

---

## Restrict CIFS Interfaces

By default, the CIFS server listens on all Data Domain system NIC-active interfaces. From the CLI:

**Procedure**

- ◆ Use this command:  

```
cifs option set interfaces value
```

.

**Results**

The *value* is a list of interfaces, such as Ethernet port names. Multiple interfaces must be separated by a space and enclosed within double quotation marks; for example, "**eth0 eth2**".

## Set CIFS Options

**Procedure**

1. Select the Data Domain system in the Navigational tree and click the **Data Management** › **CIFS** › **Configuration** tabs.
2. In the Options area, click **Configure Options**.  
 The Configure Options dialog box is displayed.
3. To restrict anonymous connections, click the checkbox of the **Enable** option in the Restrict Anonymous Connections area.
4. In the LogLevel area, click the drop-down list to select the level number.

The level is an integer from 0 (zero) to 10 (ten). One is the default system level that sends the least-detailed level of CIFS-related log messages; ten results in the most detail. Log messages are located in the files `/ddvar/log/debug/cifs/clients.log` and `/ddvar/log/debug/cifs/cifs.log`.

**Note**

A log level of 10 degrades system performance. Click the **Default** in the Log Level area after debugging an issue. This sets the level back to 1.

## Disable CIFS Services

To prevent clients from accessing the Data Domain system:

**Procedure**

1. Select the Data Domain system in the Navigational tree and click the **Data Management** › **CIFS** tabs.
2. In the Status area, click **Disable**.  
 The Disable CIFS dialog box is displayed.
3. Click **OK**.

Even after disabling CIFS access, CIFS authentication services continue to run on the Data Domain system. This continuation is required to authenticate active directory domain users for management access.

## Working with Shares

To share data, create shares on the Data Domain system. Shares are administered on the Data Domain system and the CIFS systems.

## Creating Shares on the Data Domain System

When creating shares, you have to assign client access to each directory separately and remove access from each directory separately. For example, a client can be removed from `/ddvar` and still have access to `/data/col1/backup`.

---

### Note

If Replication is to be implemented, a Data Domain system can receive backups from both CIFS clients and NFS clients as long as separate directories are used for each. Do not mix CIFS and NFS data in the same directory.

---

### Procedure

1. From the Navigation panel, select a Data Domain system to configure shares.
2. Click **Data Management** > **CIFS** tabs to navigate to the CIFS view.
3. Ensure authentication has been configured, as described in [Setting Authentication Parameters on page 207](#).
4. On the CIFS client, set shared directory permissions or security options.
5. On the CIFS view, click the Shares tab.
6. Click **Create**.

The Create Shares dialog box is displayed.

7. In the Create Shares dialog box, enter the following information:

Item	Description
Share Name	A descriptive name for the share.
Directory Path	The path to the target directory (for example, <code>/data/col1/backup/dir1</code> ).
	<hr/> <b>Note</b> <hr/> <code>col1</code> uses the lower case letter L followed by the number 1.
Comment	A descriptive comment about the share.

8. Add a client by clicking Add (+) in the Clients area. The Client dialog box appears. Enter the name of the client in the Client text box and click **OK**.

No blanks or tabs (white space) characters are allowed.

Repeat this step for each client that you need to configure.

---

**Note**

- It is not recommended to use both an asterisk (\*) and individual client name or IP address for a given share. When an asterisk (\*) is present, any other client entries for that share are not used.
  - It is not required to use both client name and client IP address for the same client on a given share. Use client names when the client names are defined in the DNS table.
  - To make share available to all clients, specify an asterisk (\*) as the client. All users in the client list can access the share, unless one or more user names are specified, in which case only the listed names can access the share.
- 

9. In the Max Connections area, select the text box and enter the maximum number of connections to the share that are allowed at one time. The default value of zero (also settable via the Unlimited button) enforces no limit on the number of connections.

10. Click **OK**.

The newly created share appears at the end of the list of shares, located in the center of the Shares panel.

## Modify a Share on a Data Domain System

To modify the setup of an existing share:

**Procedure**

1. In the Shares tab, click the checkbox next the share you wish to modify in the Share Name list.

2. Click **Modify**.

The Modify Share dialog box is displayed.

3. Modify share information:

a. To change the comment, enter new text in the Comment text field.

b. To modify a User or Group names, in the User/Group list, click the checkbox of the user or group and click **Edit** (pencil icon) or **Delete** (X). To add a user or group, click (+), and in the User/Group dialog box select the Type for User or Group, and enter the user or group name.

Group names must be preceded by the *at* (@) symbol. For example, @group1.

c. To modify a client name, in the Client list click the checkbox of the client and click **Edit** (pencil icon) or **Delete** (X). To add a client, click the Add (+) and add the name in the Client dialog box.

---

**Note**

To make share available to all clients specify an asterisk (\*) as the client. All users in the client list can access the share, unless one or more user names are specified, in which case only the listed names can access the share.

---

d. Click **OK**.

4. In the Max Connections area, in the text box, change the maximum number of connections to the share that are allowed at one time. Or select Unlimited to enforce no limit on the number of connections.

5. Click **OK**.

## Creating a Share from an Existing Share

To create a share from an existing share, use the following procedure and modify the new share if necessary.

---

### Note

User permissions from the existing share are carried over to the new share.

---

### Procedure

1. In the CIFS Shares table, click the checkbox of the share you wish to use as the source.
2. Click **Create From**.  
The Create From Existing Share dialog box appears.
3. Modify the share information, as described in [Modify a Share on a Data Domain System on page 212](#).

## Disable a Share on a Data Domain System

To disable one or more existing shares:

### Procedure

1. In the Shares tab, click the checkbox of the share you wish to disable in the Share Name list.
2. Click **Disable**.  
The Disable Shares Status dialog box is displayed.
3. Click **Close**.

## Enable a Share on a Data Domain System

To enable one or more existing share:

### Procedure

1. In the Shares tab, click the checkbox of the shares you wish to enable in the Share Name list.
2. Click **Enable**.  
The Enable Shares Status dialog box is displayed.
3. Click **Close**.

## Delete a Share on a Data Domain System

To delete one or more existing shares:

### Procedure

1. In the Shares tab, click the checkbox of the shares you wish to delete in the Share Name list.
2. Click **Delete**.  
The Warning dialog box is displayed.

3. Click **OK**.

The shares are removed.

## Performing MMC Administration

You can use the Microsoft Management Console (MMC) for administration. DD OS supports these MMC features:

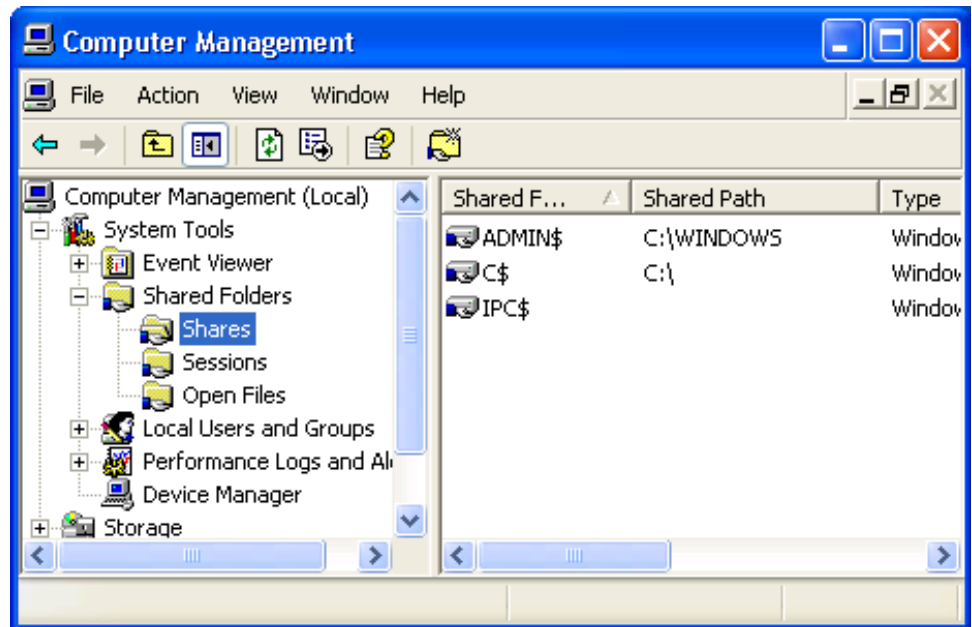
- ◆ Share management, except for browsing when adding a share, or the changing of the offline settings default, which is a manual procedure.
- ◆ Session management.
- ◆ Open file management, except for deleting files.

## Connecting to a Data Domain System from a CIFS Client

### Procedure

1. On the Data Domain system CIFS page, verify that CIFS Status shows that CIFS is enabled and running.
2. In the Control Panel, open Administrative Tools and select **Computer Management**.
3. In the Computer Management dialog box, right-click **Computer Management (Local)** and select **Connect to another computer** from the menu.
4. In the Select Computer dialog box, select **Another computer** and enter the name or IP address for the Data Domain system.
5. Create a `\backup` subfolder as read-only. See [Create a \data\col1\backup Subfolder as Read-Only on page 214](#).

**Figure 7** Computer Management Dialog Box



## Create a \data\col1\backup Subfolder as Read-Only

### Procedure

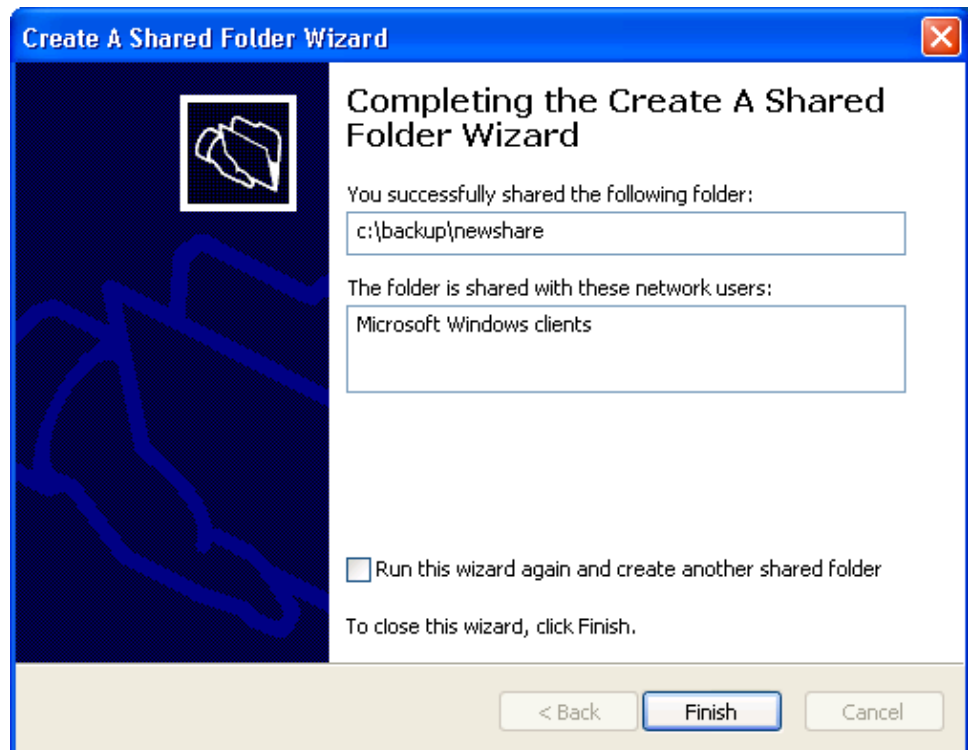
1. In the Control Panel, open Administrative Tools and select **Computer Management**.

2. Right-click **Shares** in the Shared Folders directory.
3. Select **New File Share** from the menu.

The **Create a Shared Folder** wizard opens. The computer name should be the name or IP address of the Data Domain system.

4. Enter the path for the Folder to share; for example, enter `C:\data\coll\backup\newshare`.
5. Enter the Share name; for example, enter `newshare`. Click **Next**.
6. For the Share Folder Permissions, selected Administrators have full access; other users have read-only access. Click **Next**.

**Figure 8** Completing the Create a Shared Folder Wizard



7. The Completing screen shows that you have successfully shared the folder with all Microsoft Windows clients in the network. Click **Finish**.

The newly created shared folder is listed in the Computer Management dialog box.

## Display CIFS Information

### Procedure

1. In the Control Panel, open Administrative Tools and select **Computer Management**.
2. Select one of the Shared Folders (**Shares, Sessions, or Open Files**) in the System Tools directory.

Information about shared folders, sessions, and open files is shown in the right pane.

## Managing Access Control

Managing Access Control topics in this section.

## Accessing Shares from a Windows Client

### Procedure

- ◆ From the Windows client use this DOS command:  
`net use drive: backup-location`

For example, enter:

```
net use H: \\dd02\backup /USER:dd02\backup22
```

This command maps the backup share from Data Domain system dd02 to drive H on the Windows system and gives the user named backup22 access to the \\DD\_sys\backup directory.

## Provide Domain Users Administrative Access

### Procedure

- ◆ Enter: `adminaccess authentication add cifs`

The SSH, Telnet, or FTP command that accesses the Data Domain system must include, in double quotation marks, the domain name, a backslash, and the user name. For example:

```
C:> ssh "domain2\djones" @ddr22
```

## Allow Access from Trusted Domain Users

You do not need to set this option because trusted domain users are always allowed to access shares from the Data Domain system.

### Procedure

- ◆ Enter: `cifs option set allowtrusteddomains enabled`

To disable access from trusted domain users, substitute `disabled` for `enabled`.

---

### Note

These are domains that are trusted by the domain that includes the Data Domain system.

---

## Allowing Administrative Access to a Data Domain System for Domain Users

### Procedure

1. To map a Data Domain System default group number to a Windows group name that differs from the default group name, use the `cifs option set "dd admin group2" ["windows grp-name"]` command.

The Windows group name is a group (based on one of the user roles—admin, user, or back-up operator) that exists on a Windows domain controller.



---

**Note**

For a description of DD OS user roles and Windows groups, see [Managing Data Domain Systems on page 31](#).

---

2. Enable CIFS administrative access by entering:

```
adminaccess authentication add cifs
```

- The default Data Domain System group `dd admin group1` is mapped to the Windows group Domain Admins.
- You can map the default Data Domain System group `dd admin group2` to a Windows group named Data Domain that you create on a Windows domain controller.
- Access is available through SSH, Telnet, FTP, HTTP, and HTTPS.
- After setting up administrative access to the Data Domain system from the Windows group `Data Domain`, you must enable CIFS administrative access using the `adminaccess` command.

## Restrict Administrative Access from Windows

### Procedure

1. Enter: `adminaccess authentication del cifs`

This command prohibits Windows users access to the Data Domain system if they do not have an account on the Data Domain system.

## File Access

File Access topics.

## NT Access Control Lists

Access control lists (ACLs) are enabled by default on the Data Domain system.

---

**Note**

When CIFS ACLs are disabled via the command `cifs option set ntfs-acls disabled`, the Data Domain system generates an ACL that approximates the UNIX permissions, whether or not there were previously set CIFS ACLs. For more detailed information about ACLs than is provided in this section, see your Windows Operating System documentation.

---


**CAUTION**

Data Domain recommends that you do not disable NTFS ACLs once they have been enabled. Contact Data Domain Support prior to disabling NTFS ACLs.

---

### Default ACL Permissions

The default permissions, which are assigned to new objects created through the CIFS protocol when ACLs are enabled, depend on the status of the parent directory. There are three different possibilities:

- ◆ The parent directory has no ACL because it was created through NFS protocol.

- ◆ The parent directory has an inheritable ACL, either because it was created through the CIFS protocol or because ACL had been explicitly set. The inherited ACL is set on new objects.
- ◆ The parent directory has an ACL, but it is not inheritable. The permissions are as follows:

Type	Name	Permission	Apply To
Allow	SYSTEM	Full control	This folder only
Allow	CREATOR OWNER	Full control	This folder only

#### Note

CREATOR OWNER is replaced by the user creating the file/folder for normal users and by Administrators for administrative users.

#### Permissions for a New Object when the Parent Directory Has No ACL

The permissions are as follows:

- ◆ BUILTIN\Administrators:(OI)(CI)F
- ◆ NT AUTHORITY\SYSTEM:(OI)(CI)F
- ◆ CREATOR OWNER:(OI)(CI)(IO)F
- ◆ BUILTIN\Users:(OI)(CI)R
- ◆ BUILTIN\Users:(CI)(special access:)FILE\_APPEND\_DATA
- ◆ BUILTIN\Users:(CI)(IO)(special access:)FILE\_WRITE\_DATA
- ◆ Everyone:(OI)(CI)R

These permissions are described in more detail below.

Type	Name	Permission	Apply To
Allow	Administrators	Full control	This folder, subfolders, and files
Allow	SYSTEM	Full control	This folder, subfolders, and files
Allow	CREATOR OWNER	Full control	Subfolders and files only
Allow	Users	Read & execute	This folder, subfolders, and files
Allow	Users	Create subfolders	This folder and subfolders only
Allow	Users	Create files	Subfolders only
Allow	Everyone	Read & execute	This folder, subfolders, and files

#### Setting ACL Permissions and Security

Windows-based backup and restore tools such as NetBackup can be used to back up DACL- and SACL-protected files to the Data Domain system, and to restore them from the Data Domain system.

#### Granular and Complex Permissions (DACL)

You can set granular and complex permissions (DACL) on any file or folder object within the DDFS file systems, either through using Windows commands such as `cacls`,

`xcaccls`, `xcopy` and `scopy`, or through the CIFS protocol using the Windows Explorer GUI.

### Audit ACL (SACL)

You can set audit ACL (SACL) on any object in the Data Domain File System (DDFS), either through commands or through the CIFS protocol using the Windows Explorer GUI.

## Set DACL Permissions Using the Windows Explorer

### Procedure

1. Right-click the file or folder and select **Properties** from the menu.
2. In the Properties dialog box, click the Security tab.
3. Select the group or user name, such as **Administrators**, from the list. The permissions appear, in this case for `Administrators`, `Full Control`.
4. Click the **Advanced** button, which enables you to set special permissions.
5. In the Advanced Security Settings for acl dialog box, click the Permissions tab.
6. Select the permission entry in the list.
7. To view more information about a permission entry, select the entry and click **Edit**.
8. Select the Inherit from parent option to have the permissions of parent entries inherited by their child objects, and click **OK**.

## Set SACL Permissions Using the Windows Explorer

### Procedure

1. Right-click the file or folder and select **Properties** from the menu.
2. In the Properties dialog box, click the Security tab.
3. Select the group or user name, such as **Administrators**, from the list, which displays its permissions, in this case, `Full Control`.
4. Click the **Advanced** button, which enables you to set special permissions.
5. In the Advanced Security Settings for ACL dialog box, click the Auditing tab.
6. Select the auditing entry in the list.
7. To view more information about special auditing entries, select the entry and click **Edit**.
8. Select the Inherit from parent option to have the permissions of parent entries inherited by their child objects, and click **OK**.

## View or Change the Current Owner Security ID (Owner SID)

### Procedure

1. In the Advanced Security Settings for ACL dialog box, click the Owner tab.
2. To change the owner, select a name from the Change owner list, and click **OK**.

## Controlling ID Account Mapping

The CIFS option `idmap-type` controls ID account mapping behavior. It has two values: `rid` (the default) and `none`. When the option is set to `rid`, the ID-to-id mapping is performed internally. When the option is set to `none`, all CIFS users are mapped to a local UNIX user named “`cifsuser`” belonging to the local UNIX group `users`.

---

**Note**

CIFS must be disabled to set this option. If CIFS is running, disable CIFS services. The idmap-type can be set to none only when ACL support is enabled.

---

Whenever the idmap type is changed, a file system metadata conversion might be required for correct file access. Without any conversion, the user might not be able to access the data. To converted the metadata, consult your contracted support provider.

## Monitoring CIFS Operation

Monitoring CIFS Operation topics.

### Display CIFS Status

#### Procedure

1. In the Data Domain System Manager, select **Data Management** > **CIFS**.
2. Check CIFS information, as follows:
  - Status is either enabled and running, or disabled but CIFS authentication is running.  
To enable CIFS, see [Enabling CIFS Services on page 207](#). To disable CIFS, see [Disable CIFS Services on page 210](#).
  - **Connections** lists the tally of open connections and open files.  
Click **Connection Details** to see more connection information.
  - Configuration details are described in [Display CIFS Information on page 215](#).
  - Share information is described in [Display Share Information on page 221](#).

### Display CIFS Configuration

Display CIFS Configuration topics.

#### Authentication Configuration

The information in the Authentication pane changes, depending on the type of authentication that is configured.

#### Active Directory Configuration

The following Authentication information is displays for Active Directory configuration:

Item	Description
Mode	The Active Directory mode displays.
Realm	The configured realm displays.
DDNS	The status of the DDNS Server displays: either enabled or disabled.
Domain Controller	The name of the configured domain controller displays or a * if all controllers are permitted.
Organizational Unit	The name of the configured organizational units displays.
CIFS Server Name	The name of the configured CIFS server displays.
WINS Server Name	The name of the configured WINS server displays.

Item	Description
Short Domain Name	The short domain name displays.

### Workgroup Configuration

The following Authentication information is displays for Workgroup configuration:

Item	Description
Mode	The Workgroup mode displays.
Workgroup Name	The configured workgroup name displays.
DDNS	The status of the DDNS Server displays: either enabled or disabled.
CIFS Server Name	The name of the configured CIFS server displays.
WINS Server Name	The name of the configured WINS server displays.

## Display Share Information

Display Share Information topics.

### Viewing Configured Shares

By default, the list of configured shares displays, showing the following:

Item	Description
Share Name	The name of the share (for example, share1).
Share Status	The status of the share: either enabled or disabled.
Directory Path	The directory path to the share (for example, /data/col1/backup/dir1).
	<p><b>Note</b></p> <p>col1 uses the letter ell followed by the number 1.</p>
Directory Path Status	The status of the directory path.

- ◆ To list information about a specific share, enter the share name in the Filter by Share Name text box and click **Update**.
- ◆ Click **Update** to return to the default list.
- ◆ To page through the list of shares, click the ◀ and ▶ arrows at the bottom right of the view to page forward or backward. To skip to the beginning of the list, click |◀ and to skip to the end, click ▶|.
- ◆ Click the **Items per Page** drop-down arrow to change the number of share entries listed on a page. Choices are 15, 30, or 45 entries.

### Viewing Detailed Share Information

To see detailed information about a share, click the share name in the share list. The following detailed information displays:

Item	Description
Share Name	The name of the share (for example, share1).
Directory Path	The directory path to the share (for example, /data/col1/backup/dir1).
	<b>Note</b> col1 uses the letter ell followed by the number 1.
Comment	The comment that was configured when the share was created.
Share Status	The status of the share: either enabled or disabled.
Number of ACE's	The number of Access Control Entries.

- ◆ The Clients area lists the clients that are configured to access the share, along with a client tally beneath the list.
- ◆ The User/Groups area lists the names and type of users or groups that are configured to access the share, along with a user or group tally beneath the list.
- ◆ The Options area lists the name and value of configured options.

## Display CIFS Statistics

### Procedure

- ◆ Enter: `cifs show detailed-stats`

The output shows number of various SMB requests received and the time taken to process them.

## Performing CIFS Troubleshooting

This section provides basic troubleshooting procedures.

### Note

The `cifs troubleshooting` commands provide detailed information about CIFS users and groups.

## Display Clients Current Activity

### Procedure

- ◆ Enter: `cifs show active`

### Results

The output shows shares accessed from a client system, current data transfer, and locked files.

PID	Username	Group	Machine
568	sysadmin	admin	svr24 (192.168.1.5)
566	sysadmin	admin	svr22 (192.168,1,6)

Services	PID	Machine	Connected at
ddvar	566	server22	Tues Jan 13 12:11:03 2009
backup	568	server24	Tues Jan 13 12:09:44 2009

The output for locked files provides the following information by file name or date.

PID	Deny Mode	Access	Read/Write	Oplock
566	DENY_WRITE	0x20089	RONLY	NONE
566	DENY_ALL	0x30196	WRONLY	NONE

## Set the Maximum Open Files on a Connection

### Procedure

- ◆ Use this command:  
`cifs option set maxopenfiles value`

The *value* for the maximum number of files that can be open concurrently on a given connection is an integer from 128 to 59412. The default is 10000.

If the system runs out of open files, increase the value's number.

Because each open file requires a certain amount of memory, the server may run out of memory if you set the value to the maximum. If a value is not within the accepted range, the system automatically resets it to 128 or 59412, depending on whether the value was below 128 or above 59412.

## Data Domain System Clock

When using active directory mode for CIFS access, the Data Domain System clock time can differ by no more than five minutes from that of the domain controller. The Data Domain System Manager System Settings > General Configuration Configure Time Settings option synchronizes the clock with a time server.

Because the Windows domain controller obtains the time from an external source, NTP must be configured. See the Microsoft documentation on how to configure NTP for the Windows operating system version or service pack that is running on your domain controller.

In active directory authentication mode, the Data Domain system periodically synchronizes the clock with a Windows Active Directory Domain Controller.

## Synchronizing from a Windows Domain Controller

### Note

This example is for Windows 2003 SP1; substitute your domain server for the NTP server's name (*ntpservername*).

### Procedure

1. On the Windows system, enter commands similar to the following:

```
C:\>w32tm /config /syncfromflags:manual /manualpeerlist: ntp-
server-name C:\>w32tm /config /update C:\>w32tm /resync
```

2. After NTP is configured on the domain controller, configure the time server synchronization, as described in [Working with Time and Date Settings on page 82](#).

## Synchronize from an NTP Server

To synchronize from an NTP server, configure the time server synchronization, as described in [Working with Time and Date Settings on page 82](#).



# CHAPTER 11

## Working with NFS

This chapter includes:

- ◆ [About NFS](#)..... 226
- ◆ [Managing NFS Client Access to the Data Domain System](#)..... 226
- ◆ [Displaying NFS Information](#)..... 230

## About NFS

Network File System (NFS) clients can have access to the system directories or MTrees on the Data Domain system.

- ◆ The `/backup` directory is the default destination for non-MTree compressed backup server data.
- ◆ The `/data/col1/backup` path is the root destination when using MTrees for compressed backup server data.
- ◆ The `/ddvar/core` directory contains Data Domain System core and log files (remove old logs and core files to free space in this area).

---

### Note

You can also delete core files from the `/ddvar` or the `/ddvar/ext` directory if it exists.

---

Clients, such as backup servers that perform backup and restore operations with a Data Domain System, need access to the `/backup` or `/data/col1/backup` areas. Clients that have administrative access need to be able to access the `/ddvar/core` directory to retrieve core and log files.

As part of the initial Data Domain system configuration, NFS clients were configured to access these areas. This chapter describes how to modify these settings and how to manage data access.

---

### Note

- ◆ For information about the initial system configuration, see the *EMC DD OS Initial Configuration Guide*.
  - ◆ The `nfs` command manages backups and restores between NFS clients and Data Domain systems, and it displays NFS statistics and status. For complete information about the `nfs` command, see the *EMC DD OS Command Reference Guide*.
  - ◆ For information about setting up third-party clients to use the Data Domain system as a server, see the related tuning guide, such as the *Solaris System Tuning*, which is available from the Data Domain support web site. From the Documentation > Integration Documentation page, select the vendor from the list and click **OK**. Select the tuning guide from the list.
- 

## Managing NFS Client Access to the Data Domain System

The topics in this section describe how to manage NFS client access to a Data Domain System.

### Enable NFS Services

To enable NFS services, which allows the client to access the system using the NFS protocol:

#### Procedure

1. Select the Data Domain system from the Navigational tree.

- The Summary page for the system displays.
2. Select the **Data Management** › **NFS** tabs.  
The NFS view showing the Exports tab appears.
  3. Click **Enable**.

## Disable NFS Services

To disable NFS services, which prevents the client access to the system using the NFS protocol.

### Procedure

1. Select the Data Domain system from the Navigational tree.  
The Summary page for the system displays.
2. Select the **Data Management** › **NFS** tabs.  
The NFS view showing the Exports tab appears.
3. Click **Disable**.

## Create an Export

You can use Data Domain System Manager's Create button on the NFS view or the Configuration Wizard to specify the NFS clients that can access the `/backup`, `/data/coll/backup`, and `/ddvar` areas.

---

### Note

You have to assign client access to each export separately and to remove access from each export separately. For example, a client can be removed from `/ddvar` and still have access to `/data/coll/backup`.

---

### Note

You can delete core files from the `/ddvar`, the `/dvar/core` or the `/ddvar/ext` directory if it exists.

---

### CAUTION

If Replication is to be implemented, a single destination Data Domain system can receive backups from both CIFS clients and NFS clients as long as separate directories or MTrees are used for each. Do not mix CIFS and NFS data in the same area.

---

### Procedure

1. Select the Data Domain system from the Navigational tree.  
The Summary page for this system is displayed.
2. Select the **Data Management** › **NFS** tabs.  
The NFS view showing the Exports tab appears.
3. Click **Create**.  
The Create NFS Exports dialog box is displayed.
4. Enter the pathname in the Directory Path text box (for example, `/data/coll/backup/dir1`).

---

**Note**

`coll` uses the lower-case letter L followed by the number 1.

---

5. In the Clients area, select an existing client or click the + icon to create a client.

The Clients dialog box is displayed.

- a. Enter a server name in the text box.

Enter fully qualified domain names, hostnames, or IP addresses. A single asterisk (\*) as a wild card indicates that all backup servers are to be used as clients.

---

**Note**

Clients given access to the `/data/coll/backup` directory have access to the entire directory. A client given access to a subdirectory of `/data/coll/backup` has access only to that subdirectory.

---

- A client can be a fully-qualified domain hostname, class-C IP addresses, IP addresses with either netmasks or length, an NIS netgroup name with the prefix @, or an asterisk (\*) wildcard with a domain name, such as **\*.yourcompany.com**.

A client added to a subdirectory under `/data/coll/backup` has access only to that subdirectory.

- Enter an asterisk (\*) as the client list to give access to all clients on the network.

- b. Select the checkboxes of the NFS options for the client.

- Read-only permission.
- (Default) Requires that requests originate on an Internet port that is less than IPPORT\_RESERVED (1024).
- Map requests from uid or gid 0 to the anonymous uid or gid.
- Map all user requests to the anonymous uid or gid.
- Use default anonymous UID or GID.

- c. Click **OK**.

6. Click **OK** to create the export.

## Modify an Export

**Procedure**

1. Select the Data Domain system from the Navigational tree.

The Summary page for this system is displayed.

2. Select the Data Management > NFS tabs.

The NFS view showing the Exports tab appears.

3. Click the checkbox of an export in the NFS Exports table.

4. Click **Modify**.

The Modify NFS Exports dialog box is displayed.

5. Modify the pathname in the Directory Path text box.

6. In the Clients area, select another client or click the + icon to create a client.

The Clients dialog box is displayed.

- a. Enter a server name in the text box.

Enter fully qualified domain names, hostnames, or IP addresses. A single asterisk (\*) as a wild card indicates that all backup servers are to be used as clients.

---

#### Note

Clients given access to the `/data/col1/backup` directory have access to the entire directory. A client given access to a subdirectory of `/data/col1/backup` has access only to that subdirectory.

- A client can be a fully-qualified domain hostname, class-C IP addresses, IP addresses with either netmasks or length, an NIS netgroup name with the prefix @, or an asterisk (\*) wildcard with a domain name, such as **\*.yourcompany.com**.  
A client added to a subdirectory under `/data/col1/backup` has access only to that subdirectory.
- Enter an asterisk (\*) as the client list to give access to all clients on the network.

- b. Select the checkboxes of the NFS options for the client.

- Read-only permission.
- (Default) Requires that requests originate on an Internet port that is less than IPPORT\_RESERVED (1024).
- Map requests from uid or gid 0 to the anonymous uid or gid.
- Map all user requests to the anonymous uid or gid.
- Use default anonymous UID or GID.

- c. Click **OK**.

7. Click **OK** to modify the export.

## Creating an Export from an Existing Export

To create an export from an existing export, and then modify as needed:

#### Procedure

1. In the NFS Exports table, click the checkbox of the export you wish to use as the source.
2. Click **Create From**.  
The Create NFS Export From dialog box appears.
3. Modify the export information, as described in [Modify an Export on page 228](#).

## Delete an Export

To delete an export:

#### Procedure

1. In the NFS Exports table, click the checkbox of the export you wish to delete.
2. Click **Delete**.  
The Warning dialog box is displayed.
3. Click **OK** and **Close** to delete the export.

## Displaying NFS Information

You can use the DD System Manager to monitor NFS client status and NFS configuration:

- ◆ [View NFS Status on page 230](#)
- ◆ [View NFS Exports on page 230](#)
- ◆ [View Active NFS Clients on page 230](#)

### View NFS Status

#### Procedure

1. Log into the Data Domain System Manager.
2. Select the Data Domain system in the Navigational tree.
3. Click the Data Management › NFS tabs.

The top pane shows the operational status of NFS; for example, NFS is currently active and running.

### View NFS Exports

To see the list of clients allowed to access the Data Domain System:

#### Procedure

1. Log into the Data Domain System Manager.
2. Select the Data Domain system in the Navigational pane.
3. Select the Data Management › NFS tabs.

The Exports view shows a table of NFS exports that are configured for Data Domain System and the mount path, status, and NFS options for each export.

4. Click an export in the table to populate the Detailed Information area, below the Exports table.

In addition to the export's directory path, configured options, and status, a list of clients displays.

### View Active NFS Clients

#### Procedure

1. Log into the Data Domain System Manager.
2. Select the Data Domain system in the Navigational pane.
3. Select the Data Management › NFS › Active Clients tabs.

The Active Clients view displays, showing all clients that have been connected in the past 15 minutes and their mount path.

Use the Filter By text boxes to sort by mount path and client name.

# CHAPTER 12

## Working with DD Boost

This chapter includes:

- ◆ [About Data Domain DD Boost Software](#)..... 232
- ◆ [Managing DD Boost with DD System Manager](#)..... 232
- ◆ [About Interface Groups](#)..... 240
- ◆ [Destroy DD Boost](#)..... 243
- ◆ [Managing Fibre Channel Transport](#)..... 243
- ◆ [Monitoring DD Boost](#)..... 245

## About Data Domain DD Boost Software

DD Boost software is an optional product that requires a separate license to operate on the Data Domain system. You can purchase a DD Boost software license key for a Data Domain system directly from EMC Data Domain.

DD Boost software enables backup servers to communicate with storage systems without the need for Data Domain systems to emulate tape. There are two components to DD Boost: one component that runs on the backup server and another that runs on the Data Domain system.

- ◆ In the context of the EMC NetWorker backup application, the component that runs on the backup server (DD Boost libraries) is integrated into the NetWorker storage node.
- ◆ In the context of Symantec backup applications (NetBackup and Backup Exec), you need to download an appropriate version of the DD Boost plugin that is installed on each media server. The DD Boost plugin includes the DD Boost libraries for integrating with the DD Boost server running on the Data Domain system.

A Data Domain system can be a single Data Domain system or a gateway.

The backup application (for example, Avamar, NetWorker, NetBackup, or Backup Exec) sets policies that control when backups and duplications occur. Administrators manage backup, duplication, and restores from a single console and can use all of the features of DD Boost, including WAN-efficient replicator software. The application manages all files (collections of data) in the catalog, even those created by the Data Domain system.

In the Data Domain system, storage units that you create are exposed to backup applications that use DD Boost protocol. For Symantec applications, storage units are viewed as a disk pool. For NetWorker, storage units are viewed as logical storage units (LSUs). A storage unit is considered an MTree; therefore, it supports MTree quota settings. (Do not create an MTree in place of a storage unit.)

This chapter does not contain installation instructions; refer to the documentation for the product you want to install. For example, for information about setting up DD Boost with Symantec backup applications (NetBackup and Backup Exec), see the *EMC Data Domain Boost for OpenStorage Administration Guide*.

Additional information about configuring and managing DD Boost on the Data Domain system can be found in the *EMC Data Domain Boost for OpenStorage Administration Guide*.

## Managing DD Boost with DD System Manager

To start managing DD Boost using DD System Manager:

### Procedure

1. Select the Data Domain system in the navigation panel.
2. Verify that the file system is enabled and running by clicking **Data Management** > **File System** and checking the state.
3. Select **Data Management** > **DD Boost** to access the DD Boost view.

If you go to the DD Boost page without a license, the Status states that DD Boost is not licensed. Click Add License and enter a valid license in the Add License Key dialog box.

- Use the DD Boost tabs— Settings, Active Connections, IP Network, Fibre Channel, Storage Units, and Statistics—to manage DD Boost. For more information, see [About the DD Boost Tabs on page 245](#).



## Results

The major DD Boost administration tasks include:

- ◆ [Set or Modify a DD Boost User Name on page 233](#)
- ◆ [Enable DD Boost on page 234](#)
- ◆ [Disable DD Boost on page 234](#)
- ◆ [Create a Storage Unit on page 235](#)
- ◆ [Delete a Storage Unit on page 237](#)
- ◆ [Clear DD Boost Statistics on page 238](#)
- ◆ [DD Boost Options on page 238](#)
- ◆ [Create Interface Groups on page 241](#)
- ◆ [Destroy DD Boost on page 243](#)
- ◆ [Managing Fibre Channel Transport on page 243](#)
- ◆ [Checking Interface Groups and Clients on page 247](#)

## Set or Modify a DD Boost User Name

A DD Boost user is a Data Domain user. Before setting the DD Boost user access, the username and password must have already been set up on the Data Domain system. See [Managing Access to the System on page 63](#) for more information.

- ◆ Backup applications use this user's credentials to connect to the Data Domain System. You must add the credentials to each backup server that connects to this Data Domain System. For complete information about setting up DD Boost with Symantec backup applications, NetBackup and Backup Exec, see the *EMC Data Domain Boost for OpenStorage Administration Guide*. For NetWorker, see the EMC NetWorker documentation.

Only one DD Boost user can operate DD Boost at a time. When DD Boost is enabled, a DD OS administrative user can set up another DD Boost user. When DD Boost is disabled, the administrative user can change the user.

To set or modify the DD Boost user name:

### Procedure

1. In the DD Boost Settings tab, click **Add** or **Modify** in the DD Boost User area.  
The Set or Modify DD Boost User Name dialog box is displayed.
2. To add a user, select **New User**.
  - a. Enter the user name in the DD Boost User Name text field.  
The user must be configured in the backup application to connect to the Data Domain system.
  - b. Enter the password twice in the text fields.
  - c. Select the role for the user: **admin**, **user**, **backup-operator**, or **data-access**.
3. To switch to a user that has already been added, select **Existing User** and select the user name from the menu.
4. Click **OK**.

## Enable DD Boost

DD Boost cannot be enabled without a DD Boost user. If you try to enable DD Boost without a user, you are asked to select one.

To enable DD Boost:

### Procedure

1. In the DD Boost Settings tab, click **Enable** in the DD Boost Status area.  
The Enable DD Boost dialog box is displayed.
2. Select an existing user name from the menu, or add a new user by supplying the name, password, and role.

## Disable DD Boost

Disabling DD Boost drops all active connections to the backup server. When you disable or destroy DD Boost, the DD Boost FC service is also disabled.

### Note

Ensure there are no jobs running from your backup application before disabling. File replication started by DD Boost between two Data Domain restores is not canceled.

To disable DD Boost:

### Procedure

1. In the DD Boost Settings tab, click **Disable** in the DD Boost Status area.
2. Click **OK** in the Disable DD Boost confirmation dialog box.

## View DD Boost Storage Unit

The DD Boost Storage Unit page:

- ◆ Lists the storage units and provides the following information for each storage unit:

Item	Description
Storage Unit	The name of the storage unit.
Quota Hard Limit	Percentage of hard limit quota used.
Last 24 hr Pre-Comp	Amount of raw data from the backup application that has been written in the last 24 hours.
Last 24 hr Post-Comp	Amount of storage used after compression in the last 24 hours.
Last 24 hr Comp Ratio	The compression ratio for the last 24 hours.
Weekly Avg Post-Comp	Average amount of compressed storage used in the last five weeks.
Last Week Post-Comp	Average amount of compressed storage used in the last seven days.
Weekly Avg Comp Ratio	The average compression ratio for the last five weeks.
Last Week Comp Ratio	The average compression ratio for the last seven days.

- ◆ Allows you to create a new storage unit and to delete an existing one selected from the list.
- ◆ Displays four related tabs for a storage unit selected from the list: Storage Unit, Space Usage, Daily Written, and Data Movement.
- ◆ Takes you to the Replication > DD Boost > File Replication tab when you click the **View DD Boost Replications** link.

---

#### Note

A DD Replicator license is required for DD Boost to display tabs other than the File Replication tab.

---

## Create a Storage Unit

You need to create at least one storage unit on the Data Domain system.

Each storage unit is a top-level subdirectory of the `/data/coll` directory; there is no hierarchy among storage units.

To create a storage unit:

#### Procedure

1. Go to the **Data Management > DD Boost > Storage Unit** tab.
2. Click the **Create** button.

The Create Storage Unit dialog box is displayed.

3. Enter the storage unit name in the Name text box.

Each storage unit name must be unique. Limit names to 50 characters maximum, with only 0-9, a-z, A-Z, and the n dash (-), and underscore (\_) allowed.

4. To set storage space restrictions to prevent a storage unit from consuming excess space: enter either a soft or hard limit quota setting, or both a hard and soft limit. With a soft limit an alert is sent when the storage unit size exceeds the limit, but data can still be written to it. Data cannot be written to the storage unit when the hard limit is reached.

---

#### Note

Quota limits are pre-compressed values.

To set quota limits, select **Set to Specific Value** and enter the value. Select the unit of measurement: MiB, GiB, TiB, or PiB.

---



---

#### Note

When setting both soft and hard limits, a quota's soft limit cannot exceed the quota's hard limit.

---

5. Click **OK**.
6. Repeat the above steps for each Data Domain Boost-enabled system.

## View Storage Unit Information

After selecting a storage unit from the list, the following tabs provide more information about it.

#### Storage Unit Tab

The Storage Unit tab shows detailed information for a selected storage unit in its Summary and Quota pane. The Snapshot pane shows snapshot details, allows you to

create new snapshots and schedules, and provides a link to the **Data Management > Snapshots** tab.

◆ **Summary**

Summary Item	Description
Total Files	The total number of file images on the storage unit. For compression details that you can download to a log file, click the Download Compression Details link. The generation can take up to several minutes. After it has completed, click Download.
Full Path	/data/coll/filename
Status	R: read; W: write; Q: quota defined
Pre-Comp Used	The amount of pre-compressed storage already used.

◆ **Quota**

Quota Item	Description
Quota Enforcement	Enabled or disable. Clicking Quota takes you to the <b>Data Management &gt; Quota</b> tab where you can configure quotas.
Pre-Comp Soft Limit	Current value of soft quota set for the storage unit.
Pre-Comp Hard Limit	Current value of hard quota set for the storage unit.
Quota Summary	Percentage of Hard Limit used.

To modify the pre-comp soft and hard limits shown in the tab:

1. Click the **Configure** button in the Quota pane.
2. In the Configure Quota dialog box, enter values for hard and soft quotas and select the unit of measurement: MiB, GiB, TiB, or PiB. Click **OK**.

◆ **Snapshot**

The Snapshot pane shows details about the storage unit's snapshots.

Item	Description
Total Snapshots	The total number of snapshots created for this MTree. A total of 750 snapshots can be created for each MTree.
Expired	The number of snapshots in this MTree that have been marked for deletion, but have not been removed with the clean operation as yet.
Unexpired	The number of snapshots in this MTree that are marked for keeping.
Oldest Snapshot	The date of the oldest snapshot for this MTree.
Newest Snapshot	The date of the newest snapshot for this MTree.
Next Scheduled	The date of the next scheduled snapshot.
Assigned Snapshot Schedules	The name of the snapshot schedule assigned to this MTree.

- Do one of the following:
  - Assign a snapshot schedule to a selected storage unit: Click **Assign Snapshot Schedules**. Select the schedule's checkbox; click **OK** and **Close**.

- Create a new schedule: Click **Assign Snapshot Schedules**. Enter the new schedule's name.

---

#### Note

Snapshot name pattern can be composed only of letters, numbers, `_`, `-`, `%d`, `%a`, `%m`, `%b`, `%y`, `%Y`, `%H`, and `%M`, following the pattern shown in the dialog box.

Enter the new pattern and click **Validate Pattern & Update Sample**. Click **Next**.

- Select when the schedule is to be executed: weekly, every day (or selected days), monthly on specific days that you select by clicking that date in the calendar, or on the last day of the month. Click **Next**.
- Enter the times of the day when the schedule is to be executed: Either select **At Specific Times** or **In Intervals**. If you select a specific time, select the time from the list. Click **Add (+)** to add a time (24-hour format). For intervals, select **In Intervals** and set the start and end times and how often (Every), such as every eight hours. Click **Next**.
- Enter the retention period for the snapshots in days, months, or years. Click **Next**.
- Review the Summary of your configuration. Click **Back** to edit any of the values. Click **Finish** to create the schedule.

- 
- Clicking the Snapshots link takes you to the **Data Management > Snapshots** tab.

#### Space Usage Tab

The Space Usage tab graph displays a visual representation of data usage for the storage unit over time. There are three views of space usage that you can select from the menu: MTree, Active, and Archive (Extended Retention).

- ◆ Click a point on a graph line to display a box with data at that point.
- ◆ Click **Print** (at the bottom on the graph) to open the standard Print dialog box.
- ◆ Click **Show** in new window to display the graph in a new browser window.

There are two types of graph data:

- ◆ **Pre-comp Written**—The total amount of data sent to the storage unit by backup servers. Pre-compressed data is what a backup server sees as the total uncompressed data held by the storage unit. Shown with the Space Used (left) vertical axis of the graph.
- ◆ **Post-comp Used**—The total amount of disk storage in use on the storage unit. Shown with the Space Used (left) vertical axis of the graph.

#### Daily Written Tab

The Daily Written view contains a graph that displays a visual representation of data that is written daily to the system over a period of time, selectable from 7 to 120 days. The data amounts are shown over time for pre- and post-compression amounts.

See [About the Space Usage View on page 138](#) for more information.

#### Data Movement Tab

A graph in the same format as the Daily Written graph that shows the amount of disk space moved to the Extended Retention storage area (if the Archive license is enabled).

## Delete a Storage Unit

Deleting a storage unit removes all images contained in the storage unit.

To delete a storage unit:

**Procedure**

1. Go to the **Data Management** › **DD Boost** › **Storage Unit** tab.
2. Select the storage unit to be deleted from the list.
3. Click **Delete (X)**.
4. Enter the system administration password, and click **OK**.

**Results**

The storage unit is removed from your Data Domain system. You must also manually remove the corresponding backup application catalog entries.

## Clear DD Boost Statistics

When this option is used, all DD Boost statistics are removed from the system and cannot be recovered.

**Note**

DD Boost must be enabled for statistics to be removed.

To clear DD Boost statistics:

**Procedure**

1. Go to the **Data Management** › **DD Boost** › **Statistics** tab.
2. In the Statistics pane, click the **Clear Statistics** button.

The Clear Statistics dialog box warns you that all previous statistics will be lost and asks you to confirm that you want to reset all DD Boost statistics.

3. Click **OK** to clear the statistics.

## DD Boost Options

The following DD Boost options are available from **Data Management** › **DD Boost** › **More Tasks**:

- ◆ [Distributed Segment Processing on page 239](#)
- ◆ [Virtual Synthetics on page 239](#)
- ◆ [Low Bandwidth Optimization on page 305](#)
- ◆ [File Replication Encryption on page 239](#)

The current status (enabled or disabled) of these options is shown in the Advanced Option area of the DD Boost Settings tab. Expand Advanced Options to view the list.

**Note**

You can manage distributed segment processing via the `ddboost` option commands, which are described in detail in the *EMC Data Domain Operating System Command Reference Guide*.

To enable or disable a DD Boost option:

**Procedure**

1. From the More Tasks menu on the DD Boost page, select **Set Options**.
2. Select any option to be enabled.
3. Deselect any option to be disabled.

4. Click **OK**.

## Distributed Segment Processing

Distributed segment processing increases backup throughput in almost all cases by eliminating duplicate data transmission between the media server and the Data Domain system.

---

### Note

- ◆ You can manage distributed segment processing via the `ddboost option` commands, which are described in detail in the *EMC Data Domain Operating System Command Reference Guide*.
  - ◆ Distributed segment processing is enabled by default with EMC Data Domain Extended Retention (formerly Data Domain Archiver) configurations and cannot be disabled.
- 

## Virtual Synthetics

The virtual synthetic full backup is the combination of the last full (synthetic or full) backup and all subsequent incremental backups.

Virtual synthetics are disabled by default. Enable this option before you configure the backup application for use.

## Low-Bandwidth Optimization

If you utilize file replication over a low-bandwidth network (WAN) you can increase replication speed by using low bandwidth optimization. This feature provides additional compression during data transfer. Low bandwidth compression is available to Data Domain systems with an installed Replication license.

Low-bandwidth optimization, which is disabled by default, is designed for use on networks with less than 6 Mbps aggregate bandwidth. Do not use this option if maximum file system write performance is required.

---

### Note

You can also manage low bandwidth optimization via the `ddboost file-replication` commands, which are described in detail in the *EMC Data Domain Operating System Command Reference Guide*.

---

## File Replication Encryption

You can encrypt the data replication stream by enabling its DD Boost Option.

---

### Note

For encryption other than for systems with the Data at Rest option: If DD Boost file-replication encryption is set to on, it must be set to on for both the source and destination systems.

---

### Managed File Replication TCP Port Setting

For DD Boost managed file replication, set the global listen port the same on both the source and target Data Domain systems. Use the `replication option` command for the listen-port to manage this setting as described in the *EMC Data Domain Operating System Command Reference Guide*.

## About Interface Groups

Configuring an interface group creates a private network within the Data Domain system, comprised of the IP addresses designated as a group. Clients are assigned to a single group by specifying client name (`client.emc.com`) or wild card name (`*.emc`). The group interface uses the Advanced Load Balancing and Failover feature to improve data transfer performance and increase reliability.

For example, in the Symantec NetBackup environment, media server clients use a single public network IP address to access the Data Domain system. All communication with the Data Domain system is initiated via this administered IP connection, which is configured on the NetBackup server.

If an interface group is configured, when the Data Domain system receives data from the media server clients, the data transfer is load-balanced and distributed on all the interfaces in the group, providing higher input/output throughput, especially for customers who use multiple 1 GigE connections.

The data transfer is load-balanced based on the number of connections outstanding on the interfaces. Only connections for backup and restore jobs are load-balanced. Check the Active Connections for more information on the number of outstanding connections on the interfaces in a group (see [Checking Activities on page 246](#) for details).

Should an interface in the group fail, all the in-flight jobs to that interface are automatically resumed on healthy operational links (unknown to the backup applications). Any jobs that are started subsequent to the failure are also routed to a healthy interface in the group. If the group is disabled or an attempt to recover on an alternate interface fails, the administered IP is used for recovery. Failure in one group will not utilize interfaces from another group.

---

### Note

- ◆ The IP address must be configured on the Data Domain system, and its interface enabled. To check the interface configuration, see the Network Settings tab in the **Hardware > Network** page, and check for free ports. See the `net` chapter of the *EMC Data Domain Operating System Command Reference Guide* or the *EMC Data Domain Operating System Initial Configuration Guide* for information about configuring an IP address for an interface.
- ◆ A client or wild card client must be configured on the Data Domain system for each group. See the `net` chapter of the *EMC Data Domain Operating System Command Reference Guide* or the *EMC Data Domain Operating System Initial Configuration Guide* for information about how to configure a client.
- ◆ You can also manage Advanced Load Balancing and Failover via the `ddbboost ifgroup` commands, which are described in detail in the *EMC Data Domain Operating System Command Reference Guide*.

---

Configured interfaces are listed in Active Connections, on the lower portion of the Activities page (see [Checking Activities on page 246](#)).

The following management options are available:

- ◆ [Create Interface Groups on page 241](#)
- ◆ [Enable/Disable an Interface Group on page 241](#)
- ◆ [Modify an Interface Group's Name/Interfaces on page 242](#)



- ◆ [Delete a Client from the Interface Group on page 242](#)
- ◆ [Modify a Client's Name or Interface Group on page 242](#)

## Create Interface Groups

Use this option to select the interfaces that are used in interface groups. Multiple interface groups improve the efficiency of DD Boost by allowing you to do the following:

- ◆ Configure DD Boost to use specific interfaces configured into groups.
- ◆ Assign clients to one of those interface groups.
- ◆ Monitor which interfaces are active with DD Boost clients.

First create interface groups. then add clients (as new media servers become available) to an interface group:

Follow these steps:

### Procedure

1. Select the Add (+) button associated with interface groups.
2. Enter the interface group name.
3. Select one or more interfaces. A maximum of 32 interfaces can be configured.

---

#### Note

Depending upon aliasing configurations, some interfaces may not be selectable if they are sharing a physical interface with another interface in the same group. This is because each interface within the group must be on a different physical interface to ensure fail-over recovery.

4. Click **OK**.
5. Select Add (+) associated with clients.
6. Enter a fully qualified client name or **\*.mydomain.com**.

---

#### Note

The \* client is initially available to the default group. The \* client may only be a member of one ifgroup.

7. Select a previously configured interface group, and click **OK**.

## Delete an Interface Group

To delete the interface group, which also deletes all associated interfaces and clients:

### Procedure

1. On the IP Network page, select the interface group in the list. The default group cannot be deleted.
2. Click the associated Delete (X) button.
3. Confirm the deletion.

## Enable/Disable an Interface Group

### Procedure

1. On the IP Network page, select the interface group in the list.

---

**Note**

If the interface group does not have both clients and interfaces assigned, you cannot enable the group.

---

2. Click the associated **Edit** (pencil) button.
3. Select the **Enabled** button to enable; deselect to disable.
4. Click **OK**.

## Modify an Interface Group's Name/Interfaces

**Procedure**

1. On the IP Network page, select the interface group in the list.
2. Click the associated **Edit** (pencil) button.
3. Retype the name to modify the name.

The group name must be one to 24 characters long and contain only letters, numbers, underscores, and dashes. It cannot be the same as any other group name and cannot be "default", "yes", "no", or "all."

4. Select or deselect client interfaces in the Interfaces list.
- 

**Note**

If you remove all interfaces from the group, it will be automatically disabled.

---

5. Click **OK**.

## Delete a Client from the Interface Group

**Procedure**

1. On the IP Network page, select the client in the list.
  2. Click the associated Delete (X) button.
- 

**Note**

If the interface group to which it belongs has no other clients, the interface group is disabled.

---

3. Confirm the deletion.

## Modify a Client's Name or Interface Group

**Procedure**

1. On the IP Network page, select the client in the list.
2. Click the associated **Edit** (pencil) button.
3. Type a new client name.

Client names must be unique and either **FQDN**, **\*.domain**, or for the default group only, an asterisk (\*). Client names have a maximum length of 128 characters.

4. Select a new interface group from the menu.

---

**Note**

The old interface group is disabled if it has no clients.

---

5. Click **OK**.

## Destroy DD Boost

This option permanently removes all of the data (images) contained in the storage units. When you disable or destroy DD Boost, the DD Boost FC service is also disabled.

Only an administrative user can destroy DD Boost.

**Procedure**

1. Manually remove (expire) the corresponding backup application catalog entries.
- 

**Note**

If multiple backup applications are using the same Data Domain system, then remove all entries from each of those applications' catalogs.

---

2. From the More Tasks menu, select **Destroy DD Boost**.
3. Enter your administrative credentials when prompted.
4. Click **OK**.

## Managing Fibre Channel Transport

You can use Fibre Channel transport with DD Boost by using the DD Boost over Fibre Channel service.

Go to the **Data Management > DD Boost > Fibre Channel** tab to view the Fibre Channel page. Use this page to manage Fibre Channel transport and access groups. An access group is created to hold a collection of initiator WWPNs and the drives and changers they are allowed to access. The WWPN is the unique World-Wide Port Name of the Fibre Channel port in the media server. There are separate access groups for VTL and DD Boost protocols.

To view information about access groups configured for the Data Domain system, click **View All Access Groups**.

---

**Note**

Avoid making access group changes on a Data Domain system during active backup or restore jobs. A change may cause an active job to fail. The impact of changes during active jobs depends on a combination of backup software and host configurations.

---

To enable or disable the Fibre Channel transport, click the **Enable/Disable** Status button.

You can also perform these tasks via the Fibre Channel tab:

- ◆ [Set Fibre Channel Server Name on page 243](#)
- ◆ [Create Access Group on page 244](#)
- ◆ [Delete Access Groups on page 244](#)

### Set Fibre Channel Server Name

To modify the server name so that the Data Domain system is uniquely identified on the fibre channel transport, click **Data Management > DD Boost > Fibre Channel**. The Server

Name field displays the default server name, which is the hostname by default. To change the default server name, click **Edit**, enter a new server name, and click **OK**.

## Create Access Group

---

### Note

Initiators and endpoints for access points are managed and configured via the Hardware > Fibre Channel tab.

---

### Procedure

1. In the Fibre Channel page's DD Boost Access Groups area, click Add (+) to add a group.
  2. Enter a unique name. Duplicate access groups are not supported.
  3. Select one or more initiators. Optionally, replace the initiator name by entering a new one. The WWPN is the unique World-Wide Port Name of the Fibre Channel port in the media server.
- 

### Note

An initiator is a backup client that connects to the system for the purpose of reading and writing data using the Fibre Channel protocol. A specific initiator can support DD Boost over FC or VTL, but not both.

---

4. Devices: The number of DD Boost devices to be used by the group, which determines which devices the initiator can discover and, therefore, the amount of I/O paths to the Data Domain system. The default is one, the minimum is one, and the maximum is 64.
- 

### Note

If you are working with Linux clients, you do not need to change the default. If you are working with Windows clients see the *EMC Data Domain Boost for OpenStorage Administration Guide* for the recommended value.

---

5. Indicate whether endpoints apply to all, none, or select from the list of endpoints.
- 

### Note

Devices can be mapped to one or more endpoints through an access group.

---

6. Summary: Review the Summary and make any modifications. Click **Finish** to create the access group, which is displayed in the DD Boost Access Groups list.
- 

### Note

To change settings for an existing access group, select it from the list and click **Edit** (pencil).

---

## Delete Access Groups

### Procedure

1. Go to the **Data Management > DD Boost > Fibre Channel** tab.
2. Select the group to be deleted from the DD Boost Access Groups list.

---

**Note**

You cannot delete a group that has initiators assigned to it. Edit the group to remove the initiators first.

---

3. Click Delete (X).

## Monitoring DD Boost

To check DD Boost status and activity, look at the Settings and other tabs, described in the following sections.

### About the DD Boost Tabs

Tabs include:

- ◆ Settings
- ◆ Active Connections (see [Checking Activities on page 246](#))
- ◆ IP Network (see [Checking Interface Groups and Clients on page 247](#)).
- ◆ Fibre Channel (see [Checking Storage Units on page 247](#)).
- ◆ Storage Units (see [Checking Storage Units on page 247](#))
- ◆ Statistics (see [Checking DD Boost Statistics on page 248](#))

### Settings

The Settings tab shows the DD Boost status (Enabled or Disabled) and the name of the DD Boost user. Use the **Status** button to switch between **Enabled** or **Disabled**. Use the **Edit** (pencil) button to select another authorized user.

The Settings tab lists the allowed clients and shows whether or not the Advanced Options are enabled or disabled. You can change the status of these options by selecting **More Tasks** > **Set Options**.

Also use this tab to set up the media servers that have access to DD Boost protocol. See [Set Up Media Servers on page 245](#).

### Set Up Media Servers

Use the Allowed Clients section of the Settings tab to control the number of clients with access to the DD Boost protocol.

To create access clients, or modify existing client's names:

#### Procedure

1. From the DD Boost tab, click the Settings tab.
2. Delete the \* client by selecting it and then clicking the associated Delete (X) button. Click **OK**.
3. Click Add (+) to add a new client.
4. Enter the client name and click **OK** to add the client. The list refreshes to show the new client.

Client names must be unique and either FQDN, \*.domain, or for the default group only, an asterisk (\*). Client names have a maximum length of 128 characters.

5. To modify an existing client name, click **Edit** (pencil) and enter a new name.

6. Add the client names for the other two media servers.

## Checking Activities

The Active Connections page lists the following information:

- ◆ **Clients**—Shows the following information for a connected client.

Item	Description
Client	The name of the connected client.
Idle	Whether the client is idle (Yes) or not (No).
CPUs	The number of CPUs that the client has, such as 8.
Memory (GiB)	The amount of memory (in GiB) the client has, such as 7.8.
Plug-In Version	The DD Boost plug-in version installed, such as 2.2.1.1.
OS Version	The operating system version installed, such as Linux 2.6.17-1.2142_FC4smp x86_64.
Application Version	The backup application version installed, such as NetBackup 6.5.6.

- ◆ **Interfaces**—Shows the following information about configured interface connections:

Item	Description
Interface	The IP address of the interface.
Interface Group	One of the following: <ul style="list-style-type: none"> <li>• The name of the interface group.</li> <li>• None, if not a member of one.</li> </ul>
Backup	The number of active backup connections.
Restore	The number of active restore connections.
Replication	The number of active replication connections.
Total	The total number of connections for the interface.

- ◆ **Out-Bound File Replications**—Shows the following information for out-bound files:

Out-bound Files Item	Description
File Name	The name of the out-going image file.
Target Host	The name of the host receiving the file.
Logical Bytes to Transfer	The number of logical bytes to be transferred.
Logical Bytes Transferred	The number of logical bytes already transferred.
Low Bandwidth Optimization	The number of low-bandwidth bytes already transferred.

## Checking Interface Groups and Clients

The IP Network tab lists configured interface groups. Details include whether or not a group is enabled, and any configured client interfaces. The administrator can also use the Interface Group menu to view which clients are associated with an interface group.

## Checking Storage Units

The Storage Unit page provides a button to create a storage unit and a button to delete one or more selected storage units. It lists the names of the storage units that have been created at the top of the page. For more information about any storage unit, select it in the list, which displays its details.

Details for a Storage Unit include:

Item	Description
Existing Storage Units	
Storage Unit Name	The name of the storage unit.
Pre-Comp Used	The amount of pre-compressed storage already used.
Pre-Comp Soft Limit	Current value of soft quota set for the storage unit.
% of Pre-Comp Soft Limit Used	Percentage of hard limit quota used.
Pre-Comp Hard Limit	Current value of hard quota set for the storage unit.
% of Pre-Comp Hard Limit Used	Percentage of hard limit quota used.
Storage Unit Details	Select the storage unit in the list.
Total Files	The total number of file images on the storage unit.
Download Files	Link to download storage unit file details in .tsv format. You must allow pop-ups to use this function.
Compression Ratio	The compression ratio achieved on the files.
Metadata Size	The amount of space used for metadata information.
Storage Unit Status	The current status of the storage unit (combinations are supported). Status can be: <ul style="list-style-type: none"> <li>• D—Deleted</li> <li>• RO—Read-only</li> <li>• RW—Read/write</li> <li>• RD—Replication destination</li> <li>• RLE—Retention lock enabled</li> <li>• RLD—Retention lock disabled</li> </ul>
Quota Enforcement	Click Quota to go to the Data Management Quota page, which lists hard and soft quota values/percentage used by MTrees.
Quota Summary	Percentage of Hard Limit used.
Original Size	The size of the file before compression was performed.
Global Compression Size	The total size after global compression of the files in the storage unit when they were written.

Item	Description
Locally Compressed Size	Total size after local compression of the files in the storage unit when they were written.

## Checking DD Boost Statistics

The DD Boost Statistics page has a button in the Statistics area you can click to **Clear Statistics**.

This page covers three categories of statistics.

### Histogram Statistics

This area displays the latencies of DD Boost operations in the form of a histogram that Data Domain can use to analyze performance.

### Statistics

This area displays the statistics list count and error amounts for DD Boost operations.

### File Statistics

This area displays the file statistics list count and error amounts for the following:

- ◆ File creates
- ◆ File deletes
- ◆ Pre-compressed bytes received
- ◆ Bytes after filtering
- ◆ Bytes after local compression
- ◆ Network bytes received
- ◆ Compression Ratio
- ◆ Total bytes read



# CHAPTER 13

## Working with DD Virtual Tape Library

This chapter includes:

- ◆ [About EMC Data Domain Virtual Tape Library](#)..... 250
- ◆ [Planning a VTL](#)..... 250
- ◆ [About the DD System Manager VTL View](#)..... 254
- ◆ [Setting Up a VTL](#)..... 255
- ◆ [Working with the VTL Service Operations](#)..... 255
- ◆ [Working with Storage Pools](#)..... 276
- ◆ [Working with a Single Storage Pool](#)..... 280

## About EMC Data Domain Virtual Tape Library

EMC Data Domain Virtual Tape Library (VTL) is a disk-based backup system that emulates the use of physical tapes. It enables backup applications to connect to and manage Data Domain system storage using functionality almost identical to a physical tape library.

These virtual tape drives are accessible to backup software in the same way as physical tape drives. After you create these drives in a VTL, they appear to the backup software as SCSI tape drives. The VTL, itself, appears to the backup software as a SCSI robotic device accessed through standard driver interfaces. However, the movement of the media changer and backup images is managed by the backup software – not by the Data Domain system configured as a VTL.

The following terms have a special meaning when used with VTL:

- ◆ *Library* – A library emulates a physical tape library with tape drives, changer, CAPs (cartridge access ports), and slots (cartridge slots).
- ◆ *Tapes* – Tapes are represented as files. Tapes can be imported from the vault to a library. Tapes can be exported from a library to the vault. Tapes can be moved within a library across drives, slots, and CAPs.
- ◆ *Pool* – A pool is a collection of tapes that map to a directory on the file system. A pool is used to replicate tapes to a destination.

---

### Note

You can convert directory-based VTL pools to MTrees to take advantage of the greater functionality of MTrees.

- ◆ *Vault* – The vault holds tapes not being used by any library. Tapes reside in either a library or the vault.

VTL has been tested with, and is supported by, specific backup software and hardware configurations listed in the VTL matrices. For specific backup software and hardware configurations tested and supported by Data Domain, see [Application Compatibility Matrices and Integration Guides on page 21](#).

VTL supports simultaneous use of the tape library and file system (NFS/CIFS/DD Boost) interfaces.

---

### Note

The maximum supported I/O size for any Data Domain system using the VTL is 1 MB.

When DR (disaster recovery) is needed, pools and tapes can be replicated to a remote Data Domain system using the Data Domain Replicator. See [Replicating Storage Pools on page 279](#).

To protect data on tapes from modification, tapes can be locked using the Retention Lock Governance software. See [Changing a Tape's Write or Retention Lock State on page 270](#).

## Planning a VTL

Before using a VTL (Virtual Tape Library), you need:

- ◆ The appropriate VTL license.  
VTL is a licensed feature, and it is required to use NDMP over IP or to use VTL directly over Fibre Channel. An additional license is required for IBM i systems. This license is the I/OS license.

(See your EMC Data Domain sales representative to purchase licenses.)

To activate a license, see [Managing System Licenses on page 37](#).

Adding a VTL license through the Data Domain System Manager automatically disables and enables the VTL feature.

- ◆ An installed FC (Fibre Channel) interface card or VTL configured to use NDMP (Network Data Management Protocol).
  - If the VTL communication between a backup server and a Data Domain system is through an FC interface, the Data Domain system must have an FC interface card installed.
  - If the VTL communication between a backup server and a Data Domain system is through NDMP, no FC interface card is required. However, you must configure the Tape Server Access Group. Also, when using NDMP, all initiator and port functionality does not apply. See [Configuring the NDMP Device TapeServer Group on page 288](#).
- ◆ A backup software minimum record (block) size.  
EMC Data Domain strongly recommends that backup software be set to use a minimum record (block) size of 64 KiB or larger. Larger sizes usually give faster performance and better data compression.

**CAUTION**

**Depending on your backup application, if you change the size after the initial configuration, data written with the original size might become unreadable.**

---

- ◆ Appropriate user access to the system.  
For basic tape operations and monitoring, only a user login is required. To enable and configure the VTL service and perform other configuration tasks, a sysadmin login is required.

## Limitations of the VTL Feature

The following limitations should be considered when setting up a VTL (Virtual Tape Library):

- ◆ Libraries – VTL supports a maximum of 64 libraries per system (that is, 64 VTL instances on each Data Domain system).
- ◆ Tape Drives and Data Streams – See the following sections for these numbers.
- ◆ Initiators – VTL supports a maximum of 128 initiators or WWPNs (world-wide port names) in an access group.
- ◆ Slots – VTL supports a maximum of:
  - 32,000 slots per library
  - 64,000 slots per system

**Note**

The system automatically adds slots to keep the number of slots equal to or greater than the number of drives. However, some hosts do not support as many slots as the Data Domain system within a library.

---

- ◆ CAPs (cartridge access ports) – VTL supports a maximum of:
  - 100 CAPs per library
  - 1000 CAPs per system

## Number of Supported Tape Drives

The maximum number of tape drives supported by a VTL depends on the number of CPU Cores and the amount of memory installed (both RAM and NVRAM, if applicable) on your Data Domain system.

### Note

There are no references to model numbers in this table because there are many combinations of CPU Cores and memories for each model, and the number of supported drives depends *only* on the CPU Cores and memories and not on the particular model, itself.

**Table 23** Number of Tape Drives Supported by VTL

Number of CPU Cores	RAM (in GB)	NVRAM (in GB)	Maximum Number of Supported Tape Drives
Fewer than 32	4 or less	NA	64
	More than 4, up to 38	NA	128
	More than 38, up to 128	NA	256
	More than 128	NA	540
32 to 39	Up to 128	Less than 4	270
	Up to 128	4 or more	540
	More than 128	NA	540
40 or more	NA	NA	540

## Number of Supported Data Streams

See the table, “Data Streams Sent to a Data Domain System,” in [Data Streams Sent to a Data Domain System on page 131](#), for the maximum stream limit for each Data Domain system.

## About Tape Barcodes

When a tape is created, you assign a barcode that is a unique identifier of that tape. You should avoid using duplicate barcodes, which could cause unpredictable behavior in backup applications.

The eight-character barcode must start with six numeric or uppercase alphabetic characters (from the set {0-9, A-Z}) and end in a two-character tag for the supported tape type.

Tape Code	Capacity	Tape Type
L1	100 GiB	LTO-1
L2	200 GiB	LTO-2
L3	400 GiB	LTO-3
L4	800 GiB	LTO-4

Tape Code	Capacity	Tape Type
LA <sup>a</sup>	50 GiB	LTO-1
LB	30 GiB	LTO-1
LC	10 GiB	LTO-1

a. For TSM, use the L2 tape code if the LA code is ignored.

These capacities are the default sizes used if the capacity option is not included when creating the tape cartridge. If a capacity value is included, it overrides the two-character tag.

The numeric characters immediately to the left of L set the number for the first tape created. To make use of automatic incrementing of the barcode when creating more than one tape, Data Domain numbering starts at the sixth character position, just before L. If this is a digit, the system increments it. If an overflow occurs (9 to 0), numbering moves one position to the left. If the next character to increment is an alphabetic character, incrementation stops.

For example, a barcode of ABC100L1 starts numbering the tapes at 100 and can go to a maximum of 999.

Here are a few sample barcodes:

- ◆ 000000L1 creates tapes of 100 GiB capacity and can accept a count of up to 100,000 tapes (from 000000 to 99999).
- ◆ AA0000LA creates tapes of 50 GiB capacity and can accept a count of up to 10,000 tapes (from 0000 to 9999).
- ◆ AAAA00LB creates tapes of 30GiB capacity and can accept a count of up to 100 tapes (from 00 to 99).
- ◆ AAAAAALC creates one tape of 10 GiB capacity. Only one tape can be created with this name.
- ◆ AAA350L1 creates tapes of 100 GiB capacity and can accept a count of up to 650 tapes (from 350 to 999).
- ◆ 000AAALA creates one tape of 50 GiB capacity. Only one tape can be created with this name.
- ◆ 5M7Q3KLB creates one tape of 30 GiB capacity. Only one tape can be created with this name.

## About LTO Tape Drive Compatibility

The following table shows the levels of compatibility among the different generations of LTO (Linear Tape-Open) technology.

In this table:

- ◆ RW - read and write compatible
- ◆ R - read-only compatible
- ◆ — - not compatible

**Table 24** LTO Tape Drive Compatibility

Tape Format	LTO-4	LTO-3	LTO-2	LTO-1
LTO-4	RW	—	—	—

**Table 24** LTO Tape Drive Compatibility (continued)

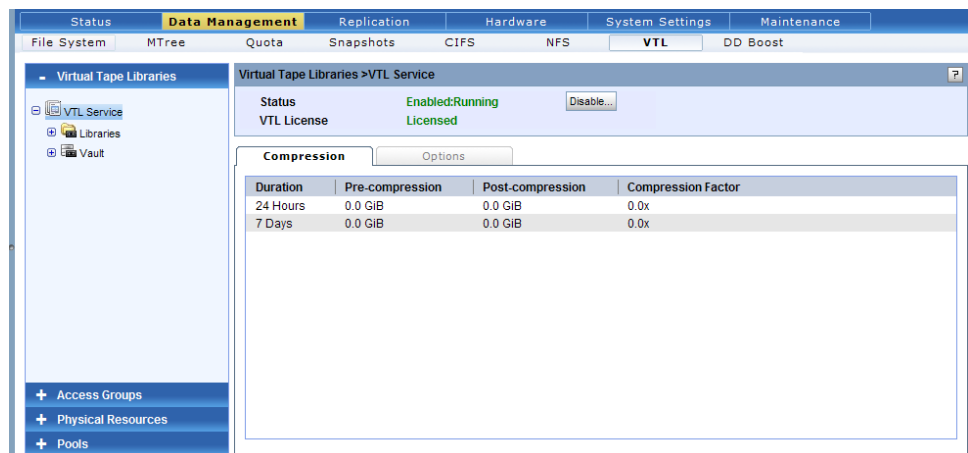
Tape Format	LTO-4	LTO-3	LTO-2	LTO-1
LTO-3	RW	RW	—	—
LTO-2	R	RW	RW	—
LTO-1	—	R	RW	RW

## About the DD System Manager VTL View

Here is how to get to the VTL view in the EMC Data Domain System Manager (DD System Manager).

### Procedure

1. Select a system from the Navigation Panel.
2. Select Data Management > VTL.



### Results

The VTL view provides the following areas for working with tape storage, which are accessed by selecting a Stack menu button in the left frame:

- ◆ Virtual Tape Libraries
- ◆ Access Groups
- ◆ Physical Resources
- ◆ Pools

The Stack menu is a stack of individual menus; selecting a button brings it to the top of the stack and reveals its elements. Selecting an element (for example, a library within the Virtual Tape Libraries or a drive within that library) displays, in the right panel, content specific to the element.

The More Tasks menu (when it is available) lists operations that can be performed on the object selected in the Stack menu. Available operations differ, depending on the selected item.

## Setting Up a VTL

To set up a simple VTL (Virtual Tape Library), use the Configuration Wizard. For more detail, see the *EMC DD OS Initial Configuration Guide*.

Then, follow these procedures:

- ◆ [Enabling VTL on page 256](#)
- ◆ [Creating Libraries on page 258](#)
- ◆ [Creating Tapes on page 270](#)
- ◆ [Importing Tapes on page 265](#)

## Working with the VTL Service Operations

In the stack menu, select **Virtual Tape Libraries** > **VTL Service**.

In this area, you can perform the following basic VTL operations:

- ◆ [Viewing the VTL Service Information Panel on page 255](#)
- ◆ [Enabling VTL on page 256](#)
- ◆ [Disabling VTL on page 256](#)
- ◆ [Configuring VTL Options on page 257](#)
- ◆ [Working with Libraries on page 257](#)
- ◆ [Working with the Vault on page 274](#)

### Viewing the VTL Service Information Panel

The Virtual Tape Libraries > VTL Service page includes the following.

#### VTL License

If the VTL license has not been applied, select **Add License**. Enter the license key in the Add License Key dialog box. Select **Next**, and select **OK** to close the dialog box after the license has been added.

#### Operational Status

At the top left of the Information Panel, a two-part status code displays, for example, `Enabled:Running` shown in a color coded to the status.

The license status is also displayed. `Licensed` is color-coded in green, and `Not Licensed` is color-coded in red. If not licensed, a link labeled **Add License** displays to allow adding the license directly from this area.

The first part of the status code can be Enabled (On) or Disabled (Off). The possible states after the colon are:

State	Description
Running	The VTL process is enabled and active. The status color is green.
Starting	The VTL process is starting.
Stopping	The VTL process is being shut down.
Stopped	The VTL process is disabled. The status color is red.
Timing out	The VTL process crashed and is attempting an automatic restart.

State	Description
Stuck	After a number of VTL process automatic restarts fail, the process is not able to shut down normally and attempts to kill the failed process.

### Options Tab

The Options tab displays the following information (see [Configuring VTL Options on page 257](#) to set options):

Item	Description
Property	Configured options. For example: <ul style="list-style-type: none"> <li>• auto-eject</li> <li>• auto-offline</li> <li>• I/OS License (for IBM i customers)</li> </ul>
Value	Value of the configured option. For example: <ul style="list-style-type: none"> <li>• auto-eject – The state, either enabled or disabled.</li> <li>• auto-offline – The state, either enabled or disabled. When enabled, automatically takes a drive offline before a tape move operation is performed.</li> <li>• I/OS License – Displays license. Select Add License to add a license for I/OS.</li> </ul>

## Enabling VTL

Here is how to enable Fibre Channel services and VTL – and to enable all libraries and library drives:

### Procedure

1. Select VTL, and select the **Virtual Tape Libraries** menu at the far left.
2. Select **VTL Services** from the Virtual Tape Libraries list.
3. In the Status area of the Information Panel, select **Enable**.

When VTL is enabled, the status in the Information Panel displays **Enabled: Running** in green text. (See [Viewing the VTL Service Information Panel on page 255](#) for details.) Options configured are shown in the Options tab.

## Disabling VTL

Here is how to disable VTL and shut down the VTL operation:

### Procedure

1. Select the VTL tab, and select the **Virtual Tape Libraries** menu on the far left.
2. Select **VTL Service** from the Virtual Tape Libraries list.
3. In the Status area of the Information Panel, select **Disable**.

The status in the Information Panel changes to **Disabled: Stopped** in red text.



## Configuring VTL Options

VTL configuration options include enabling/disabling auto-eject and auto-offline.

Enabling auto-eject causes any tape put into a CAP (cartridge access port) to automatically move to the virtual vault, unless:

- ◆ the tape came from the vault, in which case the tape stays in the CAP.
- ◆ an `ALLOW_MEDIUM_REMOVAL` command with a 0 value (false) has been issued to the library to prevent the removal of the medium from the CAP to the outside world.

Enabling auto-offline automatically takes a drive offline before a tape move operation is performed.

---

### Note

Customers of IBM i must enter a valid I/OS license in either of these formats: `xxxx-xxxx-xxxx-xxxx-xxxx` or `xxxx-xxxx-xxxx-xxxx-xxxx-xxxx`. This I/OS license must be installed before creating the library and drives to be used on the IBM i system.

---

Here is how to add a license and configure VTL options.

### Procedure

1. Select **Add License** and enter a valid license key number.
  2. Select **Next**.
  3. Select **Configure** to display the Configure Option dialog.
    - a. In the auto-eject menu, select **Enable**.
    - b. In the auto-offline menu, select **Enable**.
    - c. Select **OK**.
- 

### Note

To disable all of these options, select **Reset to Default** in the Configure Option dialog.

---

## Working with Libraries

In the menu at the far left, expand Virtual Tape Libraries > VTL Service > Libraries.

From the Libraries page, you can perform the following functions:

- ◆ [Viewing the Libraries Information Panel on page 257](#)
- ◆ [Creating Libraries on page 258](#)
- ◆ [Deleting Libraries on page 259](#)
- ◆ [Searching for Tapes on page 269](#)

### Viewing the Libraries Information Panel

The Virtual Tape Libraries > VTL Service > Libraries page displays the following information:

Item	Description
Name	Name of a configured library.
Drives	Number of drives configured in the library.

Item	Description
Slots	Number of slots configured in the library.
CAPs	Number of CAPs (cartridge access ports) configured in the library.

## Creating Libraries

VTL supports a maximum of 64 libraries per system (that is, 64 concurrently active virtual tape library instances on each Data Domain system).

Here is how to create a library.

### Procedure

1. From the More Tasks menu, select **Library** > **Create** to display the Create Library dialog.
2. Enter the following information:

Field	User Input
Library Name	Enter a name from 1 to 32 alphanumeric characters.
Number of Drives	See <a href="#">Number of Supported Tape Drives on page 252</a>
Drive Model	Select the desired model from the drop-down list: <ul style="list-style-type: none"> <li>• IBM-LTO-1</li> <li>• IBM-LTO-2</li> <li>• IBM-LTO-3</li> <li>• IBM-LTO-4</li> <li>• HP-LTO-3</li> <li>• HP-LTO-4</li> </ul>
Number of Slots	Enter the number of slots in the library: <ul style="list-style-type: none"> <li>• Up to 32,000 slots per library</li> <li>• Up to 64,000 slots per system</li> <li>• This should be equal to or greater than the number of drives.</li> </ul>
Number of CAPs	(Optional) Enter the number of cartridge access ports (CAPs): <ul style="list-style-type: none"> <li>• Up to 100 CAPs per library</li> <li>• Up to 1000 CAPs per system</li> </ul>
Changer Model Name	Select the desired model from the drop-down list: <ul style="list-style-type: none"> <li>• L180</li> <li>• RESTORER-L180</li> <li>• TS3500</li> <li>• I2000</li> </ul> <p>Check the backup software application documentation on the Data Domain support site for the model name that you should use.</p>

### 3. Select **OK**.

After the Create Library status dialog shows *Completed*, select **OK**.

The new library appears under the Libraries icon in the VTL Service tree, and the options you have configured appear as icons under the library. Selecting the library displays details about the library in the Information Panel.

#### Note

Access to VTLs and tape drives is managed with the Access Group feature. See [Working with Access Groups on page 283](#).

## Deleting Libraries

Here is how to delete a library from a VTL.

**Procedure**

1. In the More Tasks menu, select **Library** > **Delete** to display the Delete Libraries dialog.
2. Select or confirm the checkbox of the items to delete:
  - The name of each library, or
  - Library Names, to delete all libraries
3. Select **Next**.
4. Verify the libraries to delete, and select Submit in the confirmation dialogs.
5. After the Delete Libraries Status dialog shows *Completed*, select **Close**.

The selected libraries are deleted from the VTL.

---

**Note**

If there are any tapes in a library when it is deleted, they are moved to the vault.

---

## Working with a Library

**Procedure**

1. In the menu at the far left, select **Virtual Tape Libraries** > **VTL Service** > **Libraries**.
2. Select the icon of a specific library.

**Results**

From the library page, the available tasks include:

- ◆ [Viewing the Library Information Panel on page 260](#)
- ◆ [Viewing Changer Information on page 261](#)
- ◆ [Deleting Libraries on page 259](#)
- ◆ [Creating Tapes on page 270](#)
- ◆ [Deleting Tapes on page 271](#)
- ◆ [Importing Tapes on page 265](#)
- ◆ [Exporting Tapes on page 266](#)
- ◆ [Moving Tapes on page 267](#)
- ◆ [Adding Slots on page 273](#)
- ◆ [Deleting Slots on page 273](#)
- ◆ [Adding CAPs on page 274](#)
- ◆ [Deleting CAPs on page 274](#)

## Viewing the Library Information Panel

The Virtual Tape Libraries > VTL Service > Libraries / *library* page displays detailed information about the library named *library*:

**Devices**

Item	Description
Device	Elements in the library, such a drives, slots, and CAPs (cartridge access ports).
Loaded	Number of devices with media loaded.

Item	Description
Empty	Number of devices with no media loaded.
Total	Total of loaded and empty devices.

### Tapes

Item	Description
Pool	Name of the pool where tapes are located.
Tape Count	Number of tapes in the pool.
Capacity	Total configured data capacity of the tapes in that pool, in GiB (Gibibytes, the base-2 equivalent of GB, Gigabytes).
Used	Amount of space used on the virtual tapes in that pool.
Average Compression	Average amount of compression achieved on the data on the tapes in that pool.

## Viewing Changer Information

### Procedure

1. From the stack menu, select **Virtual Tape Libraries** > **VTL Service** > **Libraries** .
2. Select the name of a specific library.
3. Select the library Add (+) button to open the library, and select a Changer element to display the Changer Information Panel, containing the following information:

Item	Description
Vendor	Name of the vendor who manufactured the changer
Product	Model name
Revision	Revision level
Serial Number	Changer serial number

## Working with Tape Drives

### Procedure

1. In the stack menu, select **Virtual Tape Libraries** > **VTL Service** > **Libraries**.
2. Select the name of a specific library.
3. Select the library Add (+) button to open the library, and select the Drives icon.

### Results

From the Drives page, the available tasks include:

- ◆ [Viewing Drives Information on page 262](#)
- ◆ [Creating Tape Drives on page 262](#)
- ◆ [Removing Tape Drives on page 263](#)

To work with tape drives, you must use the tape and library drivers supplied by your backup software vendor that support the IBM LTO-1, IBM LTO-2, IBM LTO-3 (default), IBM

LTO-4, HP-LTO-3, or HP-LTO-4 drives and the IBM TS3500, StorageTek L180, RESTORER-L180, or I2000 libraries.

For more information, see the *Application Compatibility Matrices and Integration Guides* for your vendors.

Because the Data Domain system treats the LTO drives as virtual drives, you can set a maximum capacity to 4 TiB (4000 GiB) for each drive type.

The default capacities for each LTO drive type are as follows:

- ◆ LTO-1 drive: 100 GB
- ◆ LTO-2 drive: 200 GB
- ◆ LTO-3 drive: 400 GB
- ◆ LTO-4 drive: 800 GB

When configuring tape drives, keep in mind the limits on backup streams, which are determined by the platform in use. See [Data Streams Sent to a Data Domain System on page 131](#) for details.

### Tape Full: Early Warning

You receive a warning when the remaining tape space is almost completely full, that is, greater than 99.9, but less than 100 percent. The application can continue writing until the end of the tape to reach 100 percent capacity. The last write, however, is not recoverable.

### Viewing Drives Information

The Virtual Tape Libraries > VTL Service > Libraries > *library* > Drives page displays the following information for all drives in the library named *library*, or if you select a drive from the list, information about the selected drive, as follows.

Column	Description
Drive	List of drives by name, where name is “Drive #” and # is a number between 1 and n representing the address or location of the drive in the list of drives.
Vendor	Manufacturer or vendor of the drive, for example, IBM.
Product	Product name of the drive, for example, ULTRIUM-TD1.
Revision	Revision number of the drive product, for example, 4561.
Serial Number	Serial number of the drive product, for example, 6666660001.
Status	Empty, Open, Locked, or Loaded, depending on the state of the drive. A tape must be present for the drive to be locked or loaded.
Tape	Barcode of the tape in the drive (if any).
Pool	Pool of the tape in the drive (if any).

### Creating Tape Drives

Before creating tape drives, review the information in [Number of Supported Tape Drives on page 252](#).

You cannot mix drive types (such as LTO-1 and LTO-2) in the same library.

Here is how to create tape drives.

**Procedure**

1. From the More Tasks menu, select **Drives Create** to display the Create Drive dialog.
2. Enter information about the drives being added:

Field	User Input
Location	Select a library name, or leave the name as selected.
Number of Drives	Enter the number of drives.
Model Name	Select the model from the drop-down list (and it must be the same as the existing model in the library): <ul style="list-style-type: none"> <li>• IBM-LTO-1</li> <li>• IBM-LTO-2</li> <li>• IBM-LTO-3</li> <li>• IBM-LTO-4</li> <li>• HP-LTO-3</li> <li>• HP-LTO-4</li> </ul>
<p><b>Note</b></p> <p>If a drive already exists, this option is inactive, and the existing drive type is used.</p>	

3. Select **OK** and when the status shows *Completed*, select **OK**.

The added drive entry appears in the Drives list.

**Removing Tape Drives**

Here is how to remove tape drives.

**Procedure**

1. From the More Tasks menu, select **Drives Delete** to display the Delete Drives dialog.

**Note**

If a tape is in the drive, you are prompted to remove the tape.

2. Select the checkboxes of the drives to delete, or select the **Drive** checkbox to delete all drives.
3. Select **Next**, and after verifying drive deletion, select **Submit**.
4. After the Delete Drives Status dialog box shows *Completed*, select **Close**.

The drive entry is removed from the Drives list.

**Migrating LTO-1 to Other Generation Types**

You can migrate tapes from existing LTO-1 type VTLs to VTLs that include other supported LTO-type tapes and drives.

The migration options are different for each backup application, so follow the instructions in the application-specific LTO migration guides on the Data Domain support portal.

**Accessing LTO Migration Guides**

Here is how to access LTO Migration Guides.

**Procedure**

1. Go to the EMC Support website (<https://support.emc.com>).
2. In the search text box, type in **LTO Tape Migration for VTLs**.
3. Read the generic LTO tape migration information, and then select the appropriate document for your particular application.

**Working with Tapes**

When tapes are created, they are placed into the vault. After they have been added to the vault, they can be imported, exported, moved, searched, and removed.

**Procedure**

1. In the stack menu, select **Virtual Tape Libraries > VTL Service > Libraries**.
2. Select the name of a specific library.
3. Select the **Tapes** icon.

**Results**

From the Tapes page, the available tasks include:

- ◆ [Viewing Tape Information on page 264](#)
- ◆ [Importing Tapes on page 265](#)
- ◆ [Exporting Tapes on page 266](#)
- ◆ [Moving Tapes on page 267](#)
- ◆ [Searching for Tapes on page 269](#)
- ◆ [Changing a Tape's Write or Retention Lock State on page 270](#)
- ◆ [Creating Tapes on page 270](#)
- ◆ [Deleting Tapes on page 271](#)
- ◆ [Copying Tapes Between Pools on page 272](#)

**Viewing Tape Information**

The Virtual Tape Libraries > VTL Service > Libraries > *library* > Tapes page displays the following information for the library named *library*:

Item	Description
Barcode	The unique barcode for the tape.
Pool	The name of the pool that holds the tape. The default pool holds all tapes unassigned to a user-created pool.
Location	The location of the tape - whether in a library (and which drive, CAP, or slot number) or in the virtual vault.
State	The state of the tape: <ul style="list-style-type: none"> <li>• RW – Read-writable</li> <li>• RL – Retention-locked</li> <li>• RO – Readable only</li> <li>• WP – Write-protected</li> <li>• RD – Replication destination</li> </ul>
Capacity	The total capacity of the tape.



Item	Description
Used	The amount of space used on the tape.
Compression	The amount of compression performed on the data on a tape.
Last Modified	The date of the last change to the tape's information.
	<p><b>Note</b></p> <p>Modification times used by the system for age-based policies might differ from the last modified time displayed in the tape information sections of the Data Domain System Manager.</p>
Locked Until	If a Retention Lock deadline has been set, the time set is shown. If no retention lock exists, this value is <code>Not specified</code> .

### Importing Tapes

Importing moves existing tapes from the vault to a library slot, drive, or cartridge access port (CAP).

The number of tapes you can import at one time is limited by the number of empty slots in the library. (You cannot import more tapes than the number of currently empty slots.)

- ◆ To view the available slots for a library, select the library from the stack menu. The Information Panel for the library shows the count in the Empty column.
- ◆ If a tape is in a drive, and the tape origin is known to be a slot, the slot is reserved.
- ◆ If a tape is in a drive, and the tape origin is unknown (slot or CAP), a slot is reserved.
- ◆ A tape known to have come from a CAP, and in a drive, does not get a reserved slot. (The tape returns to the CAP when removed from the drive.)
- ◆ To move a tape to a drive, see the procedure [Moving Tapes on page 267](#).

To import tapes:

#### Procedure

1. In the Tapes view, either:
  - a. Enter search information about the tapes to import, and select **Search**:

Field	User Input
Location	Select the location of the tape, for example, <b>Vault</b> .
	<p><b>Note</b></p> <p>Only tapes with the location Vault selected will be imported.</p>
Pool	Select the name of the pool where the tapes reside. If no pools have been created, use the default pool.
Barcode	<ol style="list-style-type: none"> <li>1 Leave the default (*) selected to search for a group of tapes.</li> <li>1 Specify a specific barcode to search for, and only that tape is imported.</li> </ol>

Field	User Input
	<ul style="list-style-type: none"> <li>  Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.</li> </ul>
Count	<p>Enter the maximum number of tapes the search can find.</p> <ul style="list-style-type: none"> <li>  Enter a specific maximum value</li> <li>  Leave blank to find all matching tapes (the Barcode group default (*) is used)</li> </ul>
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.
Items Selected	The number of tapes selected across multiple pages – updated automatically for each tape selection.

**Note**

Based on the previous conditions, a default set of tapes is searched to select the tapes to import. If pool, barcode, or count is changed, select Search to update the set of tapes available to choose from.

- b. Select tapes to import by selecting the checkbox next to:
  - An individual tape,
  - The Barcode column to select all tapes on the current page, or
  - The **Select All Pages** checkbox to select all tapes returned by the search query.

**Note**

Only tapes showing Vault in the Location will be imported.

- c. Select **Import from Vault**. This button is disabled by default and enabled only if the all selected tapes are from the Vault.
- 2. From the Import Tapes: library view, verify the summary information and the tape list, and select **OK**.
- 3. Select **Close** in the status window.

**Exporting Tapes**

Here is how to export tapes from a library to the vault.

**Procedure**

- 1. In the Tapes view, either:
  - a. Select the tape(s) from the list, and select **Export from Vault**.
  - b. Enter search information about the tapes to export, and select Search:

Field	User Input
Location	Select the name of the library where the tape is located.

Field	User Input
Pool	Select the name of the pool to which the tape is to be exported. If no pools have been created, use the default pool.
Barcode	<ul style="list-style-type: none"> <li>I The default (*) selected for a group of tapes.</li> <li>I Specify a specific barcode to search for, and only that tape is exported.</li> <li>I Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.</li> </ul>
Count	<p>The maximum number of tapes the search can find.</p> <ul style="list-style-type: none"> <li>I Enter a specific maximum value, or</li> <li>I Leave blank and use the Barcode group default (*).</li> </ul>
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.
Select All Pages	Select the <b>Select All Pages</b> checkbox to select all tapes returned by the search query.
Items Selected	The number of tapes that are selected across multiple pages – updated automatically for each tape selection.

- c. Select tapes to export: an individual tape; all tapes on the current page (select the Barcode column); or all tapes returned by the search query (select **Select All Pages**).

---

#### Note

Only tapes with a library name in the Location column are exported.

---

- d. Select **Export from Library**.
2. From the Export Tapes: library view, verify the summary information and the tape list, and select **OK**.
  3. Select **Close** in the status window.

## Moving Tapes

One or more tapes can be moved between physical devices from within a library, or several tapes can be moved between pools.

### Moving Tapes Between Devices within a Library

Tapes can be moved between physical devices within a library to mimic backup software procedures for physical tape libraries (which move a tape in a library from a slot to a drive, a slot to a CAP, a CAP to a drive, and the reverse).

In a physical tape library, backup software never moves a tape outside the library. Therefore, the destination library cannot change and is shown only for clarification.

#### Procedure

1. In the More Tasks menu, select **Tapes Move** to display the Move Tapes dialog.

**Note**

When started from a library, the Tapes Panel allows tapes to be moved only between devices.

2. Enter information to search for the tapes to move, and select Search:

Field	User Input
Location	Location cannot be changed.
Pool	N/A
Barcode	<ul style="list-style-type: none"> <li>• Leave the default (*) selected to search among a group of tapes, or</li> <li>• Specify a unique barcode.</li> </ul>
Count	Maximum number of tapes the search can find.
Tapes Per Page	Number of tape entries to display per page.
Items Selected	Number of tapes selected across multiple pages – updated automatically for each tape selection.

3. From the search results list, select the tape or tapes to move.
4. Do one of the following:
  - a. Select the device from the Device list (for example, a slot, drive, or CAP), and enter a starting address using sequential numbers for the second and subsequent tapes (slot address 1-32000, drive address 1-540, and CAP address 1-100). For each tape to be moved, if the specified address is occupied, the next available address is used.
  - b. Leave the address blank if the tape in a drive originally came from a slot and is to be returned to that slot; or if the tape is to be moved to the next available slot.
5. Select **Next**.
6. From the Move Tapes view, verify the summary information and the tape listing, and select **Submit**.
7. Select **Close** in the status window.

**Moving Tapes Between Pools**

Tapes can be moved between pools to accommodate replication activities. For example, pools are needed if all tapes were created in the Default pool, but you later need independent groups for replicating groups of tapes.

You can create named pools and re-organize the groups of tapes into new pools. To move tapes between pools, the tapes must be in the vault.

You cannot move tapes from a tape pool that is a directory replication source. As a workaround, you can:

- ◆ Copy the tape to a new pool, then delete the tape from the old pool.
- ◆ Use an MTree pool, which allows you to move tapes from a tape pool that is a directory replication source.

Here is how to move tapes between pools.

## Procedure

1. In the More Tasks menu, select **Tapes Move** to display the Move Tapes dialog.

### Note

When started from a pool, the Tapes Panel allows tapes to be moved only between pools.

2. Enter information to search for the tapes to move, and select **Search**:

Field	User Input
Location	Location cannot be changed.
Pool	To move tapes between pools, select the name of the pool where the tapes currently reside. If no pools have been created, use the default pool.
Barcode	<ul style="list-style-type: none"> <li>• Leave the default (*) selected to search for a group of tapes, or</li> <li>• Specify a specific barcode, and only that tape will be found.</li> </ul>
Count	Maximum number of tapes the search can find.
Tapes Per Page	Maximum number of tapes to display per page – possible values are 15, 30, and 45.
Items Selected	Number of tapes selected across multiple pages.

3. From the search results list, select the tapes to move.
4. From the Select Destination: Location list, select the location of the pool to which tapes are to be moved. This option is available only when started from the (named) Pool view.
5. Select **Next**.
6. From the Move Tapes view, verify the summary information and tape list, and select **Submit**.
7. Select **Close** in the status window.

## Searching for Tapes

Here is how to search for tapes.

### Procedure

1. In the VTL Navigation tree, select the area to search (library, vault, and Vault > Any Pool).
2. From the More Tasks menu, select **Tapes Search** to display the Search Tapes dialog.
3. Enter information about the tapes you want to find.

Field	User Input
Location	Select the location, or leave the default library selection.
Pool	Select the name of the pool in which to search for the tape. If no pools have been created, use the default pool.

Field	User Input
Barcode	<ul style="list-style-type: none"> <li>• Leave the default (*) selected to search for a group of tapes, or</li> <li>• Specify a unique barcode.</li> </ul>
Count	<p>The maximum number of tapes the search can find.</p> <ul style="list-style-type: none"> <li>• Enter a specific maximum value, or</li> <li>• Leave blank and use the Barcode group default (*).</li> </ul>

#### 4. Select **Search**.

## Changing a Tape's Write or Retention Lock State

Before changing a tape's write or retention lock state, the tape must have been created and imported. VTL tapes follow the standard Data Domain Retention Lock policy. After the retention period for a tape has expired, it cannot be written to or changed (however, it can be deleted).

### Procedure

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries** > *library* > **Tapes** where *library* is the name of the tape library.
2. In the Tapes page, select the tape to modify from the list, and select **Set State** (above the list).
3. In the Set Tape State dialog, select read-writeable, write-protected, or retention-lock.
4. If the state is retention-lock, either
  - enter the tape's expiration date in a specified number of days, weeks, months, years, or
  - select the calendar icon, and select a date from the calendar. The retention lock expires at noon on the selected date.
5. Select **Next**, and select **Submit** to change the state.

## Creating Tapes

### Note

This procedure can be performed from either a library or a pool. If initiated from a library, the system first creates the tapes, then imports them to the library.

Although the number of supported tapes is unlimited, you can create no more than 100,000 tapes at a time.

Here is how to create tapes in a specified pool, then import them to the current library.

### Procedure

1. In the More Tasks menu, select **Tapes Create** to display the Create Tapes dialog.
2. Enter information about the tape:

Field	User Input
Location	If a drop-down menu is enabled, select the library or leave the default selection.

Field	User Input
Pool Name	Select the name of the pool, from the drop-down list, where the tape will reside. If no pools have been created, use the default pool.
Number of Tapes	Select from 1 to 100,000 tapes.
Starting Barcode	Enter the initial barcode number (using the format A99000LA, for example). See <a href="#">About Tape Barcodes on page 252</a> .
Tape Capacity	(optional) Specify the number of GBs from 1 to 4000 for each tape (this setting overrides the barcode capacity setting). For efficient use of disk space, use 100 GB or less.

3. Select **OK** and **Close**.

## Deleting Tapes

### Note

This procedure can be performed from both a library and a pool. If initiated from a library, it will first export the tapes, then delete them.

To remove one or more tapes from the vault and delete all of the data in the tapes, use the **Tapes Delete** option. The tapes must be in the vault, not in a library.

On a Replication destination Data Domain system, deleting a tape is not permitted.

Here is how to delete tapes from the vault.

### Procedure

1. In the More Tasks menu, select **Tapes Delete** to display the Delete Tapes dialog.
2. Enter information about the tape to delete:

Field	User Input
Location	Select a library or leave the default <b>Vault</b> selection.
Pool	Select the name of the pool from which to delete the tape. If no pools have been created, use the default pool.
Barcode	<ul style="list-style-type: none"> <li>• Leave the default (*) selected to search for a group of tapes, or</li> <li>• Specify a specific barcode to search for.</li> <li>• Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.</li> </ul>
Count	Enter the number of tapes to delete. <ul style="list-style-type: none"> <li>• Enter a specific maximum value, or</li> <li>• Leave blank and use the Barcode group default (*).</li> </ul>
Tapes Per Page	Select the maximum number of tapes to display per page – possible values are 15, 30, and 45.

Field	User Input
Select All Pages	Select the <b>Select All Pages</b> checkbox to select all tapes returned by the search query.
Items Selected	The number of tapes that are selected across multiple pages – updated automatically for each tape selection.

3. Select the checkbox of the tape that should be deleted or the checkbox on the heading column to delete all tapes, and select **Next**.
4. Select **Submit** in the confirmation window, and select **Close**.

After a tape is removed, the physical disk space used for the tape is not reclaimed until after a file system cleaning operation.

### Copying Tapes Between Pools

Tapes can be copied between pools, or from the vault to a pool, to accommodate replication activities. This option is available only when invoked from the (named) Pool view.

#### Procedure

1. From the More Tasks menu, select **Tapes Copy** to display the Copy Tapes dialog.
2. Select the checkboxes of tapes to copy, or enter information to search for the tapes to copy, and select Search.

Field	User Input
Location	Select either a library or the <b>Vault</b> for locating the tape.
	<p><b>Note</b></p> <p>While tapes always show up in a pool (under the Pools menu), they are technically in either a library or the vault, but not both, and they are never in two libraries at the same time. Use the import/export options to move tapes between the vault and a library.</p>
Pool	To copy tapes between pools, select the name of the pool where the tapes currently reside. If no pools have been created, use the <b>Default</b> pool.
Barcode	<ul style="list-style-type: none"> <li>• Leave the default (*) selected to search for a group of tapes, or</li> <li>• Specify a unique barcode to search for a single barcode or use the wildcard * or ? to search for a set of barcodes.</li> </ul>
Count	Select the maximum number of tapes the search can find.
Tapes Per Page	Select the maximum number of tapes to display per page – possible values are 15, 30, and 45.
Items Selected	The number of tapes selected across multiple pages – updated automatically for each tape selection.

3. From the search results list, select the tapes to copy.



4. From the Select Destination: Pool list, select the pool where tapes are to be copied.

---

#### Note

If a tape with a matching barcode already resides in the destination pool, an error is displayed, and the copy aborts.

---

5. Select **Next**.
6. From the Copy Tapes Between Pools dialog, verify the summary information and the tape list, and select **Submit**.
7. Select **Close** on the Copy Tapes Between Pools Status window.

## Working with Tape Slots and CAPs

You can add and delete slots and CAPs (cartridge access ports) from a configured library to change the number of storage elements.

---

#### Note

Some backup applications do not automatically recognize that drives, slots, or CAPs have been added to a VTL. For example, when a tape drive is added to a VTL, the administrator may need to remove the VTL from the application and then add it back in before the tape drive can be detected by the application. See the application documentation for information on how to configure the application to recognize changes.

---

#### Procedure

1. Select **Virtual Tape Libraries > VTL Service > Libraries**.
2. Select a specific library.

#### Results

Tasks for working with slots and CAPs include:

- ◆ [Adding Slots on page 273](#)
- ◆ [Deleting Slots on page 273](#)
- ◆ [Adding CAPs on page 274](#)
- ◆ [Deleting CAPs on page 274](#)

### Adding Slots

The total number of slots in a library, or all libraries on a system, cannot exceed 32,000 for a library and 64,000 for a system.

#### Procedure

1. From the More Tasks menu, select **Slots Add**.
2. In Number of Slots on the Add Slots dialog, enter the number of slots to add.
3. Select **OK** and **Close** when the status shows *Completed*.

### Deleting Slots

Before deleting any slots containing tape cartridges, move those cartridges to the vault. The system will delete only empty, uncommitted slots.

From the More Tasks menu, select **Slots Delete** to display the Delete Slots dialog.

### Procedure

1. In **Number of Slots**, type the number of slots to delete – you can delete from 1 to 32,000 slots.
2. Select **OK** and **Close** when the status shows *Completed*.

## Adding CAPs

---

### Note

CAPs (cartridge access ports) are used by a limited number of backup applications. See your backup application documentation to ensure that CAPs are supported.

---

The total number of CAPs cannot exceed 100 per library or 1000 per system.

### Procedure

1. From the **More Tasks** menu, select **CAPs Add** to display the Add CAPs dialog.
2. In **Number of CAPs**, type the number of CAPs to add – you can add from 1 to 100 CAPs per library and 1,000 CAPs per system.
3. Select **OK** and **Close** when the status shows *Completed*.

## Deleting CAPs

---

### Note

If tape cartridges are loaded in CAPs to be deleted, the cartridges will be moved to the vault.

---

### Procedure

1. From the **More Tasks** menu, select **CAPs Delete** to display the Delete CAPs dialog.
2. In **Number of CAPs**, type the number of CAPs to delete – you can delete a maximum of 100 CAPs per library and 1000 CAPs per system.
3. Select **OK** and **Close** when the status shows *Completed*.

## Working with the Vault

Select **Virtual Tape Libraries** > **VTL Service** > **Vault**.

From the vault page, tasks available include:

- ◆ [Viewing Vault Information on page 274](#)
- ◆ [Creating Storage Pools on page 277](#)
- ◆ [Deleting Storage Pools on page 279](#)
- ◆ [Creating Tapes on page 270](#)
- ◆ [Deleting Tapes on page 271](#)
- ◆ [Searching for Tapes on page 269](#)

## Viewing Vault Information

The **Virtual Tape Libraries** > **VTL Service** > **Vault** page provides the following information for the Default pool and any other existing pools:

Item	Description
Location	Name of the pool.
Type	Directory or MTree.
Tape Count	Number of tapes in the pool.
Capacity	Total amount of space in the pool.
Used	Amount of space used on in the pool.
Average Compression	Average amount of compression in the pool.

## Working with Vault Pools

The vault contains storage pools that can be replicated.

To access a vault pool, select **Virtual Tape Libraries > VTL Service > Vault > pool**. Notice that pool “Default” always exists.

From a vault pool, the following tasks can be performed:

- ◆ [Viewing Vault Pool Information on page 275](#)
- ◆ [Creating Tapes on page 270](#)
- ◆ [Deleting Tapes on page 271](#)
- ◆ [Moving Tapes on page 267](#)
- ◆ [Copying Tapes Between Pools on page 272](#)
- ◆ [Searching for Tapes on page 269](#)

### Viewing Vault Pool Information

The Virtual Tape Libraries > VTL Service > Vault > *pool* page includes two tabs.

#### Pool Tab

The Pool tab contains the following information:

Item	Description
Convert to MTree Pool	Select this to convert directory pool to MTree pool.
Type	Directory or MTree.
Tape Count	Number of tapes in pool.
Capacity	Total configured data capacity of tapes in pool, in GiB (Gibibytes, base-2 equivalent of GB, Gigabytes).
Used	Amount of space used on virtual tapes in pool.
Average Compression	Average amount of compression achieved for data on tapes in pool.

#### Replication Tab

The Replication tab presents the following replication information:

Item	Description
Name	Name of storage pool.

Item	Description
Configured	Shows whether replication is configured for pool: yes or no.
Source	Path to source pool to be replicated.
Destination	Path to destination where pool will be replicated.

## Working with Storage Pools

VTL storage pools allow the replication of pools of VTL tapes. VTL tapes can be replicated from multiple replication sources to a single replication destination (many-to-one configurations).

When using storage pools, consider the following:

- ◆ A storage pool can be replicated no matter where individual tapes are located. Tapes can be in the vault or in a library (slot, cap, or drive).
- ◆ You can copy and move tapes from one storage pool to another.
- ◆ Two tapes in two different storage pools on a Data Domain system can have the same name. In this case, neither tape can be moved to the other tape's storage pool.
- ◆ Storage pools can be of two types: MTree (recommended), or Directory, which is backward-compatible.
- ◆ A storage pool sent to a replication destination must have a name that is unique on the destination.
- ◆ Data Domain storage pools are not accessible by backup software.
- ◆ No VTL configuration or license is needed on a replication destination when replicating storage pools.
- ◆ You must create tapes with unique bar codes. Duplicate bar codes may cause unpredictable behavior in backup applications and can be confusing to users.

From the stack menu, select **Pools > Pools**.

From the Pools page, tasks available include:

- ◆ [Viewing Storage Pools Information on page 276](#)
- ◆ [Creating Storage Pools on page 277](#)
- ◆ [Converting from a Directory to an MTree Replication Pair on page 319](#)
- ◆ [Deleting Storage Pools on page 279](#)
- ◆ [Searching for Tapes on page 269](#)

## Viewing Storage Pools Information

The Pools page includes two tabs.

### Pools Tab

The Pools tab displays the following information:

Item	Description
Location	Location of storage pool.
Type	Directory or MTree.
Tape Count	Number of tapes in pool.

Item	Description
Capacity	Total configured data capacity of tapes in pool, in GiB (Gibibytes base-2 equivalent of GB, Gigabytes).
Used	Amount of space used on virtual tapes in pool.
Average Compression	Average amount of compression achieved for data on tapes in pool.

### Replication Tab

The Replication tab presents the following detailed replication information:

Item	Description
Name	Name of storage pool.
Configured	Shows whether replication is configured for pool: yes or no.
Source	When configured, shows path where pool is replicating from, or both, if a cascaded configuration.
Destination	When configured, shows path where pool is replicating to, or both, if a cascaded configuration.

## Creating Storage Pools

### MTree versus Directory Pool

When you create a storage pool, you are creating an MTree, unless you specifically elect to create the older style directory pool (one that is backward-compatible).

The advantages of using an MTree (as opposed to a directory pool) include the ability to:

- ◆ make individual snapshots and schedule snapshots.
- ◆ apply retention locks.
- ◆ set an individual retention policy.
- ◆ get compression information.
- ◆ get data migration policies to the Retention Tier.
- ◆ establish a storage space usage policy (quota support) by setting hard limits and soft limits.

### Points to Consider

When creating storage pools, consider the following:

- ◆ A storage pool name:
  - cannot be one of the following: “all,” “vault,” or “summary.”
  - must contain between 1 and 32 characters (excluding the characters “ \* / < > ? : \ ”).
  - cannot have a space or period at its beginning or end.
  - is case-sensitive.
- ◆ A storage pool can be replicated no matter where individual tapes are located in a VTL provided by the Data Domain system. Tapes can be in the vault, a library, or a drive.

- ◆ A storage pool sent to a replication destination must have a name that is unique on the destination.

## How to Create a Storage Pool

Here is how to create a storage pool.

### Procedure

1. From the **More Tasks** menu, select **Pool Create** to display the Create Pool dialog.
2. In Pool Name, enter a name that conforms to [Points to Consider on page 277](#).
3. If you want the pool to be a directory and compatible with the previous version of Data Domain System Manager, select the option "Create a directory backwards compatibility mode pool."
4. Select **OK** to display the Create Pool Status dialog.
5. After the Create Pool Status dialog box shows `Completed`, select **Close**.

The storage pool is added in the Pools subtree, and you can now add virtual tapes to it.

## Converting a Directory Pool to an MTree Pool

Here is how to convert a directory pool to an MTree pool.

### WARNING

**Make sure your source and destination pools are synchronized before you convert, so that the number of tapes, and the data on each side, is fully intact before the conversion is performed.**

---

### Procedure

1. In the Pools submenu, select the directory pool you wish to convert.
2. In the Pool tab, select **Convert to MTree Pool**.
3. Select **OK** in the Convert to MTree Pool dialog.

This conversion affects replication in the following ways:

- VTL is temporarily disabled on the replicated systems during pool conversion.
- The destination data is copied to a new pool on the destination system to preserve the data until the new replication is initialized and synced. Afterward, you may safely delete this temporarily copied pool, which is named **CONVERTED-*pool***, where *pool* is the name of the pool that was upgraded (or the first 18 characters for long pool names). (This applies only to DD OS 5.4.1.0 and later.)
- The target replication directory will be converted to MTree format. (This applies only to DD OS 5.2 and later.)
- Replication pairs are broken before pool conversion and re-established afterward if no errors occur.
- If the directory pool was being replicated on multiple systems, those replicating systems must be known to the managing system for the conversion to work properly.
- If the directory pool was being replicated to an older DD OS, such as DD OS 5.2 to DD OS 5.1, this MTree conversion is not allowed. As a workaround:
  - Replicate the directory pool to a second Data Domain system.

- Replicate the directory pool from the second Data Domain system to a third Data Domain system.
- Remove the second and third Data Domain systems from the managing Data Domain system's Data Domain network.
- On any of the systems running DD OS 5.2, from the Pools submenu, select **Pools** and a directory pool. In the Pools tab, select **Convert to MTree Pool**. A message states that the conversion cannot proceed because some replication system information is missing.

Retention Lock cannot be enabled on systems involved in MTree pool conversion.

## Renaming Storage Pools

Here is how to rename a storage pool:

### Procedure

1. In the Pools submenu, select the pool you wish to rename.
2. From the More Tasks menu, select **Pool Rename** to display the Rename Pool dialog.
3. In Pool Name, enter a name that conforms to the naming conventions in [Points to Consider on page 277](#).
4. Select **OK** to display the Rename Pool status dialog.
5. After the Rename Pool status dialog box shows *Completed*, select **Close**.

The pool is renamed in the Pools subtree.

## Deleting Storage Pools

---

### Note

Before a pool can be deleted, you must delete any tapes contained within it. If replication is configured for the pool, the replication pair must also be deleted.

---

Deleting a pool corresponds to renaming the MTree and then deleting it, which occurs at the next cleaning process.

Here is how to delete a storage pool.

### Procedure

1. From the More Tasks menu, select **Pool Delete** to display the Delete Pools dialog.
2. Select the checkbox of items to delete:
  - The name of each pool, or
  - **Pool Names**, to delete all pools.
3. Select **Submit** in the confirmation dialogs.

The selected pools are deleted.

4. After the Delete Pool Status dialog shows *Completed*, select **Close**.

The pool entry is removed from the pool list.

## Replicating Storage Pools

Storage pools can be replicated and monitored using the Replication tab.

- ◆ See [Creating a Directory, MTree, or Pool Replication Pair on page 312](#).

- ◆ See [Tracking Status of a Backup Job's Replication Progress on page 322](#).

## Working with a Single Storage Pool

To access a specific storage pool, select Pools > Pools > *pool*. Notice that pool “Default” always exists.

From a single storage pool page, the following tasks can be performed:

- ◆ [Renaming Storage Pools on page 279](#)
- ◆ [Deleting Storage Pools on page 279](#)
- ◆ [Creating Tapes on page 270](#)
- ◆ [Deleting Tapes on page 271](#)
- ◆ [Moving Tapes on page 267](#)
- ◆ [Searching for Tapes on page 269](#)



# CHAPTER 14

## Working with SCSI Target

This chapter includes:

- ◆ [About SCSI Target](#).....282
- ◆ [Working with Access Groups](#)..... 283
- ◆ [Working with an Access Group](#)..... 287
- ◆ [Working with Physical Resources](#).....289

## About SCSI Target

---

### Note

SCSI (Small Computer System Interface) Target is currently supported for VTL and DD Boost over FC (Fibre Channel) services.

- ◆ See [Working with DD Boost on page 231](#) for DD Boost-related SCSI Target features of the Data Domain System Manager.
  - ◆ See the *EMC Data Domain Boost for OpenStorage Administration Guide* for all other types of information about DD Boost.
  - ◆ See the *EMC DD OS Command Reference Guide* for a description of the SCSI Target (`scsitar`) commands.
- 

SCSI Target starts when FC Ports are present or VTL is licensed. It provides unified management for all SCSI Target *services* and *transports*.

- ◆ A *service* is anything that has a target LUN (logical unit number) on a Data Domain system that uses SCSI commands, such as VTL (tape drives and changers) and DD Boost over FC (processor devices).
- ◆ A *transport* enables devices to become visible to *initiators*. An initiator is a backup client that connects to a system to read and write data using the FC protocol. A specific initiator can support DD Boost over FC or VTL, but not both.

Devices are visible on a SAN (storage area network) through physical ports. Host initiators communicate with the Data Domain system through the SAN. Access groups manage access between devices and initiators.

An endpoint is the logical target on a Data Domain system to which the initiator connects. Endpoints have the following attributes:

- ◆ port topology
- ◆ FCP2-RETRY status
- ◆ WWPN
- ◆ WWNN

You can disable, enable, and rename endpoints. You can also delete endpoints; for example, you can delete endpoints whose associated transport hardware no longer exists. Endpoints are automatically discovered and created when a new transport connection occurs.

SCSI Target functionality is managed in the following ways:

- ◆ For basic management tasks – Data Domain System Manager – for example, to disable and enable endpoints.
- ◆ For more controlled management – the `scsitar` command – for example, to rename and delete endpoints.
- ◆ For specific tasks – service commands, such as `vtl` or `ddboost`.

This chapter focuses on using the Data Domain System Manager.

After you have become familiar with the basic tasks, see the `scsitar` command in the *EMC DD OS Command Reference Guide* for more advanced management tasks.

**Note**

Avoid using the `scsitarget group use` command while under heavy VTL usage.

## Working with Access Groups

An access group is created to hold a collection of initiator WWPNs or aliases and the drives and changers they are allowed to access.

A VTL default group named TapeServer lets you add devices that will support NDMP-based backup applications. See [Configuring the NDMP Device TapeServer Group on page 288](#) for details.

Access group configuration allows initiators (in general backup applications) to read and write data to devices in the same access group.

Access groups allow clients to access only selected LUNs (media changers or virtual tape drives) on a system. A client set up for an access group can access only devices in its access group.

**Note**

Avoid making access group changes on a Data Domain system during active backup or restore jobs. A change may cause an active job to fail. The impact of changes during active jobs depends on a combination of backup software and host configurations.

To work with Access Groups, select **Access Groups > Groups**.

From the Groups page, you can perform the following tasks:

- ◆ [Viewing Access Groups Information on page 283](#)
- ◆ [Configuring an Access Group on page 283](#)
- ◆ [Deleting an Access Group on page 286](#)

## Viewing Access Groups Information

When you select Access Groups > Groups, the following information is displayed:

Item	Description
Group Name	Name of group.
Initiators	Number of initiators in group.
Devices	Number of devices in group.

## Configuring an Access Group

When you create or configure an access group on a Data Domain system, each Data Domain system device (media changer or drive) can be assigned to multiple groups. Devices assigned to other groups, however, cannot be assigned to TapeServer. A maximum of 128 groups can be created.

Here is how to configure or modify an Access Group:

**Procedure**

1. Select **Access Groups**.

The Groups icon should be highlighted, or select an existing group from the list to change the configuration.

2. In the **More Tasks** menu, select **Group Create**.

The Create Group dialog is displayed if the Groups icon is selected. If an existing group is selected, the Configure Group dialog lists the devices that have been configured for the group.

3. In the Group Name text box, enter a name for the group. (This field is required.)

The group name must be a unique name of up to 128 characters, and can contain only the characters 0-9, a-z, A-Z, underscore (\_), and hyphen (-). Group names are not case sensitive. Up to 128 groups can be created.

The names “TapeServer,” “all,” and “summary” are reserved and cannot be used as group names.

4. To configure initiators to the access group, check the box next to the initiator in the Initiators Panel. You can add initiators to the group later (see [About Initiators on page 294](#)).
5. Select **Next**.

The Devices dialog box lists devices that have been configured for the group. The name of the library, devices in the group, LUN, and primary or secondary status is displayed in the table.

6. Select Add (+) to add devices to the Access Group, as described in the steps a-e. In this dialog, you can also modify or delete a set of devices that were previously added. See [Modifying Access Group Devices on page 285](#) and [Deleting Access Group Devices on page 286](#).
  - a. Verify the correct library is selected in the Library Name drop-down list, or select another library.
  - b. In the Device area, select the checkboxes of the devices (changer and drives) to be included in the group.
  - c. Optionally, specify a starting LUN in the LUN Start Address text box.

This is the LUN that the Data Domain system returns to the initiator. Each device is uniquely identified by the library and the device name. (For example, it is possible to have drive 1 in Library 1 and drive 1 in Library 2). Therefore, a LUN is associated with a device, which is identified by its library and device name.

The initiators in the access group interact with the LUN devices that are added to the group.

The maximum LUN accepted when creating an access group is 16383.

A LUN can be used only once for an individual group. The same LUN can be used with multiple groups.

---

#### Note

Some VTL initiators (clients) have specific rules for VTL target LUN numbering; for example, requiring LUN 0 or requiring contiguous LUNs. If these rules are not followed, an initiator may not be able to access some or all of the LUNs assigned to a VTL target port.

---

Check your initiator documentation for special rules, and if necessary, alter the device LUNs on the VTL target port to follow the rules. For example, if an initiator requires LUN 0 to be assigned on the VTL target port, check the LUNs for devices assigned to ports, and if there is no device assigned to LUN 0, change the LUN of a device so it is assigned to LUN 0.

- d. In the Primary and Secondary Endpoints area, select an option to determine from which ports the selected device will be seen. The following conditions apply for designated ports:
- All – The checked device is seen from all ports.
  - None – The checked device is not seen from any port.
  - Select – The checked device is to be seen from selected ports. Select the checkboxes of the appropriate ports.  
If only primary ports are selected, the checked device is visible only from primary ports.  
If only secondary ports are selected, the checked device is visible only from secondary ports. Secondary ports can be used if the primary ports become unavailable.

---

#### Note

The switchover to a secondary port is not an automatic operation. You must manually switch the VTL device to the secondary ports if the primary ports become unavailable. See [Working with Physical Resources on page 289](#).

---

The port list is a list of physical port numbers. A port number denotes the PCI slot and a letter denotes the port on a PCI card. Examples are 1a, 1b, or 2a, 2b.

---

#### Note

A drive appears with the same LUN on all the ports that you have configured.

- e. Select **OK**.
- You are returned to the Devices dialog box where the new group is listed. To add more devices, repeat these five substeps.
7. Select **Next**.
8. Select **Close** when the `Completed` status message displays.

## Modifying Access Group Devices

Here is how to modify devices for an access group.

### Procedure

1. Display the Create Group dialog. See [Create Access Group on page 244](#).
2. Select a device in the group table, and select the edit (pencil) icon to modify devices in the access group, to display the Modify Devices dialog. Then, follow steps a-e.
  - a. Verify that the correct library is selected in the **Library Name** drop-down list, or select another library.
  - b. In the Device to Modify area, select the checkboxes of the devices (changer and drives) to be modified.
  - c. Optionally, modify the starting LUN (logical unit number) in the Starting LUN text box.

This is the LUN that the Data Domain system returns to the initiator. Each device is uniquely identified by the library and the device name. (For example, it is possible to have drive 1 in Library 1 and drive 1 in Library 2). Therefore, a LUN is associated with a device, which is identified by its library and device name.

The initiators in the access group interact with the LUN devices that are added to the group.

The maximum LUN accepted when creating an access group is 16383.

A LUN can be used only once for an individual group. The same LUN can be used with multiple groups.

- d. In the Primary and Secondary Ports area, change the option that determines the ports from which the selected device is seen. The following conditions apply for designated ports:
- All – The checked device is seen from all ports.
  - None – The checked device is not seen from any port.
  - Select – The checked device is seen from selected ports. Select the checkboxes of the ports from which it will be seen.  
If only primary ports are selected, the checked device is visible only from primary ports.  
If only secondary ports are selected, the checked device is visible only from secondary ports. Secondary ports can be used if primary ports become unavailable.

---

#### Note

The switchover to a secondary port is not an automatic operation. You must manually switch the VTL device to the secondary ports if the primary ports become unavailable, see [Working with Physical Resources on page 289](#).

The port list is a list of physical port numbers. A port number denotes the PCI slot, and a letter denotes the port on a PCI card. Examples are 1a, 1b, or 2a, 2b.

---

#### Note

A drive appears with the same LUN on all ports that you have configured.

- e. Select **OK**.

## Deleting Access Group Devices

Here is how to delete devices for an access group.

### Procedure

1. Display the Create Group dialog. See [Create Access Group on page 244](#).
2. Select a device in the group table, and select the remove (X) icon to delete it.

The device is deleted.

## Deleting an Access Group

Before a group can be removed, you must remove the initiators and LUNs from the group. Here is how to remove an access group.

### Procedure

1. Delete devices in the group using the procedure [Configuring an Access Group on page 283](#).
2. In the **More Tasks** menu, select **Group Delete** to display the Delete Group dialog.
3. Select the checkbox of the group to be removed, and select **Next**.
4. In the groups confirmation dialog box, verify the deletion, and select **Submit**.

5. Select **Close** when the Delete Groups Status displays **Completed**.

## Working with an Access Group

To start working with an access group, select an access group in the Access Groups > Groups list.

From the Access Group page, available tasks include:

- ◆ [Viewing Access Group Information on page 287](#)
- ◆ [Configuring an Access Group on page 283](#)
- ◆ [Deleting an Access Group on page 286](#)
- ◆ [Configuring the NDMP Device TapeServer Group on page 288](#)
- ◆ [Selecting Endpoints for a Device on page 289](#)

### Viewing Access Group Information

The Access Groups > Groups > group page includes these tabs:

#### LUNs Tab

The LUNs (logical unit numbers) tab contains the following information:

Item	Description
LUN	Device address, where the maximum number is 16383. A LUN can be used only once within a group, but can be used again within another group. VTL devices added to a group must use contiguous LUNs.
Library	Name of library associated with LUN.
Device	Changers and drives.
In-Use Endpoints	Set of endpoints currently being used: primary or secondary.
Primary Endpoints	Initial (or default) endpoint used by the backup application. In the event of a failure on this endpoint, the secondary endpoints may be used, if available.
Secondary Endpoints	Set of fail-over endpoints to use if a primary endpoint fails. Use the task <a href="#">Working with Physical Resources on page 289</a> to manually fail-over to the secondary endpoints.

#### Initiators Tab

The Initiators tab contains the following information:

Item	Description
Initiator	Name of initiator, which is either WWPN or alias assigned to initiator (see <a href="#">About Initiators on page 294</a> ).
WWPN	Port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel port, and which is worldwide unique.

## Configuring the NDMP Device TapeServer Group

The TapeServer group holds tape drives that interface with NDMP-based backup applications and that send control information and data streams over IP instead of FC.

### Note

A device used by the NDMP TapeServer must be in the VTL group TapeServer. That device is then available only to the NDMP TapeServer.

Here is how to configure the TapeServer group.

### Procedure

1. Add tape drives to a new or existing library (named “dd660-16” in this example), as described in the section [Creating Tape Drives on page 262](#).
2. Create slots and CAPs to the library, as described in [Adding Slots on page 273](#) and [Adding CAPs on page 274](#).
3. Add the created devices in a library (in this example, “dd660-16”) to the TapeServer access group using the procedure in [Configuring an Access Group on page 283](#).
4. Enable NDMPD by entering:

```
# ndmpd enable
Starting NDMP daemon, please wait.....
NDMP daemon is enabled.
```

5. Ensure that the NDMP daemon sees the devices in the TapeServer group:

```
sysadmin@dd660-16# ndmpd show devicenames
NDMP Device          Virtual Name          Vendor   Product          Serial Number
-----
/dev/dd_ch_c0t010    dd660-16 changer     STK      L180              6290820000
/dev/dd_st_c0t110    dd660-16 drive 1     IBM      ULTRIUM-TD3      6290820001
/dev/dd_st_c0t210    dd660-16 drive 2     IBM      ULTRIUM-TD3      6290820002
/dev/dd_st_c0t310    dd660-16 drive 3     IBM      ULTRIUM-TD3      6290820003
/dev/dd_st_c0t410    dd660-16 drive 4     IBM      ULTRIUM-TD3      6290820004
-----
```

6. Add an NDMP user (ndmp in this example) with the following command:

```
sysadmin@dd660-16# ndmpd user add ndmp
Enter password:
Verify password:
```

7. Verify the user ndmp is added correctly:

```
sysadmin@dd660-16# ndmpd user show
ndmp
```

8. Show the NDMP configuration:

```
sysadmin@dd660-16# ndmpd option show all
Name          Value
-----
authentication  text
debug          disabled
port          10000
preferred-ip
-----
```

9. Change the default user password authentication to use MD5 encryption for enhanced security, and verify the change (notice the authentication value changes from text to md5):

```
sysadmin# ndmpd option set authentication md5
sysadmin# ndmpd option show all
Name          Value
```



```

-----
authentication  md5
debug           disabled
port           10000
preferred-ip
-----

```

### Results

NDMP is now configured, and the TapeServer access group shows the device configuration. See the “ndmpd” chapter of the *EMC DD OS Command Reference Guide* for the complete command set and options.

## Working with Physical Resources

Physical Resources lets you work with Endpoints and Initiators, as described in the following sections:

- ◆ [About Endpoints on page 289](#)
- ◆ [About Initiators on page 294](#)
- ◆ [Setting a Loop ID on page 296](#)
- ◆ [FC Link Monitoring on page 296](#)

For more information about the tabs under the Physical Resources area of the EMC Data Domain System Manager, see:

- ◆ [Endpoints Tab of Physical Resources on page 293](#)
- ◆ [Initiators Tab of Physical Resources on page 295](#)
- ◆ Configure Resources link, which opens the Hardware › Fibre Channel › Physical Resources tab where you can modify endpoints and initiators.

## About Endpoints

An endpoint is the logical target on the Data Domain system to which the initiator connects.

This section contains the following topics:

- ◆ [Viewing Endpoints Information on page 291](#)
- ◆ [Viewing Endpoint Information on page 292](#)
- ◆ [Endpoints Tab of Physical Resources on page 293](#)

## Selecting Endpoints for a Device

Here is how to select endpoints for a device.

### Procedure

1. Go to Access Groups › Groups.
2. Select a specific group from the list.
3. From the More Tasks menu, select **Endpoints** › **Set In-Use** to display the Set In-Use Endpoints dialog.
4. Select only specific devices, or select **Devices** to select all devices in the list.
5. Indicate whether the endpoints are primary or secondary.
6. Select **OK**.

## Configuring an Endpoint

Here is how to configure an endpoint.

### Procedure

1. Select **Hardware > Fibre Channel > Physical Resources**. Under **Endpoints**, select an endpoint, and then select **Configure** to display the **Configure Endpoint** dialog.
2. Enter a name for the endpoint (1 to 28 characters). The field cannot be empty or be the word “all,” and cannot contain the characters asterisk (\*), question mark (?), front or back slashes (/), \), or right or left parentheses [(,)].
3. Uncheck **Enabled** next to **FCP2 Retry**, if you do *not* want the endpoint to support retransmission of data during data recovery.
4. Select one of these options for the topology.
  - **Default**, which is loop preferred
  - **Loop Only**
  - **Point to Point**
  - **Loop Preferred**
5. Select **OK**.

## Enabling Endpoints

Here is how to enable endpoints.

### Procedure

1. In the **Hardware > Fibre Channel > Physical Resources** tab, select **Endpoints > Enable** from the **More Tasks** menu to display the **Enable Endpoints** dialog.

---

#### Note

If all endpoints are already enabled, a message to that effect is displayed.

---

2. Select one or more endpoints from the list, and select **Next**.
3. After the confirmation, select **Next** to complete the task.

## Disabling Endpoints

Here is how to disable an endpoint.

### Procedure

1. In the **Hardware > Fibre Channel > Physical Resources** tab, select **Endpoints > Disable** to display the **Disable Endpoints** dialog.
2. Select one or more endpoints from the list, and select **Next**.

---

#### Note

If an endpoint is in use, you are warned that deleting it might disrupt the system.

---

3. Select **Next** to complete the task.

## Modifying an Endpoint's System Address

You can modify the active system address for a SCSI Target endpoint using the `scsitararget endpoint modify` command option. This may be useful if the endpoint is associated with a system address that no longer exists, for example after a controller

upgrade or when a controller HBA (host bus adapter) has been moved. When the system address for an endpoint is modified, all properties of the endpoint, including WWPN and WWNN (worldwide port and node names, respectively), if any, are preserved and are used with the new system address.

In the following steps, suppose that an endpoint, ep-1, was assigned to system address 5a, but this system address is no longer valid. Instead a new controller HBA has been added at system address 10a. The SCSI Target subsystem automatically creates a new endpoint, ep-new, for the newly discovered system address. Currently, only a single endpoint can be associated with a given system address, so ep-new must be deleted, and then ep-1 must be assigned to system address 10a.

Be aware that it may take some time for the modified endpoint to come online, depending on the SAN environment, since the WWPN and WWNN have moved to a different system address. SAN zoning may also need to be updated to reflect the different configuration.

### Procedure

1. Show all endpoints to verify the endpoints to be changed:

```
# scsitarget endpoint show list
```

2. Disable all endpoints:

```
# scsitarget endpoint disable all
```

3. Delete the new, unnecessary endpoint, ep-new:

```
# scsitarget endpoint del ep-new
```

4. Modify the endpoint you want to use, ep-1, by assigning it the new system address 10a:

```
# scsitarget endpoint modify ep-1 system-address 10a
```

5. Enable all endpoints:

```
# scsitarget endpoint enable all
```

## Viewing Endpoints Information

When you select Endpoints from the Physical Resources menu, you will see two tabs: Hardware and Endpoints, as follows:

**Table 25** Hardware

Item	Description
Name	Specific name of endpoint.
Model	Model of hardware.
Firmware	Data Domain system HBA (host bus adapter) firmware version.
WWPN	Port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel port, and which is worldwide unique.
WWNN	Node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel node, and which is worldwide unique.
Physical Port	Physical port number.

**Table 26** Endpoints

Item	Description
Endpoint	Specific name of endpoint.
Connection Type	Connection type, such as N-Port, loop, or SAN (storage area network).
Link Speed	Transmission speed of link, in Gbps (Gigabits per second).
Port ID	Port identifier.
Enabled	HBA (host bus adapter) port operational state, which is either <code>Yes</code> (enabled) or <code>No</code> (not enabled).
Status	Data Domain system VTL link status, which is either <code>Online</code> (capable of handling traffic) or <code>Offline</code> .

## Viewing Endpoint Information

When you select a *specific* endpoint, from Physical Resources > Endpoints, you will see four tabs: Hardware, Summary, Statistics, and Detailed Statistics, as follows:

**Table 27** Hardware

Item	Description
Name	Specific name of endpoint.
Model	Model of hardware.
Firmware	Data Domain system HBA (host bus adapter) firmware version.
WWPN	Port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel port, and which is worldwide unique.
WWNN	Node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel node, and which is worldwide unique.
Physical Port	Physical port number.

**Table 28** Summary

Item	Description
Endpoint	Specific name of endpoint.
Connection Type	Connection type, such as N-Port, loop, or SAN (storage area network).
Link Speed	Transmission speed of link, in Gbps (Gigabits per second).
Port ID	Port identifier.
Enabled	HBA (host bus adapter) port operational state, which is either <code>Yes</code> (enabled) or <code>No</code> (not enabled).

**Table 28** Summary (continued)

Item	Description
Status	Data Domain system VTL link status, which is either <i>Online</i> (capable of handling traffic) or <i>Offline</i> .

**Table 29** Statistics

Item	Description
Endpoint	Specific name of endpoint.
Library	Name of library containing endpoint.
Device	Number of device.
Ops/s	Operations per second.
Read KiB/s	Speed of reads in KiB per second.
Write KiB/s	Speed of writes in KiB per second.

**Table 30** Detailed Statistics

Item	Description
Endpoint	Specific name of endpoint.
# of Control Commands	Number of control commands.
# of Read Commands	Number of read commands.
# of Write Commands	Number of write commands.
In (MiB)	Number of MiB written (the binary equivalent of MB).
Out (MiB)	Number of MiB read.
# of Error Protocol	Number of error protocols.
# of Link Fail	Number of link failures.
# of Invalid Crc	Number of invalid CRCs (cyclic redundancy checks).
# of Invalid TxWord	Number of invalid tx (transmission) words.
# of Lip	Number of LIPs (loop initialization primitives).
# of Loss Signal	Number of signals or connections that have been lost.
# of Loss Sync	Number of signals or connections that have lost synchronization.

## Endpoints Tab of Physical Resources

When you select Endpoints under Physical Resources, you will see the following information:

Item	Description
Name	Specific name of endpoint.

Item	Description
WWPN	Port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel port, and which is worldwide unique.
WWNN	Node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel node, and which is worldwide unique.
Physical Port	Physical port number.
Enabled	HBA (host bus adapter) port operational state, which is either <code>Yes</code> (enabled) or <code>No</code> (not enabled).
Status	Data Domain system VTL link status, which is either <code>Online</code> (capable of handling traffic) or <code>Offline</code> .

## About Initiators

An *initiator* is a client system FC HBA (Fibre Channel Host Bus Adapter) WWPN (World-Wide Port Name) with which the Data Domain system interfaces. An initiator name is an alias for the client's WWPN, for ease of use.

### Note

While a client is mapped as an initiator – but before an access group has been added – the client cannot access any data on a Data Domain system.

After adding an access group for the initiator or client, the client can access only the devices in the access group. A client can have access groups for multiple devices.

An access group may contain multiple initiators (a maximum of 128), but an initiator can exist in only one access group. A maximum of 512 initiators can be configured for a Data Domain system.

This section contains the following topics:

- ◆ [Viewing Initiators Information on page 295](#)
- ◆ [Initiators Tab of Physical Resources on page 295](#)

## Adding an Initiator

Here is how to add an initiator.

### Procedure

1. Select **Configure Initiators**, which takes you to the Hardware > Fibre Channel > Physical Resources tab.
2. Under Initiators, select Add (+) to display the Add Initiator dialog.
3. Enter the port's unique WWPN in the specified format.
4. Enter a Name for the initiator.
5. Select the Address Method: **Auto** is used for standard addressing, and **VSA** (Volume Set Addressing) is used primarily for addressing virtual buses, targets, and LUNs.
6. Select **OK**.

## Modifying an Initiator

Here is how to modify an initiator.

### Procedure

1. Select **Configure Initiators**, which takes you to the Hardware > Fibre Channel > Physical Resources tab.
2. Under Initiators, select one of the initiators to modify, and select the editing icon (a **pencil**) to display the Modify Initiator dialog.
3. You can change the initiator's Name and Address Method [**Auto** is used for standard addressing, and **VSA** (Volume Set Addressing) is used primarily for addressing virtual buses, targets, and LUNs.]
4. Select **OK**.

## Deleting an Initiator

---

### Note

Before you can delete an initiator, it must be offline and not attached to any group. Otherwise, you will get an error message, and the initiator will not be deleted.

---

Here is how to delete an initiator.

### Procedure

1. Select **Configure Initiators**, which takes you to the Hardware > Fibre Channel > Physical Resources tab.
2. Under Initiators, select the initiator(s) to delete, and select the delete icon (**X**).

## Viewing Initiators Information

The Initiators page contains the following information:

Item	Description
Name	Name of initiator.
Group	Group associated with initiator.
Online Endpoints	Endpoints seen by initiator. Displays <i>none</i> or <i>offline</i> if initiator is unavailable.
WWPN	Port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel port, and which is worldwide unique.
WWNN	Node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel node, and which is worldwide unique.
Vendor Name	Model of initiator.

## Initiators Tab of Physical Resources

When you select Initiators under Physical Resources, you will see the following information:

Item	Description
Name	Name of initiator, which is either the WWPN or the alias assigned to initiator (see <a href="#">About Initiators on page 294</a> ).
WWPN	Port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel port, and which is worldwide unique.
WWNN	Node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit <i>Network Address Authority</i> identifier), of the Fibre Channel node, and which is worldwide unique.
Online Endpoints	Group name where ports are seen by initiator. Displays <code>None</code> or <code>Offline</code> if initiator is unavailable.

## Setting a Loop ID

Some backup software requires that all private-loop targets have a hard address (loop ID) that does not conflict with another node. The range for a loop ID is 0-125.

### Procedure

1. Go to the Hardware > Fibre Channel > Physical Resources tab.
2. From the More Tasks menu, select **Set Loop ID**.
3. Enter the loop ID, and select **OK**.

## FC Link Monitoring

Here is some important information on how the current and previous releases of DD OS handle FC Link Monitoring.

### DD OS 5.3 and Later

Port monitoring detects an FC (Fibre Channel) port at system startup and raises an alert if the port is enabled and offline. To clear the alert, disable an unused port using the `scsitarget` or `vtl` commands.

### DD OS 5.1 up to 5.3

If a port is offline, an alert notifies you that the link is down. This alert is managed, which means it stays active until cleared. This occurs when the VTL FC port is online or disabled. If the port is not in use, disable it unless it needs to be monitored.

### DD OS 5.0 up to 5.1

If a port is offline, an alert notifies you the link is down. The alert is not managed, which means it does not stay active and does not appear in the current alerts list. When the port is online, an alert notifies you that the link is up. If the port is not in use, disable it unless it needs to be monitored.

### DD OS 4.9 up to 5.0

An FC port must be included in a VTL group to be monitored.



# CHAPTER 15

## Working with DD Replicator

This chapter includes:

◆ <a href="#">About EMC Data Domain Replicator</a> .....	298
◆ <a href="#">Replication Types</a> .....	298
◆ <a href="#">Supported Replication Topologies</a> .....	301
◆ <a href="#">Using Encryption of Data at Rest with Replication</a> .....	304
◆ <a href="#">Encryption on the Wire</a> .....	304
◆ <a href="#">Low Bandwidth Optimization</a> .....	305
◆ <a href="#">Bandwidth Delay Settings</a> .....	305
◆ <a href="#">About the Replication View</a> .....	305
◆ <a href="#">Preparing to Configure Replication</a> .....	310
◆ <a href="#">Configuring Replication</a> .....	310
◆ <a href="#">Resynchronizing Data in a Replication Pair</a> .....	318
◆ <a href="#">Recovering Data from a Replication Pair</a> .....	320
◆ <a href="#">Replication Seeding</a> .....	321
◆ <a href="#">Monitoring Replication</a> .....	322

## About EMC Data Domain Replicator

EMC Data Domain Replicator provides automated, policy-based, network-efficient, and encrypted replication for DR (disaster recovery) and multi-site backup and archive consolidation. DD Replicator software asynchronously replicates only compressed, deduplicated data over a WAN (wide area network).

Cross-site deduplication further reduces bandwidth requirements when multiple sites are replicating to the same destination system. With cross-site deduplication, any redundant segment previously transferred by any other site, or as a result of a local backup or archive, will not be replicated again. This improves network efficiency across all sites and reduces daily network bandwidth requirements up to 99%, making network-based replication fast, reliable, and cost-effective.

In order to meet a broad set of DR requirements, DD Replicator provides flexible replication topologies, such as full system mirroring, bi-directional, many-to-one, one-to-many, and cascaded. In addition, you can choose to replicate either all or a subset of the data on your Data Domain system. For the highest level of security, DD Replicator can encrypt data being replicated between Data Domain systems using the standard SSL (Secure Socket Layer) protocol.

DD Replicator scales performance and supported fan-in ratios to support large enterprise environments. When deployed over a 10GB network, DD Replicator can mirror data between two systems at up to 52 TB/hr. In addition, DD Replicator enables up to 270 remote systems to replicate into a single DD990 system, which simplifies the administration and reduces the costs of providing DR for remote sites.

Here are some points to consider when using DD Replicator:

- ◆ DD Replicator is a licensed product. Contact Data Domain Sales to obtain a license, then install the license, as described in [Displaying Licenses on page 37](#).
- ◆ A DD Replicator license is required for DD Boost to display tabs other than the File Replication tab.
- ◆ For encryption other than for systems with the Data at Rest option: if DD Boost file replication encryption is set to on, it must be set to on for both the source and destination systems.
- ◆ It is recommended that you disable all replication throttles when using managed file replication.
- ◆ A file (or directory) may not be renamed or moved into or out of a replication source; this includes a Cut operation followed by a Paste operation in Windows.
- ◆ You can usually replicate only between machines that are within two releases of each other, for example, from 5.1. to 5.3. However, there may be exceptions to this (as a result of atypical release numbering), so check with your EMC representative.
- ◆ If you are unable to manage and monitor Replication from the current version of the EMC Data Domain System Manager, use the Replication commands described in *EMC DD OS Command Reference Guide*.

## Replication Types

Replication typically consists of a *source* Data Domain system (which receives data from a backup system) and one or more *destination* Data Domain systems.

DD Replicator performs two levels of deduplication to significantly reduce bandwidth requirements: *local* and *cross-site* deduplication. Local deduplication determines the unique segments to be replicated over a WAN (wide area network). Cross-site

deduplication avoids sending any segments that may already exist on the destination, due to replication from another site or a local backup or archive at that site.

The choice of replication type depends on your specific needs:

- ◆ *Managed file replication*, which is used by DD Boost, directly transfers a backup image from one Data Domain system to another, one at a time, at the request of the backup software. The backup software keeps track of all copies, allowing easy monitoring of replication status and recovery from multiple copies. Managed file replication offers flexible replication topologies including full system mirroring, bi-directional, many-to-one, one-to-many, and cascaded, enabling efficient cross-site deduplication. See [DD Boost Managed File Replication on page 330](#).
- ◆ *Directory replication* transfers deduplicated changes of any file or subdirectory within a Data Domain file system directory that has been configured as a replication *source* to a directory configured as a replication *target* on a different system. Directory replication also has the same flexible network deployment topologies and cross-site deduplication effects as managed file replication. See [Directory Replication on page 299](#).
- ◆ *MTree replication* is used to replicate MTrees (see also [Working with MTrees on page 185](#)) between Data Domain systems. Periodic snapshots are created on the source, and the differences between them are transferred to the destination by leveraging the same cross-site deduplication mechanism used for directory replication. This ensures that the data on the destination is always a point-in-time copy of the source, with file consistency. This also reduces replication of churn in the data, leading to more efficient utilization of the WAN. MTree replication also has the same flexible network deployment topologies and cross-site deduplication effects as managed file replication. See [MTree Replication on page 300](#).

---

#### Note

Replicating directories under an MTree is not permitted.

- ◆ *Collection replication* performs whole-system mirroring in a one-to-one topology, continuously transferring changes in the underlying collection, including all of the logical directories and files of the Data Domain file system. While collection replication does not have the flexibility of the other types, it is very simple and lightweight, so it can provide higher throughput and support more objects with less overhead, which is ideal in high-scale enterprise cases. See [Collection Replication on page 300](#).

---

#### Note

Retention Lock Compliance supports collection and MTree replication only; directory replication is not supported. For more information on using replication with Retention Lock Compliance, see [DD Retention Lock Compliance on page 176](#).

---

## Directory Replication

Directory replication replicates data at the level of individual subdirectories under `/data/coll/backup`.

With directory replication, a Data Domain System can simultaneously be the *source* of some replication contexts and the *destination* for other contexts. The Data Domain system can also receive data from backup and archive applications while it is replicating data. See [Replication Types on page 298](#).

Here are some points to consider when using directory replication:

- ◆ Renaming (moving) files or tapes *into or out of* a directory replication source directory is *not* permitted. Renaming files or tapes *within* a directory replication source directory *is* permitted.
- ◆ A destination Data Domain System must have available storage capacity of at least the post-compressed size of the expected maximum post-compressed size of the source directory.
- ◆ When replication is initialized, a destination directory is created automatically.
- ◆ After replication is initialized, ownership and permissions of the destination directory are always identical to those of the source directory. As long as the context exists, the destination directory is kept in a read-only state and can receive data only from the source directory.
- ◆ At any time, due to differences in global compression, the source and destination directory can differ in size.

## MTree Replication

MTree replication replicates data for an MTree specified by the `/data/coll/mtree` pathname. With MTree replication, a Data Domain system can simultaneously be the *source* of some replication contexts and the *destination* for other contexts. The Data Domain system can also receive data from backup and archive applications while it is replicating data.

Here are some points to consider when using MTree replication:

- ◆ A destination Data Domain System must have available storage capacity of at least the post-compressed size of the expected maximum post-compressed size of the source MTree.
- ◆ When replication is initialized, a destination MTree is created automatically.
- ◆ After replication is initialized, ownership and permissions of the destination MTree are always identical to those of the source MTree. If the context is configured, the destination MTree is kept in a read-only state and can receive data only from the source MTree.
- ◆ At any time, due to differences in global compression, the source and destination MTree can differ in size.

## Collection Replication

Collection replication replicates the entire `/data/coll` area from a source Data Domain system to a destination Data Domain system.

Collection replication requires that the storage capacity of the destination system be equal to, or greater than, that of the source system. If the destination capacity is less than that of the source, the available capacity on the source is reduced to that of the destination.

The Data Domain system to be used as the collection replication destination must be empty before configuring replication. After replication is configured, this system is dedicated to receive data from the source system, and data can be read only from this system.

With collection replication, all user accounts and passwords are replicated from the source to the destination.

## Supported Replication Topologies


This section describes the following supported replication topologies and the typical uses for those configurations:

- ◆ [One-to-One Replication on page 301](#)
- ◆ [Bi-Directional Replication on page 301](#)
- ◆ [One-to-Many Replication on page 302](#)
- ◆ [Many-to-One Replication on page 302](#)
- ◆ [Cascaded Replication on page 303](#)

### One-to-One Replication

The simplest type of replication is from a Data Domain source system to a Data Domain destination system, otherwise known as a *one-to-one* replication pair. This replication topology can be configured with directory, MTree, or collection replication types. To set up this type of configuration, see [Creating a Replication Pair on page 311](#).

**Figure 9** One-to-One Replication Pair




Data flows from the source to the destination system

### Bi-Directional Replication

In a bi-directional replication pair, data from a directory or MTree on System A is replicated to System B, and from another directory or MTree on System B to System A.

To set up this type of configuration, see [Configuring Bi-Directional Replication on page 313](#).

**Figure 10** Bi-Directional Replication




Data flows in both directions between two systems

## One-to-Many Replication

In one-to-many replication, data flows from a source directory or MTree on one system to several destination systems. You could use this type of replication to create more than two copies for increased data protection, or to distribute data for multi-site usage.

To set up this type of configuration, see [Configuring One-to-Many Replication on page 313](#).

**Figure 11** One-to-Many Replication




Data flows from a directory or MTree source system to many destination systems

## Many-to-One Replication

In many-to-one replication, whether with MTree or directory, replication data flows from several source systems to a single destination system. This type of replication can be used to provide data recovery protection for several branch offices on a corporate headquarter's IT system.

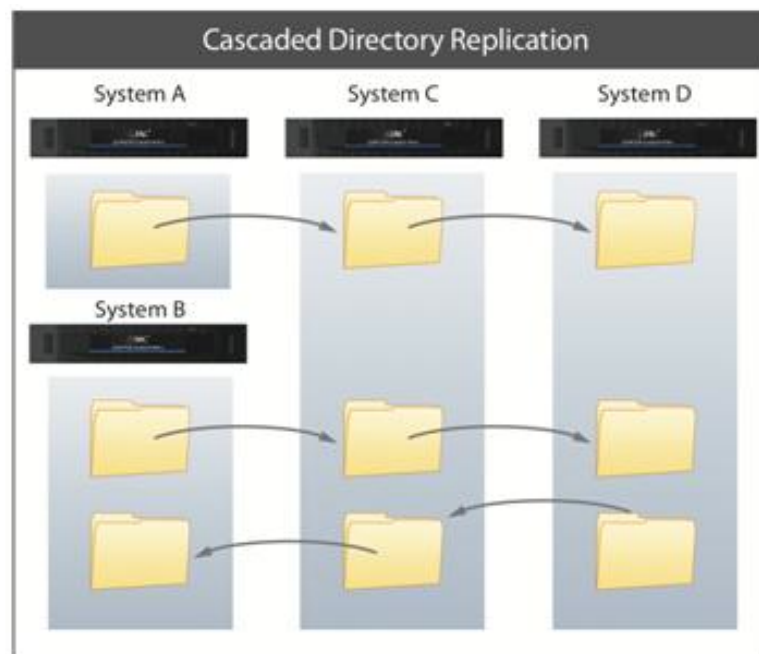
To set up this type of configuration, see [Configuring Many-to-One Replication on page 313](#).

**Figure 12** Many-to-One Replication

Data flows from many source systems to one destination system

## Cascaded Replication

In a cascaded replication topology, a source directory or MTree is chained among three Data Domain systems. The last hop in the chain can be configured as collection, MTree, or directory replication, depending on whether the source is directory or MTree. For example, DD System A replicates one or more MTrees to DD System B, which then replicates those MTrees to DD System C. The MTrees on DD System B are both a destination (from DD System A) and a source (to DD System C).

**Figure 13** Cascaded Directory Replication

Data recovery can be performed from the non-degraded replication pair context. For example:

- ◆ In the event DD System A requires recovery, data can be recovered from DD System B.
- ◆ In the event DD System B requires recovery, the simplest method is to perform a replication resync from DD System A to (the replacement) DD System B. In this case,

the replication context from DD System B to DD System C should be broken first. After the DD System A to DD System B replication context finishes resync, a new DD System B to DD System C context should be configured and resynced.

To set up this type of configuration, see [Configuring Cascaded Replication on page 313](#).

## Using Encryption of Data at Rest with Replication

Data Domain Replicator software can be used with the optional *Encryption of Data at Rest* feature, enabling encrypted data to be replicated using collection, directory, or MTree for all of the supported topologies.

To configure and work with Encryption of Data at Rest, see [Managing Encryption of Data at Rest on page 147](#).

Replication contexts are always authenticated with a *shared secret*. That shared secret is used to establish a session key using a Diffie-Hellman key exchange protocol, and that session key is used to encrypt and decrypt the Data Domain system encryption key when appropriate.

Each replication form works uniquely with encryption and offers the same level of security.

- ◆ Collection replication requires the source and target to have the same encryption configuration, because the target is expected to be an exact replica of the source data. In particular, the encryption feature must be turned on or off at both source and target, and if the feature is turned on, the encryption algorithm and the system passphrases must also match. The parameters are checked during the replication association phase.  
During collection replication, the source system transmits the encrypted user data with the encrypted system encryption key. The data can be recovered at the target because the target machine has the same passphrase and the same system encryption key.
- ◆ MTree or directory replication does not require encryption configuration to be the same at both the source and target Data Domain systems. Instead, the source and target securely exchange the target system's encryption key during the replication association phase, and the data at rest is first decrypted and then re-encrypted at the source using the target system's encryption key before transmission to the target. If the target machine has a different encryption configuration, the data transmitted is prepared appropriately. For example, if the feature is turned off at the target, the source decrypts the data, and it is sent to the target un-encrypted.
- ◆ In a cascaded replication topology, a replica is chained among three Data Domain systems. The last system in the chain can be configured as a collection, MTree, or directory. If the last system is a collection replication target, it uses the same encryption keys and encrypted data as its source. If the last system is an MTree or directory replication target, it uses its own key, and the data is encrypted at its source. The encryption key for the target at each link is used for encryption. Encryption for systems in the chain works as in a replication pair.

## Encryption on the Wire

DD Replicator supports encryption of data-in-flight by using standard SSL (Secure Socket Layer) protocol version 3, which uses the ADH-AES256-SHA cipher suite to establish secure replication connections.



## Low Bandwidth Optimization

For enterprises with small data sets and 6 Mb/s or less bandwidth networks, Data Domain Replicator can further reduce the amount of data to be sent using *low bandwidth optimization* mode. This enables remote sites with limited bandwidth to use less bandwidth or to replicate and protect more of their data over existing networks.

After enabling low bandwidth optimization on the source and target systems, both systems must undergo a full cleaning cycle to prepare the existing data. Issue the command `filesys clean start` on both systems. The duration of the cleaning cycle depends on the amount of data on the Data Domain system, but takes longer than a normal cleaning.

## Bandwidth Delay Settings

Bandwidth delay settings are used to control the TCP (transmission control protocol) buffer size. This allows the source system to send enough data to the destination while waiting for an acknowledgment.

Both the source and destination systems must have the same bandwidth delay settings. These tuning controls can benefit replication performance over higher latency links.

## About the Replication View

The Replication view lets you configure replication pairs and see the configured replicas as a list and as a topology map, check performance graphs, and configure network settings that affect performance.

The Replication view contains the following components:

- ◆ [Replication Status on page 305](#)
- ◆ [Summary View on page 306](#)
- ◆ [DD Boost View on page 308](#)
- ◆ [Topology View on page 309](#)
- ◆ [Performance View on page 309](#)
- ◆ [Advanced Settings View on page 309](#)

To display the Replication view:

### Procedure

1. Select the source system in the Navigation Panel of the EMC Data Domain System Manager.
2. Select the Replication tab to access the Replication view.  
The Replication Status and Summary view is displayed.
3. Select a replication context in the table to populate the Detailed Information area of the Summary view.

## Replication Status

Replication Status shows the system-wide count of replication contexts exhibiting a warning (yellow text) or error (red text) state, or if conditions are normal.

## Summary View

The Summary view lists the configured replication contexts for the system. Selecting a context in the table populates the content's information in Detailed Information.

The Summary view shows aggregated information about the selected Data Domain system – that is, summary information about all of the system's inbound replication pairs and all of that system's outbound replication pairs. The focus is the Data Domain system itself and the inputs to it and outputs from it.

The Detailed Information Panel, by contrast, shows the information for a selected individual replication pair.

The Summary table can be filtered by entering a Source or Destination name, or by selecting a State (Error, Warning, or Normal).

The Summary view includes the following information:

Item	Description
Source	System and path name of source context, with format <i>system.path</i> . For example, for directory <code>dir1</code> on system <code>dd120-22</code> , you would see <code>dd120-22.datadomain.com/data/coll/dir1</code> .
Destination	System and path name of destination context, with format <i>system.path</i> . For example, for MTree <code>MTree1</code> on system <code>dd120-44</code> , you would see <code>dd120-44.datadomain.com/data/coll/MTree1</code> .
Type	Type of context: MTree, directory (Dir), or Pool.
State	Current state describes replication pair status. Possible states include: <ul style="list-style-type: none"> <li>• Normal – If the replica is Initializing, Replicating, Recovering, Resyncing, or Migrating.</li> <li>• Idle – For MTree replication, this state can display if the replication process is not currently active or for network errors (such as the destination system being inaccessible).</li> <li>• Warning – If there is an unusual delay for the first five states, or for the Uninitialized state.</li> <li>• Error – Any possible error states, such as Disconnected.</li> </ul>
Synced As Of Time	Timestamp for last automatic replication sync operation performed by source. For MTree replication, this value is updated when a snapshot is exposed on the destination. For directory replication, it is updated when a sync point inserted by the source is applied. A value of unknown displays during replication initialization.
Pre-Comp Remaining	Amount of pre-compressed data remaining to be replicated.
Completion Time (Est.)	Value is either <code>Completed</code> , or the estimated amount of time required to complete the replication data transfer based on the last 24 hours' transfer rate.

## Detailed Information

Detailed Information provides the following information for the selected replication context:

Item	Description
State Description	Message about state of replica.
Source	System and path name of source context, with format <i>system.path</i> . For example, for <code>dir1</code> on system <code>dd120-22</code> , you would see <code>dd120-22.datadomain.com/data/coll/dir1</code> .
Destination	System and path name of destination context, with format <i>system.path</i> . For example, for <code>MTree MTree1</code> on system <code>dd120-44</code> , you would see <code>dd120-44.datadomain.com/data/coll/MTree1</code> .
Connection Port	System name and listen port used for replication connection. See <a href="#">Changing Host Connection Settings on page 315</a> .
<b>Completion Stats</b>	
Synced As Of Time	Timestamp for last automatic replication sync operation performed by source. For MTree replication, this value is updated when a snapshot is exposed on the destination. For directory replication, it is updated when a sync point inserted by the source is applied. A value of unknown displays during replication initialization.
Completion Time (Est.)	Value is either <code>Completed</code> or the estimated amount of time required to complete the replication data transfer based on the last 24 hours' transfer rate.
Pre-Comp Remaining	Amount of data remaining to be replicated.
Files Remaining	(Directory Replication Only) Number of files that have not yet been replicated.
<b>Status</b>	For source and destination endpoints, shows status (Enabled, Disabled, Not Licensed, etc.) of major components on the system, such as: <ul style="list-style-type: none"> <li>• Replication</li> <li>• File System</li> <li>• Replication Lock</li> <li>• Encryption at Rest</li> <li>• Encryption over Wire</li> <li>• Available Space</li> <li>• Low Bandwidth Optimization</li> <li>• Compression Ratio</li> <li>• Low Bandwidth Optimization Ratio</li> </ul>

### Performance Chart

Select **Performance Chart** to open a Replication graph for the selected context.

The Replication Performance graph shows performance over time and differs depending on whether it is for a collection or a directory pair, or for an MTree or Pool.

Item	Description for Collection	Description for a Directory, MTree, or Pool
Pre-Comp Remaining	Pre-compressed data remaining to be replicated.	Pre-compressed data remaining to be replicated.
Pre-Comp Written	Pre-compressed data written on the source.	Pre-compressed data written on the source.
Post-Comp Replicated	Post-compressed data that has been replicated.	Post-compressed data that has been replicated.

### Completion Predictor

A widget for predicting when replication will complete for the selected context. See [Tracking Status of a Backup Job's Replication Progress on page 322](#).

## DD Boost View

The DD Boost view provides configuration and troubleshooting information to NetBackup administrators who have configured their Data Domain system or systems to use DD Boost AIR (Automatic Image Replication) or any DD Boost application that uses managed file replication.

See the *EMC Data Domain Boost for OpenStorage Administration Guide* for DD Boost AIR configuration instructions.

The File Replication tab shows:

- ◆ Currently Active File Replication:
  - Direction (Out-Going and In-Coming) and the number of files in each.
  - Remaining data to be replicated (pre-compressed value in GiB) and the amount of data already replicated (pre-compressed value in GiB).
  - Total size: The amount of data to be replicated and the already replicated data (pre-compressed value in GiB).
- ◆ Most Recent Status: Total file replications and whether completed or failed
  - during the last hour
  - over the last 24 hours
- ◆ Remote Systems:
  - Select a replication from the list.
  - Select the time period to be covered from the menu.
  - Select **Show Details**. The source, destination, number of files complete/failed, pre-compressed size in MiB, and Network Throughput speed in Kbps for the selected replications over the specified time period is shown.

Storage Unit Associations displays the following information, which you can use for audit purposes or to check the status of DD Boost AIR events used for the storage unit's image replications:

- ◆ A list of all storage unit associations known to the system. The source is on the left, and the destination is on the right. This information shows the configuration of AIR on the Data Domain system.

- ◆ The Event Queue is the pending event list. It shows the local storage unit, the event ID, and the status of the event.

## Topology View

The Topology view shows how the selected Data Domain system's replication pairs are configured in the network.

- ◆ The arrow between Data Domain systems represents one or more replication pairs.
- ◆ Depending on the status of the contexts between the two systems, the arrow displays as normal (green), warning (yellow), or error (red).
- ◆ Select a context to open the Context Summary dialog, where context details can be viewed (paths, status), and links to other operations are available (**Show Summary**, **Modify Options**, **Enable/Disable Pair**, **Graph Performance**).
- ◆ Select **Collapse All** to roll-up the Expand All context view and show only the name of the system and the count of destination contexts.
- ◆ Select **Expand All** to show all the destination directory and MTree contexts configured on other systems.
- ◆ Select **Reset Layout** to return to the default view.
- ◆ Select **Print** to open a standard print dialog box.

## Performance View

The Performance view displays a graph that accurately represents the fluctuation of data during replication. However, during times of inactivity (when no data is being transferred), the shape of the graph may display a gradually descending line, instead of an expected sharply descending line.

The Performance view displays a replication's historical data for:

- ◆ Network In – Total replication network bytes entering the system (all contexts)
- ◆ Network Out – Total replication network bytes leaving the system (all contexts)

These are aggregated statistics of each replication pair for this Data Domain system. The duration (x-axis) is 21 days by default. The y-axis is in GibiBytes or MebiBytes (the binary equivalents of GigaBytes and MegaBytes).

For an accurate reading, hover the cursor over points in the chart. A tooltip displays the ReplIn, ReplOut, date/time, and amount of data for a given point in time.

## Advanced Settings View

Advanced Settings provides management of, and detailed information about, the replication settings, as described in the following sections:

- ◆ [Throttle Settings on page 309](#)
- ◆ [Network Settings on page 310](#)

### Throttle Settings

Throttle Settings shows the current settings for:

- ◆ Temporary Override – If configured, shows the throttle rate or 0, which means all replication traffic is stopped.
- ◆ Permanent Schedule – Shows the time for days of the week on which scheduled throttling occurs.

For details on configuring these options, see [Adding Throttle Settings on page 316](#).

## Network Settings

The following settings affect data transfer over the network.

### Bandwidth Settings

Shows (Default) if bandwidth has not been configured or the configured data stream rate. The average data stream to the replication destination is at least 98,304 bits per second (12 KiB). To configure bandwidth, see [Changing Global Network Settings on page 318](#).

Also see [Bandwidth Delay Settings on page 305](#).

### Delay

Shows (Default) if a network delay has not been configured or the configured network delay setting (in milliseconds). To configure network delay, see [Changing Global Network Settings on page 318](#).

Also see [Bandwidth Delay Settings on page 305](#).

### Listen Port

Shows (Default) if a listen port has not been configured or the configured global listen port. To configure the global listen port, see [Changing the Global Listen Port on page 318](#).

## Preparing to Configure Replication

Before starting Replication configuration, review these prerequisites:

- ◆ Both the source and destination Data Domain systems must be managed by the EMC Data Domain System Manager.
- ◆ Adequate storage must be available on the source and destination. At a minimum, the destination must have the *same amount of space* as the source.
- ◆ The destination directory for the context must not contain the destination directory for another context or must not be contained within the destination directory for another context.
- ◆ You must determine the type of replication configuration to use (see [Replication Types on page 298](#)).
- ◆ For Directory replication, the destination directory must be empty, or its contents no longer needed, because it will be overwritten.

## Limitations on Number of Contexts

Before configuring directory replication, determine the maximum number of contexts for your Data Domain system. This is the value shown in the *Repl Source Streams* column in the [Table 21 on page 132](#).

- ◆ If the source holds a lot of data, the replication operation can take many hours. Consider putting both Data Domain systems in the Replicator pair in the same location with a direct link to cut down on initialization time.
- ◆ A subdirectory under a source directory in a replication context cannot be used in another directory replication context. A directory can be in only one context at a time.

## Configuring Replication

To configure replication, follow the procedures for creating the replication contexts in [Creating a Replication Pair on page 311](#). After the replication pair has been created, data

replication begins, and the progression of the data copy can be monitored from the EMC Data Domain System Manager Replication Summary view (see [Monitoring Replication on page 322](#)).

Check that the capacity of the destination system is equal to or greater than the source. If the destination has less storage capacity than the source, the capacity of the source is reduced to that of the destination.

Other optional replication configuration tasks include the following:

- ◆ [Creating a Replication Pair on page 311](#)
- ◆ [Enabling and Disabling a Replication Pair on page 314](#)
- ◆ [Deleting a Replication Pair on page 314](#)
- ◆ [Converting a Directory Replication Pair to an MTree on page 314](#)
- ◆ [Changing Host Connection Settings on page 315](#)
- ◆ [Managing Bandwidth with Throttling on page 316](#)
- ◆ [Changing Network Settings on page 317](#)

## Creating a Replication Pair

To create and start initiation of a replication pair, use the following procedure (which is available to administrative users only). Before starting this procedure, make sure:

- ◆ The destination system has the same amount of space as the source system.
- ◆ For collection replication only, the destination file system must have been destroyed and subsequently created, but not enabled.

### Procedure

1. Select the source system in the Navigation Panel of the EMC Data Domain System Manager.
2. Select Replication to access the Replication view.
3. Under Summary, select **Create Pair** to display the Create Pair dialog.

Add specific information to create a collection, directory, MTree, or pool replication pair, as described in the following sections:

- [Creating a Collection Replication Pair on page 311](#)
- [Creating a Directory, MTree, or Pool Replication Pair on page 312](#)

---

### Note

For information on configuring parameters under Advanced, see [Changing Host Connection Settings on page 315](#).

---

## Creating a Collection Replication Pair

Here is how to create a collection replication pair.

### Procedure

1. In the Create Pair dialog, select **Collection** from the **Replication Type** menu.
2. Select the source system hostname from the **Source System** menu.
3. Select the destination system hostname from the **Destination System** menu.

The list includes only those hosts in the DD-Network list.

4. Select **OK**.

Replication from the source to the destination begins.

**Results**

Test results from Data Domain returned the following performance guidelines for replication initialization. These are guidelines *only*, and actual performance seen in production environments may vary.

- ◆ Over a gibibit LAN: With a high enough shelf count to drive maximum input/output and ideal conditions, collection replication can saturate a 1GigE link (modulo 10% protocol overhead), as well as 400-900 MB/sec on 10GigE, depending on the platform.
- ◆ Over a WAN, performance is governed by the WAN link line speed, bandwidth, latency, and packet loss rate.

## Creating a Directory, MTree, or Pool Replication Pair

Here is how to create a basic directory, MTree, or pool replication pair.

**Procedure**

1. In the Create Pair dialog box, select **Directory**, **MTree** (default), or **Pool** from the **Replication Type** menu.
2. Select the source system hostname from the **Source System** menu.
3. Select the destination system hostname from the **Destination System** menu.
4. Enter the source path in the Source Path text box (notice the first part of the path is a constant that changes based on the type of replication chosen).
5. Enter the destination path in the Destination Directory text box (notice the first part of the path is a constant that changes based on the type of replication chosen).
6. Select **OK**.

The Replication from the source to the destination begins.

**Note**

- If you try to create an MTree replication context using the name of an existing MTree, an error message is displayed; however, you can rename MTrees. Deleting the MTree will not take effect until the next garbage collection is run; therefore, this may not be the best choice.
- The following occurs only during the first phase of directory replication initialization: When a new directory, MTree, or pool replication pair is being created, the source directory cannot be written to until the replication relationship between the source and destination systems has been established. Attempts to write to the newly configured replication source directory will fail until the replication relationship has been established. Instead, schedule the replication configuration at a time when backups are not occurring.
- (Directory Replication Only) Replicated files can arrive on the destination system in a order different from how they were closed (or last written) on the source. If file order is important to your site, see the release notes for additional details and a workaround, or contact technical support.



## Results

Test results from Data Domain returned the following guidelines for estimating the time needed for replication initialization. These are guidelines *only* and may not be accurate in specific production environments.

- ◆ Using a T3 connection, 100ms WAN, performance is about 40 MiB/sec of pre-compressed data, which gives data transfer of:  
40 MiB/sec = 25 seconds/GiB = 3.456 TiB/day
- ◆ Using the base-2 equivalent of gigabit LAN, performance is about 80 MiB/sec of pre-compressed data, which gives data transfer of about double the rate for a T3 WAN.

For advanced directory or MTree replication pair configurations that build on this procedure, see:

- ◆ [Configuring Bi-Directional Replication on page 313](#)
- ◆ [Configuring One-to-Many Replication on page 313](#)
- ◆ [Configuring Many-to-One Replication on page 313](#)
- ◆ [Configuring Cascaded Replication on page 313](#)

## Configuring Bi-Directional Replication

To create the configuration described in [Bi-Directional Replication on page 301](#), use the procedure [Creating a Directory, MTree, or Pool Replication Pair on page 312](#) to create a replication pair (for example, using mtree2) from host A to host B.

Use the same procedure to create a replication pair (for example, using mtree1) from host B to host A. For this configuration, destination pathnames cannot be the same. Guidelines for directory and MTree replication are applicable.

## Configuring One-to-Many Replication

To create the configuration described in [One-to-Many Replication on page 302](#), use the procedure [Creating a Directory, MTree, or Pool Replication Pair on page 312](#) to create pairs (for example, using mtree1) on host A to:

- ◆ mtree1 on host B
- ◆ mtree1 on host C
- ◆ mtree1 on host D

---

### Note

A replication recovery cannot be done to a source context whose path is the source path for other contexts; the other contexts must be broken and resynced after the recovery.

---

## Configuring Many-to-One Replication

To create the configuration described in [Many-to-One Replication on page 302](#), use the procedure [Creating a Directory, MTree, or Pool Replication Pair on page 312](#) to create a pair, for example:

- ◆ mtree1 from host A to mtree1 on host C
- ◆ mtree2 on host B to mtree2 on host C

## Configuring Cascaded Replication

To create the configuration described in [Cascaded Replication on page 303](#), use the procedure [Creating a Directory, MTree, or Pool Replication Pair on page 312](#) to create a pair for:

- ◆ mtree1 on host A to mtree1 on host B
- ◆ On host B, create a pair for mtree1 to mtree1 on host C  
The final destination context (on host C in this example, but more than three hops are supported) can be a collection replica or a directory or MTree replica.

## Enabling and Disabling a Replication Pair

Disabling a replication pair temporarily pauses the active replication of data between a source and a destination. The source stops sending data to the destination, and the destination stops serving an active connection to the source.

To disable a replication pair, from either the source or the destination:

### Procedure

1. Select one or more replication pairs in the Summary table, and select **Disable Pair** to display the Display Pair dialog.
2. Select **Next** and then **OK**.
3. To resume operation of a disabled replication pair, select one or more replication pairs in the Summary table, and select **Enable Pair** to display the Enable Pair dialog.
4. Select **Next** and then **OK**.

### Results

Replication of data is resumed.

## Deleting a Replication Pair

---

### Note

With collection replication, the file system is disabled when deleting the replication pair.

---

Here is how to delete a replication pair:

### Procedure

1. Select one or more replication pairs in the Summary table, and select **Delete Pair** to display the Delete Pair dialog.
2. Select **Next** and then **OK**.

### Results

The Replication pairs are deleted.

When a directory or MTree replication context is deleted, the destination directory or MTree, respectively, becomes writeable. When a collection replication pair is broken, the destination Data Domain system becomes a stand-alone read/write system.

## Converting a Directory Replication Pair to an MTree

When a directory replication pair is converted to an MTree, the directory data is initialized in an MTree, and the directory replication configuration is deleted.

Here is how to convert a directory context to an MTree.

### Procedure

1. Select the directory replication pair in the Summary table, and select **Convert to MTree** to display the Convert to MTree dialog.

The directory name is now the MTree name.

2. Select **OK**.

A Warning dialog is displayed, indicating the directory data is being initialized in the new MTree, and the old directory replication configuration is being deleted.

3. Select **OK**.

The Convert to MTree Status dialog is displayed, showing the progress of the conversion.

4. Select **Close**.

## Changing Host Connection Settings

Here is how to change replication pair connection settings.

### Procedure

1. Select the replication pair in the Summary table, and select **Modify Settings** to display the Modify Connection Settings dialog.

2. Choose one of the following options:

- [Configuring Low Bandwidth Optimization on page 315](#)
- [Configuring Encryption Over Wire on page 315](#)
- [Configuring a Non-Default Connection Port on page 315](#)

3. Select **Next** and then **Close**.

The replication pair settings are updated, and replication resumes.

### Configuring Low Bandwidth Optimization

Low bandwidth optimization can be enabled on a per-context basis. Low bandwidth optimization must be enabled on both the source and destination Data Domain systems. If the source and destination have incompatible low bandwidth optimization settings, low bandwidth optimization will be inactive for that context.

- ◆ To configure low bandwidth optimization, in the Modify Connection Settings dialog box, select the checkbox for **Use Low Bandwidth Optimization**.
- ◆ In Create Pair, Start Resync, or Start Recover (either for collection or directory), select Advanced, and select the checkbox for **Use Low Bandwidth Optimization**.

### Configuring Encryption Over Wire

To encrypt data sent over the replication network connection, in the Modify Connection Settings dialog, select the checkbox for **Enable Encryption Over Wire**.

Both sides of the connection must enable this feature for encryption to proceed. Encrypted replication will use the ADH-AES256-SHA cipher suite.

### Configuring a Non-Default Connection Port

The source system transmits data to a destination system listen port. As a source system can have replication configured for many destination systems (each of which can have a different listen port), each context on the source can configure the connection port to the corresponding listen port of the destination.

To change the connection port:

- ◆ In the Modify Connection Settings dialog box, under Details, select the checkbox for **Use Non-default Connection Host**, and in the source Connection Port text box, change the listen port to a new value.

- ◆ In Create Pair, Start Resync, or Start Recover (either for collection or directory), select Advanced, and under Connection, select the checkbox for **Use Non-default Connection Host**, and change the listen port to a new value.

## Managing Bandwidth with Throttling

To modify the amount of bandwidth used in the network, modify the throttle setting for the replication data stream. Throttling can be set to occur at certain times by establishing a schedule.

The average data stream to the replication destination is at least 98,304 bits per second (12 KiB).

Throttle options:

- ◆ Apply equally to all replication pairs and all network interfaces on a system.
- ◆ Affect only outbound network traffic.
- ◆ Calculate the proper TCP (transmission control protocol) buffer size for replication usage, using bandwidth settings.

## Adding Throttle Settings

Here is how to add throttle settings.

### Procedure

1. Select Replication > Advanced Settings, and select **Add Throttle Setting** to display the Add Throttle Setting dialog.
2. Set the days of the week for which throttling is to be active by selecting the checkboxes next to the days.
3. Set the time that throttling is to start with the **Start Time** drop-down selectors for the hour:minute and AM/PM.
4. Under Throttle Rate:
  - Select **Unlimited** to set no limits.
  - Enter a number in the text entry box (for example, 20000) and select the rate from the menu (bps, Bps, Kibps, or KiBps).
  - Select the **0 Bps (Disabled)** option to disable all replication traffic.
5. Select **OK** to set the schedule.

The new schedule is shown in Throttle Settings Permanent Schedule.

### Results

Replication runs at the given rate until the next scheduled change or until a new throttle setting forces a change.

## Deleting Throttle Settings

Here is how to delete throttle settings.

### Procedure

1. Select Replication > Advanced Settings, and select **Delete Throttle Setting** to display the Delete Throttle Setting dialog.
2. Select the checkbox for the throttle setting to delete or the heading checkbox to delete all settings. This list can include settings for the “disabled” state.
3. Select **OK** to remove the setting.

4. On the Delete Throttle Setting Status dialog, select **Close**.

## Temporarily Overriding a Throttle Setting

A throttle override temporarily changes a throttle setting. The current setting is listed at the top of the window.

With the clear option enabled, the setting is in effect until a scheduled change or a system reboot. If the clear option is disabled, the change is in effect indefinitely.

Here is how to temporarily override a throttle setting.

### Procedure

1. Select Replication > Advanced Settings, and select **Set Throttle Override** to display the Throttle Override dialog.
2. Select how the current throttle setting is to be overridden.
  - **Unlimited** – Reverts to the system-set throttle rate (no throttling performed).
  - Set the throttling bit and rate in the text entry box (for example, 20000) and (bps, Bps, Kibps, or KiBps).
  - **0 Bps (Disabled)** – Sets the throttle rate to 0, effectively stopping all replication network traffic.
3. To enforce the change temporarily, check the box **Clear at next scheduled throttle event**.
4. Select **OK** to save the setting.

---

### Note

To clear an override that has been set, select **Clear Throttle Override**, and select **OK**.

---

## Working with Low Bandwidth Optimization

Low bandwidth optimization can be used to improve data transfer over low bandwidth links. Using low bandwidth optimization adds increased data compression to optimize network bandwidth. More compression directly translates to more throughput on low bandwidth links. On high bandwidth links, the computational overhead of low bandwidth optimization may actually reduce throughput. For this reason, low bandwidth optimization is recommended on T2 and lower bandwidth links.

- ◆ To configure low bandwidth optimization, see [Changing Host Connection Settings on page 315](#).
- ◆ To check the status of a low bandwidth optimization configuration, see [Detailed Information on page 307](#).

## Changing Network Settings

Here is how to change network settings for bandwidth, network delay, and global IP listen port.

### Procedure

1. Select the source system in the Navigation Panel of the EMC Data Domain System Manager.
2. Select Replication > Advanced Settings.
3. Beside Network Settings, select **Change Network Settings** to display the Network Settings dialog.

4. Change either the network settings or listen port, referring to the following sections:
  - [Changing Global Network Settings on page 318](#)
  - [Changing the Global Listen Port on page 318](#)
5. Select **OK**.
 

The new settings appear in the Network Settings table.

## Changing Global Network Settings

Using the bandwidth and network-delay settings together, replication calculates the proper TCP (transmission control protocol) buffer size for replication usage.

---

### Note

- ◆ You can determine the actual bandwidth and the actual network delay values for each server by using the `ping` command.
- ◆ The default network parameters in a restorer work well for replication in low latency configurations, such as a local 100Mbps or 1000Mbps Ethernet network where the latency round trip time (as measured by the `ping` command) is usually less than 1 millisecond. The defaults also work well for replication over low- to moderate-bandwidth WANs where the latency may be as high as 50-100 milliseconds. However, for high-bandwidth high-latency networks, some tuning of the network parameters is necessary.
 

The key number for tuning is the bandwidth-delay number produced by multiplying the bandwidth and round-trip latency of the network. The number is a measure of how much data can be transmitted over the network before any acknowledgments can return from the far end. If the bandwidth-delay number of a replication network is more than 100,000, then replication performance benefits from setting the network parameters in both restorers.

---

These network settings are global to the Data Domain system and should be set only once per system.

### Procedure

1. In the Network Settings dialog, select **Custom Values**.
2. Enter Delay and Bandwidth values in the text boxes.
 

The network delay setting is in milliseconds, and bandwidth is in bytes per second.

## Changing the Global Listen Port

The default IP Listen Port for a replication destination for receiving data streams from the replication source is 2051. This is a global setting for the Data Domain system.

In the Network Settings dialog, simply enter a new value in the Listen Port text box.

## Resynchronizing Data in a Replication Pair

Resynchronization is the process of recovering (or bringing back into sync) the data between a source and destination replication pair after a manual break. The replication pair are resynchronized so both endpoints contain the same data.

A replication resynchronization can also be used:

- ◆ To recreate a context that has been deleted.

- ◆ When a directory replication destination runs out of space, but the source destination still has data to replicate.
- ◆ To convert a collection replication to directory replication.

## Resyncing a Directory, MTree, or Pool Replication Pair

Here is how to resync a directory, MTree, or pool replication pair.

### Procedure

1. Delete the context on both the source and destination directory replication systems.
2. From either the source or the destination directory replication system, select the **More** menu, and select **Start Resync** to display the Start Resync dialog.
3. Select the Replication Type to be resynced: **Directory, MTree, or Pool**.
4. Select the source system hostname from the Source System menu.
5. Select the destination system hostname from the Destination System menu.
6. Enter the directory path in the Source Path text box.
7. Enter the directory path in the Destination Path text box.
8. Select **OK**.

## Converting from a Directory to an MTree Replication Pair

A directory replication pair can be converted to an MTree replication pair.

A conversion is started with a replication resync that filters all data from the source Data Domain system to the destination Data Domain system. The filtering performance over a T3, 100ms WAN is about 100 MiB/sec, which gives data transfer of:

$$100 \text{ MiB/sec} = 10 \text{ seconds/GiB} = 8.6 \text{ TiB/day}$$

---

### Note

MiB = MibiBytes, the base-2 equivalent of Megabytes. GiB = GibiBytes, the base-2 equivalent of Gigabytes. TiB = TibiBytes, the base-2 equivalent of Terabytes.

Over a gibibit LAN, performance is about 120 MiB/sec, which gives data transfer of:

$$120 \text{ MiB/sec} = 8.3 \text{ seconds/GiB} = 10.3 \text{ TiB/day}$$

Here is how to convert from a directory to an MTree replication pair.

### Procedure

1. Create a new MTree on both the source and the destination (see [Create an MTree on page 192](#) for details).
2. From Replication > Summary, select the **Directory context to convert**.
3. From the More menu, select **Convert to MTree** to display the Convert to MTree dialog.
4. Add the new MTree paths to the source and destination text fields.
5. Select **OK**.

MTree replication is established after data is copied to the source MTree.

## Aborting a Resync of a Directory Replication Pair

Here is how to abort a resync in progress.

**Procedure**

1. From either the source or destination directory replication system, select the **More** menu, and select **Abort Resync** to display the Abort Resync dialog, which lists all contexts currently performing resynchronization.
2. Select the checkboxes of one or more contexts to abort their resync.
3. Select **OK**.

**Results**

Replication resynchronization is aborted.

## Recovering Data from a Replication Pair

If source replication data becomes inaccessible, it can be *recovered* from the replication pair destination. Either collection or directory can be recovered to the source, as described in the following procedures:

- ◆ [Recovering Directory Pool Data on page 320](#)
- ◆ [Recovering Collection Replication Pair Data on page 320](#)
- ◆ [Recovering Directory Replication Pair Data on page 321](#)  
You can also stop a recovery, as described in:
- ◆ [Aborting a Replication Pair Recovery on page 321](#)

**Note**

The source must be empty before the recovery can proceed.

Recovery can be performed for *most* types of replication topologies; however, there is *no option* for recovery when using MTree replication.

### Recovering Directory Pool Data

Here is how to recover directory pool data.

**Procedure**

1. Select the **More** menu, and select **Start Recover** to display the Start Recover dialog.
2. Select **Pool** from the **Replication Type** menu.
3. Select the source system hostname from the **System to recover to** menu.
4. Select the destination system hostname from the **System to recover from** menu.
5. Select the context on the destination from which data is recovered.
6. Select **OK** to start the recovery.

### Recovering Collection Replication Pair Data

**Note**

The source file system must be in a pristine state for the recovery to proceed. The destination context must be fully initialized for the recovery to be successful.

Here is how to recover a replication pair source.



**Procedure**

1. Select the **More** menu, and select **Start Recover** to display the Start Recover dialog.
2. Select **Collection** from the **Replication Type** menu.
3. Select the source system hostname from the **System to recover to** menu.
4. Select the destination system hostname from the **System to recover from** menu.
5. Select the context on the destination from which data is recovered. Only one collection will exist on the destination.
6. Select **OK** to start the recovery.

## Recovering Directory Replication Pair Data

---

**Note**

The same directory used in the original context must be created (but left empty) in order for the recovery to work.

---

Here is how to recover one or more directory replication pairs.

**Procedure**

1. Select the **More** menu, and select **Start Recover** to display the Start Recover dialog.
2. Select **Directory** from the **Replication Type** menu.
3. Select the hostname of the system that needs to have data restored to it from the **System to recover to** menu.
4. Select the hostname of the system that will be the data source from the **System to recover from** menu.
5. Select the context to restore from the context list.
6. Select **OK** to start the recovery.

## Aborting a Replication Pair Recovery

If a recovery fails or must be terminated, here is how to stop the replication recovery.

**Procedure**

1. Select the **More** menu, and select **Abort Recover** to display the Abort Recover dialog, which shows the contexts currently performing recovery.
2. Select the checkbox of one or more contexts to abort from the list.
3. Select **OK**.

**After you finish**

Recovery on the source should be restarted again, as soon as possible, by restarting the recovery.

## Replication Seeding

If the source has a lot of data, the initial replication seeding can take some time over a slow link.

To expedite the initial seeding, bring the destination system to the same location as the source system to use a high-speed, low-latency link. After data is initially seeded using

the high-speed network, move the system back to its intended location. After this "initial seeding," only new data will be sent.

All replication topologies are supported, and the procedures to start the replication are found in [Configuring Replication on page 310](#).

## Monitoring Replication

This section describes how to use the EMC Data Domain System Manager to check the status of Replication operations.

For an overview of the visual components of the Replication view, see [About the Replication View on page 305](#).

### Checking Replication Status

Replication status is available at all levels of the Replication page, as described in the following procedures:

- ◆ [Checking Replication Pair Status on page 322](#)
- ◆ [Tracking Status of a Backup Job's Replication Progress on page 322](#)
- ◆ [Tracking Status of a Replication Process on page 323](#)
- ◆ [Checking Performance of a Replication Context on page 323](#)

#### Checking Replication Pair Status

Here is how to check replication pair status.

##### Procedure

1. Select the Data Domain system to be checked in the Navigation Panel.  
The content of the Replication > Topology view changes to focus on that system. The system should have a blue arrow pointing to it in the Topology Panel.
2. In the Topology Panel, check the colors of the arrows showing the status of the context (for more information, see [Topology View on page 309](#)).
3. In Replication > Summary, from the **Filter By** drop-down list (under the Create Pair button), select **State**, and select **Error**, **Warning**, or **Normal** from the state menu.  
The Replication contexts are sorted according to the selection.

#### Tracking Status of a Backup Job's Replication Progress

Here is how to check the progress of a replication for a point in time.

##### Procedure

1. Select the Summary tab, and select a Replication context to display the Detailed Information area.
2. In the Completion Predictor area, select options from the **Source Time** drop-down list for a replication's completion time, and select **Track**.  
The estimated time displays, in the Completion Time area, for when a particular backup will finish its replication to the destination. If the replication is finished, the area shows **Completed**.

## Tracking Status of a Replication Process

To display the progress of a replication initialization, resync, or recovery operation, use the Replication > Summary view to check the current state.

## Checking Performance of a Replication Context

To check the performance of a replication context over time, select a Replication context in the Summary view, and select **Performance Graph** in the Detailed Information area.



# CHAPTER 16

## Working with DD Extended Retention

This chapter includes:

- ◆ [About DD Extended Retention Software](#)..... 326
- ◆ [Getting Started](#)..... 327
- ◆ [Initial Setup](#)..... 334
- ◆ [Administration](#)..... 339
- ◆ [Upgrades and Recovery](#)..... 343

## About DD Extended Retention Software

The DD Extended Retention software option (formerly Data Domain Archiver) addresses the need to cost-effectively retain and access deduplicated data for longer periods. Instead of storing long-term data on tape, it can be stored locally on the Data Domain system using Data Domain Extended Retention software option. The result is lower operational cost as the deduplication engine can be shared across more shelves and managed as a single namespace and a single system.

### Two-Tiered File System

DD Extended Retention provides an internal two-tiered file system—*active* and *retention*. (The file system appears to users as a single entity.) Incoming data to a Data Domain system with the DD Extended Retention software option is first placed in the active tier of the file system. The data (in the form of complete files) is later moved to the retention tier of the file system as specified by data movement policies. For example, the active tier might retain weekly full and daily incremental backups for 90 days while the retention tier might retain the monthly fulls for seven years.

### Transparency of Operation

Data Domain systems with DD Extended Retention software support existing backup applications using simultaneous data access methods through NFS and CIFS file service protocols over Ethernet, VTL for open systems and IBMi, or as a disk-based target using application-specific interfaces such as EMC Data Domain Boost (for use with EMC Avamar®, EMC NetWorker®, EMC GreenPlum, Symantec OpenStorage, and Oracle RMAN).

EMC Data Domain Extended Retention software extends the Data Domain architecture with automatic transparent data movement from the active to retention tier. It is designed to enable cost-effective long-term retention of data on deduplicated disk. All of the data in the two tiers is accessible, although there might be a slight delay on initial access to data in the retention tier. The namespace of the system is global, and not affected by data movement. No partitioning of the file system is necessary to take advantage of the two-tiered file system.

### Data Movement Policies

The data movement policies are based on when a file was last modified. You can set policies with dates and times that differ for subsets of data. Files that may be updated for a period of time need a different policy than ones that never change. Data movement policies are set on a per-MTree basis.

### Deduplication

The retention tier is comprised of one or more retention units, each of which may consist of one or more entire shelf's worth of storage. For fault isolation purposes, deduplication occurs within each retention unit. There is no cross-deduplication among retention units, or between active and retention tiers.

### Storage

The concept of tiering extends to the storage level. The active tier of the file system draws storage from the active tier of storage. The retention tier of the file system uses storage from the retention tier of storage. For both tiers, DD OS 5.2 and later support ES20 and ES30 shelves.

---

**Note**

A Data Domain system with the DD Extended Retention software can use both ES20 and ES30 models of shelves. Different models cannot be mixed in the same shelf set, and the shelf sets must be balanced according to the configuration rules specified in the *EMC ES30 Shelf Hardware Guide*.

---

You can attach significantly more storage to the same controller, up to a maximum of 56 shelves on a DD990 with the DD Extended Retention feature. The active tier must include at least one shelf worth of storage. As one retention unit fills up, simply add an additional retention unit. For the best deduplication, each retention unit must be configured with the maximum storage allowed for a given system before adding another retention unit. DD Extended Retention allows incremental capacity growth for the long-term retention of data.

**Data Protection**

Data is protected with built-in fault isolation features, disaster recovery capability, and Data Invulnerability Architecture (DIA). DIA checks files when they are moved from the active to the retention tier. After data is copied into the retention tier, the container and file system structures are read back and verified. The location of the file is updated, and the space on the active tier is reclaimed only after the file is verified to have been correctly written to the retention tier.

When a retention unit is filled up, namespace information and system files are copied into it so that the data in the retention unit may be recovered even when other parts of the system are lost. A retention unit is completely filled up and sealed before another one is used.

---

**Note**

Encryption of data at rest, Sanitization, and some forms of replication are not supported for the DD OS 5.5 release.

---

**Reclaiming Space in the Retention Tier**

The data is moved from the active tier to the retention tier as set by your data movement policies. Data in the retention tier can be deleted. Prior to DD OS 5.2, you could not reclaim this freed-up space until all data on a retention unit was completely deleted, and then you had to reclaim the entire unit.

Starting with DD OS 5.3, you can reclaim this freed-up space. For instructions, see [Reclaiming Space in the Retention Tier on page 341](#).

Space reclamation runs in the background as a low-priority activity. It suspends itself when there are higher priority activities, such as data movement and cleaning. It reclaims space from deleted data across all units in the retention tier, migrates data between retention units to keep the older retention units full, and moves the free space to the target retention unit.

## Getting Started

The sections that follow explain the capabilities and requirements for using the Extended Retention feature. Read these descriptions before continuing to set up and configure the feature.

## Accessing Data

Systems with DD Extended Retention software support the protocols NFS, CIFS, DD Boost, and, as of DD OS 5.2, VTL. Support for NDMP was added in DD OS 5.3. For a list of applications supported with DD Boost, see the DD Boost Compatibility List on the Data Domain support site.

Data first lands in the active tier. Files are moved in their entirety into retention units in the retention tier, as specified by data movement policies. All files appear in the same namespace. There is no need to partition data, and you can continue to expand the file system as desired.

All data is visible to users, and all file system metadata is present in the active tier. All data is accessible via CIFS, NFS, DD Boost, VTL, and NDMP.

The trade-off in moving data to retention tiers is larger capacity versus slightly slower access time if the unit to be accessed is not currently ready for access.

## Supported Replication Types

---

### Note

An Extended Retention system cannot replicate to a non-Extended Retention system.

---

The supported replication types depend on the data to be protected:

- ◆ To protect data on the system as a source, Data Domain with DD Extended Retention supports collection replication, MTree replication, and DD Boost managed file replication.
  - ◆ To protect data from other systems as a destination, Data Domain with DD Extended Retention supports collection replication, directory replication, MTree replication, and DD Boost managed file replication.
- 

### Note

For instructions on configuring a replication type, see [Configuring Replication on page 310](#).

---

## Performing Collection Replication

Collection replication takes place between corresponding active tier and retention units. If the active tier or a retention unit at the source fails, the data can be copied from the corresponding unit at the remote site onto a new unit, which is shipped to your site as a replacement unit.

Prerequisites for setting up collection replication:

- ◆ Both the source and destination systems must be configured as controllers with DD Extended Retention enabled.
- ◆ The destination system must have the same number (or more) corresponding retention units as at the source.
- ◆ Each unit at the destination must be equivalent or larger in capacity than its corresponding source unit.
- ◆ The file system must not be enabled on the destination until the retention units have been added to it, and replication has been configured.
- ◆ When a retention unit is added to the source system, a corresponding retention unit needs to be added to the destination system to establish a collection replication pair.



Use the `replication status detailed` command to verify that the pairing is complete and that replication is progressing.

## Performing Directory Replication

With directory replication, the system with DD Extended Retention is used as a replication target and supports one-to-one and many-to-one topologies from any supported Data Domain system.

---

### Note

A system with DD Extended Retention does not support bi-directional directory replication.

A system with DD Extended Retention cannot be a source of directory replication.

---

To copy data into a system with DD Extended Retention using directory replication, the source has to be using DD OS 5.0 or a later version. Replicating from releases prior to DD OS 5.0 requires that the data first be imported into an intermediate Data Domain system running DD OS 5.0. For example, replication from a system that runs DD OS 4.9 is made into a non-DD Extended Retention Data Domain system that runs DD OS 5.0. Then replication is made from the DD OS 5.0 Data Domain system into the system with DD Extended Retention.

Complete the steps below to copy data into a system with DD Extended Retention using directory replication.

### Procedure

1. Set up directory replication into a directory within `/backup`.
  2. On the controller with DD Extended Retention enabled, perform a fast copy from `/backup` into the target MTree. A fast copy ensures that data on a system with DD Extended Retention is not deleted when the data on the source Data Domain system, which has a shorter retention period, is deleted. For more information, see [About Fast Copy Operations on page 329](#).
  3. Manage each MTree with per-MTree data movement policies.
- 

### Note

Ingest the data in age-order (oldest data first) to maximize deduplication and to improve performance.

---

## About Fast Copy Operations

A fast copy operation clones files and directory trees of a source directory to a target directory on a Data Domain system. The `force` option allows the destination directory to be overwritten if it exists. Executing the fast copy operation displays a progress status dialog box.

A fast copy operation makes the destination equal to the source, but not at a specific time. There are no guarantees that the two are or were ever equal if you change either folder during this operation.

## Performing MTree Replication

You can set up MTree replication between two controllers with DD Extended Retention software. Replicated data is first placed in the active tier on the destination system. The data movement policies on the destination system determine when this replicated data is moved to the retention tier.

As of DD OS 5.1 data can be replicated from a regular Data Domain system to one with DD Extended Retention using MTree replication.

As of DD OS 5.2 you can protect the data within the active tier of by replicating it to the active tier of another controller with DD Extended Retention.

As of DD OS 5.3 you can protect any data on a Data Domain system with DD Extended Retention by using MTree replication to another Data Domain system with DD Extended Retention. The destination system for MTree replication from a DD Extended Retention system cannot be a Data Domain system without Extended Retention, except for recovery purposes.

---

#### Note

With DD OS 5.3 and 5.4, when using Extended Retention do not set up replication for the /backup MTree on the source machine.

The following is recommended for an MTree replication configuration (either an MTree source or an MTree destination) and with a missing or deleted collection-partition.

Before using the file system, you must remove deleted and missing files from the namespace. Follow these steps:

#### Procedure

1. Identify the files that belong to the deleted or missing unit by entering:

```
# archive report generate file-location filename
```

---

#### Note

The `archive report` command is the only Extended Retention command that does not have a DD System Manager equivalent.

---

#### Note

The DD System Manager was formerly known as the Enterprise Manager.

2. Expire all snapshots that contain any of the identified files.
3. Remove identified file from MTree if MTree is not replica.

## DD Boost Managed File Replication

With DD Boost managed file replication, supported topologies are one-to-one, many-to-one, bi-directional, one-to-many, and cascaded.

With DD Boost 2.3 or later you can specify how multiple copies are to be made and managed within the backup application.

## Licenses

EMC Data Domain Extended Retention is a licensed software option installed on a supported EMC Data Domain controller. A separate shelf capacity license is needed for each storage shelf, for shelves installed in both the active tier and retention tier.

The appropriate shelf capacity license is required for any new shelf that you add. The shelf capacity license is specific to either an active or retention tier shelf. An Expanded-Storage license is required to expand the active tier storage capacity beyond the entry capacity, which varies by controller model:

- ◆ DD860: greater than 48 TB of usable capacity; for example, more than two shelves of 2 TB disks.

- ◆ DD990: greater than 285 TB of usable capacity; for example, more than 12 shelves of 2 TB disks or 8 shelves of 3 TB disks.
- ◆ DD4200: greater than 96 TB of usable capacity; for example, more than 4 shelves of 2 TB disks or 2 shelves of 3 TB disks.
- ◆ DD4500: greater than 142 TB of usable capacity; for example, more than 6 shelves of 2 TB disks or 4 shelves of 3 TB disks.
- ◆ DD7200: greater than 285 TB of usable capacity; for example, more than 12 shelves of 2 TB disks or 8 shelves of 3 TB disks.

You cannot use the additional storage without first applying the appropriate licenses.

## Using the DD System Manager

The DD System Manager incorporates the functionality of all of the `archive` commands, except for the `archive report` command. For `archive` command information, see the *EMC Data Domain Operating System Command Reference Guide*.

---

### Note

The DD System Manager was formerly known as the Enterprise Manager.

The sections that follow describe how to use DD System Manager to perform Extended Retention functions.

## File System Overview Panel for DD Extended Retention

The File System Overview panel displays the file system Clean Status, Data Movement Status, and Space Reclamation Status for Data Domain systems.

- ◆ Clean Status shows the time the last cleaning operation finished, or the current cleaning status if the cleaning operation is currently running. If cleaning can be run, it shows a **Start Cleaning** button. When cleaning is running, the **Start Cleaning** button changes to a **Stop Cleaning** button.
- ◆ Data Movement Status shows the time the last data movement finished. If data movement can be run, it shows a **Start** button. When data movement is running, the **Start** button changes to a **Stop** button.
- ◆ Space Reclamation Status shows the amount of space reclaimed after deleting data in the retention tier. If space reclamation can be run, it shows a **Start** button. If it is already running, it shows **Stop** and **Suspend** buttons. If it was running previously and was suspended, it shows **Stop** and **Resume** buttons.
- ◆ Selecting the **More Tasks > Expanded Capacity** menu item brings up a dialog where you can choose to expand the active or archive tier. If you choose the archive tier (also referred to as the retention tier), you can select the option to add a new archive unit (also referred to as a retention unit).

## The Data Management File System Tab

To access this tab, go to **Data Management > File System**.

### Summary Tab

The File System Summary tab displays information about disk space usage and compression for both the active and retention tiers.

**Disk Space Usage:** This panel shows the total size, amount of space used, and amount of available space and combined totals for active and retention tiers. The amount of cleanable space is shown for the active tier.

Compression: Shows the pre-compression and post-compression values that are currently used and those that have been written in the last 24 hours. Also shows the global, local, and total compression (reduction percentage) factors.

#### Archive Units Tab

The Archive Units tab on the File System page lists each retention unit. It shows the unit's state (new, sealed, target, or cleaning), its status (disabled, ready, or stand-by), and its size. A unit will be in the cleaning state if space reclamation is running on that unit. If the unit has been sealed, that is, no more data can be added, then the date that it was sealed is given. The Extended Retention start time is shown in the Start Date column.

There are two buttons: **Delete** (for deleting units) and **Expand** (for adding storage to a unit). Only units in the new or target states can be expanded.

Click the diamond symbol to the right of a column heading to sort the order of the values in reverse.

#### Configuration Tab

Clicking this option's **Edit** button allows you to change the Modify Settings: Local Compression setting. Compression options are none, lz, gz (the default), and gzfast.

The Data Movement Policy section allows you to set two parameters: File Age Threshold and Schedule. File Age Threshold is a system-wide default and applies to all MTrees which are using the default File Age Threshold value. Schedule allows setting when Data Movement will be run. The recommended schedule is every two weeks.

#### Space Usage Tab

You can select one of the three chart types to view space usage over time in MiB for the entire file system, or for an archive or active tier. Select one of the duration values in the upper-right of the tab. The data is presented (color-coded) as pre-compression written (blue), post-compression used (red), and the compression factor (black).

#### Consumption Tab

You can select one of the three chart types to view the amount of post-compression storage used and the compression ratio over time, which enables you to view consumption trends. An option allows you to display the post-compression storage against total system capacity.

#### Daily Written Tab

Shows the amount of data written per day for a selected duration (7, 30, 60, or 120 days). The data is presented (color-coded) in both graph and table format as pre-compression written (blue), post-compression used (red), and the compression factor (black).

#### Data Management MTree (Data Movement Policy)

Go to **Data Management > MTree** to access this tab.

The Data Movement Policy shows the current threshold for the age of a file before it is archived. Fourteen days is the default. To change the value, click **Edit**, enter the new value, and click **OK**.

## Data Domain Provided Hardware

The following controller models support the DD Extended Retention software option.

#### DD860

- ◆ 72 GB of RAM
- ◆ 1 - NVRAM IO module (1 GB)
- ◆ 3 - Quad-port SAS IO modules
- ◆ 2 - 1 GbE ports on the motherboard

- ◆ 0 to 2 - 1/10 GbE NIC IO cards for external connectivity
- ◆ 0 to 2 - Dual-Port FC HBA IO cards for external connectivity
- ◆ 0 to 2 - Combined NIC and FC cards
- ◆ 1 to 24 - ES20 or ES30 shelves (1 TB or 2 TB disks), not to exceed the system maximum usable capacity of 142 TB

If DD Extended Retention is enabled on the DD860, the maximum usable storage capacity of an active tier is 142 TB. The retention tier can consist of one or more retention units, each of which can have a maximum usable capacity of 142 TB. The active and retention tiers have a total usable storage capacity of 570 TB.

#### **DD990**

- ◆ 256 GB of RAM
- ◆ 1 - NVRAM IO module (2 GB)
- ◆ 4 - Quad-port SAS IO modules
- ◆ 2 - 1 GbE ports on the motherboard
- ◆ 0 to 4 - 1 GbE NIC IO cards for external connectivity
- ◆ 0 to 3 - 10 GbE NIC cards for external connectivity
- ◆ 0 to 3 - Dual-Port FC HBA cards for external connectivity
- ◆ 0 to 3 - Combined NIC and FC cards, not to exceed three of any one specific IO module
- ◆ 1 to 56 - ES20 or ES30 shelves (1, 2, or 3 TB disks), not to exceed the system maximum usable capacity of 570 TB

If DD Extended Retention is enabled on the DD990, the maximum usable storage capacity of the active tier is 570 TB. The retention tier can consist of one or more retention units, each of which can have a maximum usable capacity of 570 TB. The active and retention tiers have a total usable storage capacity of 1344 TB.

#### **DD4200**

- ◆ 128 GB of RAM
- ◆ 1 - NVRAM IO module (4 GB)
- ◆ 4 - Quad-port SAS IO modules
- ◆ 1 - 1 GbE port on the motherboard
- ◆ 0 to 6 - 1/10 GbE NIC cards for external connectivity
- ◆ 0 to 6 - Dual-Port FC HBA cards for external connectivity
- ◆ 0 to 6 - Combined NIC and FC cards, not to exceed four of any one specific IO module
- ◆ 1 to 16 - ES30 SAS shelves (2 or 3 TB disks), not to exceed the system maximum usable capacity of 192 TB. ES30 SATA shelves (1, 2, or 3 TB disks) are supported for system controller upgrades.

If DD Extended Retention is enabled on the DD4200, the maximum usable storage capacity of the active tier is 192 TB. The retention tier can consist of one or more retention units, each of which can have a maximum usable capacity of 192 TB. The active and retention tiers have a total usable storage capacity of 576 TB. External connectivity is supported for DD Extended Retention configurations up to 16 shelves.

#### **DD4500**

- ◆ 192 GB of RAM
- ◆ 1 - NVRAM IO module (4 GB)

- ◆ 4 - Quad-port SAS IO modules
- ◆ 1 - 1 GbE port on the motherboard
- ◆ 0 to 6 - 1/10 GbE NIC IO cards for external connectivity
- ◆ 0 to 6 - Dual-Port FC HBA cards for external connectivity
- ◆ 0 to 5 - Combined NIC and FC cards, not to exceed four of any one specific IO module
- ◆ 1 to 20 - ES30 SAS shelves (2 or 3 TB disks), not to exceed the system maximum usable capacity of 285 TB. ES30 SATA shelves (1 TB, 2 TB, or 3 TB) are supported for system controller upgrades.

If DD Extended Retention is enabled on the DD4500, the maximum usable storage capacity of the active tier is 285 TB. The retention tier can consist of one or more retention units, each of which can have a maximum usable capacity of 285 TB. The active and retention tiers have a total usable storage capacity of 1152 TB. External connectivity is supported for DD Extended Retention configurations up to 32 shelves.

### DD7200

- ◆ 128 GB of RAM for entry capacity; optional upgrade to 256 GB RAM for expanded capacity
- ◆ 1 - NVRAM IO module (4 GB)
- ◆ 4 - Quad-port SAS IO modules
- ◆ 1 - 1 GbE port on the motherboard
- ◆ 0 to 6 - 1/10 GbE NIC cards for external connectivity
- ◆ 0 to 6 - Dual-Port FC HBA cards for external connectivity
- ◆ 0 to 5 - Combined NIC and FC cards, not to exceed four of any one specific IO module
- ◆ 1 to 20 - ES30 SAS shelves (2 or 3 TB disks), not to exceed the system maximum usable capacity of 432 TB. ES30 SATA shelves (1 TB, 2 TB, or 3 TB) are supported for system controller upgrades.

If DD Extended Retention is enabled on the DD7200, the maximum usable storage capacity of the active tier is 432 TB. The retention tier can consist of one or more retention units, each of which can have a maximum usable capacity of 432 TB. The active and retention tiers have a total usable storage capacity of 1728 TB. External connectivity is supported for DD Extended Retention configurations up to 56 shelves.

## Customer-Provided Infrastructure

- ◆ Specifications and Site Requirements  
See the *EMC Data Domain Installation and Setup Guide* for your controller model.
- ◆ Rack Space and Interconnect Cabling  
See the *EMC DD Extended Retention Installation and Setup Guide* for your controller model.
- ◆ Racking and Cabling  
Rack the system with future expansion in mind. All of the shelves are attached to the single Data Domain system. It is recommended that the initial system configuration be racked as shown in the *EMC Data Domain Expansion Shelf Hardware Guide* for a system consisting of only ES20 shelves, or the *EMC ES30 Shelf Hardware Guide*, for a system with ES30 shelves.

## Initial Setup

See also [Using the DD System Manager on page 331](#).

The sections that follow provide procedures for setting up and configuring the Extended Retention feature. Complete these procedures before using the Extended Retention feature.

## Initial Configuration

You must perform Extended Retention software configuration procedures that follow before setting up data handling.

### Adding Shelf Capacity Licenses

Each storage shelf requires a shelf capacity license. This license is specific for either active or retention tier usage of the shelf. An expanded storage license to increase the size of the active tier might also be required.

To add shelf capacity licenses:

#### Procedure

1. Within the DD System Manager, open the **System Settings > Licenses** tab.
2. Click **Add Licenses**.
3. Enter one or more licenses, one per line.

---

#### Note

If there are any errors, a summary of the added licenses and those not added because of error are listed. Click **Fix Errors** and reenter the licenses keys. Click **Add**.

---

#### Results

The licenses for the Data Domain system are displayed in two groups:

- ◆ Software option licenses, required for options such as DD Extended Retention and DD Boost.
- ◆ Shelf Capacity Licenses, which displays shelf capacity in TiB, the shelf model, such as ES30, and the shelf's storage tier (active or archive).

To delete a license, select the license in the Licenses list and click **Delete Selected Licenses**. If prompted to confirm, read the warning and click **OK** to proceed.

### Configuring Storage

Additional storage requires the appropriate license or licenses, and the Data Domain system must have enough installed memory to support it. Error messages display if more licenses or memory is needed.

To configure storage:

#### Procedure

1. Within the DD System Manager, open the **Hardware > Storage** tab.
2. In the Overview tab, click **Configure Storage**.
3. In the Configure Storage tab, select the storage to be added from the Available Storage list.
4. Select the appropriate Tier Configuration (**Archive** or **Active**) from the menu. The active tier is analogous to a standard Data Domain system and should be sized similarly. The maximum amount of storage that can be added to the active tier depends on the DD controller used. See [Data Domain Provided Hardware on page 332](#).

---

**Note**

The two bars show the portion of licensed capacity used/remaining for each shelf model (ES20 and ES30).

---

5. Select the check box for the Shelf to be added.
  6. Click the **Add to Tier** button.
  7. Click **OK** to add the storage.
- 

**Note**

To remove an added shelf, select it in the Tier Configuration list and click **Remove from Configuration**. Click **OK**.

---

## Creating a Two-Tiered File System

Follow these steps to create a new file system:

**Procedure**

1. Add appropriate licenses as required.
  2. Within the DD System Manager, open the **Data Management > File System** tab.
  3. Select **Create file system** from the **More Tasks** menu. **Create file system** is not available as an option if a file system exists.
  4. Storage is classified in tiers—either active tier or archive tier (also referred to as the retention tier). Select the type of file system to be created: either an archive-capable file system (the default), or a non-archive file system (to implement DD Extended Retention at a later date).
  5. Click **Next**.
  6. Storage must be configured before the file system can be created. Click **Configure** in the File System Create dialog box.
  7. When the Configure Storage Status dialog box shows the disk enclosures addition has completed, click **Close**.
  8. The storage in the active tier is used to create the active file system tier. The storage in the archive tier (also referred to as retention tier) is used to create archive units (also referred to as retention units). Using the File System Create dialog box, you can choose to create an active tier and the first archive unit. Click **Next** to confirm.
- 

**Note**

To ensure optimal performance of a Data Domain system with the Extended Retention software option, you should not expand archive units in 1-shelf increments. You should also not wait till the archive unit is nearly full before expanding it (see [Expanding an Archive Unit on page 340](#)).

---

Each archive unit also is a fault-isolation domain. If a unit is damaged, only the files on that unit are affected. A larger unit size enhances deduplication and performance, but increases the amount of data compromised should a unit fail.

---

**Note**

You want the file system to be enabled (the default) after the file system has been created.

---

9. Select the size of the first archive unit from the menu.



10. Click **Next**.

A Summary page shows the size of the active and archive tiers in the new file system. You have the option of returning to the previous step to change the allocation size, or of cancelling the procedure.

11. Click **Finish** to create the file system.

The progress of each step is shown, and a progress bar monitors overall status.

## Configuring Data Movement

A file is moved from the active to the retention tier based on the date it was last modified. For data integrity, the entire file is moved at this time.

You can specify the data movement age threshold, such as nine weeks, after which data that has not been modified is to be moved from the active to the retention tier.

You can specify different age thresholds for each of your defined MTrees. An MTree is a subtree within the namespace that is a logical set of data for management purposes. For example, you might place financial data, emails, and engineering data in separate MTrees.

To take advantage of the space reclamation feature, introduced with DD OS 5.3, EMC recommends that you schedule data movement and file system cleaning on a bi-weekly basis. Also, update existing data movement schedules to bi-weekly.

Before changing the data movement schedule from weekly to bi-weekly, provision storage in the active tier to hold one additional week of data.

Schedule cleaning to run after data movement completes. Do not schedule cleaning separately.

## Data Movement Policy

Set the data movement age threshold to distinguish between data that is changing and data that is static and should be archived.

Segregate data with different degrees of change into different MTrees and set the data movement age threshold to move the data soon after the data stabilizes. For example, create MTree A for the daily-incremental backups and create MTree B for the weekly fulls. Set the age threshold for MTree A such that the data in MTree A is never moved and the age threshold for MTree B to one week.

You should not set an aggressive migration threshold for B, such as one day. At least one full backup from MTree B should always be retained on the active to obtain good de-duplication and performance for the new incoming backup.

---

### Note

Ensure that the active tier is large enough to hold the data that is changing.

In another example, the data cannot be separated into MTrees A and B. The retention period of daily-incremental backups is eight weeks and the retention period of weekly fulls is three years. In this case, set the age threshold to nine weeks. If you set it lower, you would be moving daily incremental data that is soon to be deleted.

## Starting or Stopping Data Movement on Demand

The File System page for a DD Extended Retention system shows the status of data as it is moved from the active to the retention tier. The target unit in the retention tier is the recipient of the data. The status includes when the data movement completed, the number of files copied, and the amount of data copied in GiB.

Clicking the Data Movement Status Start button starts the data movement based on your defined data movement policy.

#### Procedure

1. Go to **Data Management** › **File System**.
2. Click the **Start** button to the right of Data Movement Status.
3. The Start Data Movement dialog box warns that data is to be moved from the active to the retention tier as defined in your data movement policy, followed by a file system clean. Click **Finish** to start the data movement.

---

#### Note

If cleaning is already in progress, starting the data movement schedules the data movement to run after the clean completes.

The data movement status is shown in the File System tab. The **Start** button is replaced by a **Stop** button.

4. Clicking **Stop** stops the data movement. Click **OK** in the Stop Data Movement dialog box to confirm.

## Modifying the Data Movement Schedule

#### Procedure

1. Go to **Data Management** › **File System** and select the **Configuration** tab.
2. Click the Data Movement Policy's **Edit** button.
  - a. In the Data Movement Policy dialog box, specify the system-wide default age threshold value in number of days. (This applies to newly created MTrees and MTrees that have not been assigned a per-MTree age threshold value using the File Age Threshold per MTree link). When data movement starts, all files that have not been modified for the specified threshold number of days are moved from the active tier to the retention tier.
  - b. Set a schedule for when the data movement should begin; for example, bi-weekly, weekly, daily, specific days, and a time in hours and minutes.
 

EMC recommends that you schedule data movement and file system cleaning on a bi-weekly basis. To use the space reclamation feature, select bi-weekly and update existing data movement schedules to the bi-weekly basis.
  - c. A file system clean is recommended after the data movement. De-select this option if you do not want the clean to occur.

---

#### Note

Schedule cleaning to run after data movement completes. Do not schedule cleaning separately.

3. If desired, specify age threshold values for individual MTrees using the File Age Threshold per MTree link on the Configuration panel.

## Using Data Packing

As of DD OS 5.2, data is compacted in the target partition after every file migration. By default, this feature, which is called data movement packing, is enabled. When this feature enabled, the overall compression of the retention tier improves, but there is a slight increase in migration time.

To see if this feature is enabled, go to **Data Management** › **File System** › **Configuration**.

The current value for Packing data during archive data movement can be either Enabled or Disabled. Consult with a system engineer to enable or disable data movement packing.

## Administration

After setting up and configuring the Extended Retention feature, you can perform the ongoing administrative tasks described in the sections that follow.

### Avoiding Common Sizing Errors

Avoid these common sizing errors:

- ◆ Setting a data movement policy that is overly aggressive in which data is moved too soon.
- ◆ Setting a data movement policy that is too conservative. Once the active tier fills up, no more data can be written to the system.
- ◆ Configuring an undersized active tier, such that the active tier fills up prematurely.
- ◆ Creating an overly aggressive movement policy to compensate for an undersized active tier.

Space is not always reclaimed in the retention tier so moving files that are to be deleted or updated soon into the retention tier results in wasted space.

### Cleaning and Snapshots

Cleaning is performed on the active tier either as scheduled, or by default immediately after files have been moved from the active to the retention tier.

---

#### Note

Files in snapshots are not cleaned, even after they have been moved to the retention tier. The space cannot be reclaimed until the snapshots have been deleted. Set the retention for snapshots to less than two weeks.

---

### Expanding an Active or Archive Tier

When a file system exists, you can expand the active or archive tier (also referred to as the retention tier).

To expand the active tier, complete these steps:

#### Procedure

1. Select the Data Domain system and go to **Data Management > File System > More Tasks > Expand Capacity**.
2. In the Expand File System Capacity dialog box, select Expand File System Active Tier, then click **Next**.
3. In the Expand File System Capacity dialog, click **Configure**.
4. In the Configure Storage dialog, make sure that Active Tier is displayed as the Configure selection and click **OK**.
5. After the configuration completes, you are returned Expand File System Capacity dialog. Click **Finish** to complete the active tier expansion.

To expand the archive tier (also referred to as the retention tier), complete these steps:

#### Procedure

1. Select the Data Domain system and go to **Data Management > File System > More Tasks > Expand Capacity**.
2. In the Expand File System Capacity dialog box, select Expand File System Archive Tier, then click **Next**.
3. There are two methods to expand the archive tier: expand an existing archive unit (also referred to as a retention unit) or add a new archive unit. In the Expand File System Capacity dialog, select Add New Archive Unit or Expand on Existing Archive Unit then click **Next**.

---

#### Note

To ensure optimal performance of a Data Domain system with the Extended Retention software option, you should not expand archive units in 1-shelf increments. You should also not wait till the archive unit is nearly full before expanding it (see [Expanding an Archive Unit on page 340](#)).

- a. If you selected Add New Archive Unit, select a value for the size of the new archive unit and click **Configure**. In the Configure Storage dialog, make sure that Archive Tier is displayed as the Configure selection and click **OK**.
- b. If you selected Expand on Existing Archive Unit, select the archive unit you want to expand and click **Next**. Then select the size to expand the archive unit and click **Configure**.
4. After the configuration completes, you are returned Expand File System Capacity dialog. Click **Finish** to complete the archive tier expansion.

## Deleting a Retention Unit

Only empty retention units can be deleted.

#### Procedure

1. Disable the file system.
2. Go to **Data Management > File System** and select the Archive Units tab.
3. Select the retention (archive) unit from the list.
4. Click **Delete**.

## Expanding an Archive Unit

Only unsealed archive units (also referred to as retention units) can be expanded. You can expand a file system by expanding an existing unit, or by adding new units.

---

#### Note

Storage cannot be moved from the active tier to the archive tier (also referred to as the retention tier) after the file system has been created. Only unused enclosures can be added to the archive tier.

---

**Procedure**

1. Go to **Data Management > File System** and select the **Archive Units** tab.
2. Select the archive unit from the list.
3. Click **Expand**. (The **Expand** button is disabled if the unit is sealed.)
4. The dialog box that displays shows the current archive unit size and allows you to select the size to expand the selected unit.
5. To expand an existing unit, select an estimated size to increase the unit from the menu and click **Next**.

If there is insufficient configured storage to expand this unit, a message to this effect is displayed. Once the file system is expanded, however, you cannot revert it to its original size.

**Note**

To ensure optimal performance of a Data Domain system with the Extended Retention software option, you should not expand archive units in 1-shelf increments. You should also not wait till the archive unit is nearly full before expanding it.

6. To expand the file system using available non-configured storage, Click **Configure**.

## Reclaiming Space in the Retention Tier

You can reclaim space from the deleted data in the retention tier by running space reclamation.

**Procedure**

1. Go to **Data Management > File System**. Space Reclamation Status is one of the Statuses listed at the top.
2. Based on the current status of the Space Reclamation process, you can start, stop, suspend, or resume the process.
3. Click **More Information** for details about the current and last, if any, cycle of space reclamation.

The following is shown: cycle number; start and end times; effective run time; percent complete (if in progress), or if stopped, who stopped the process); units reclaimed; space freed on target unit, and total space freed.

## Upgrading Data Domain Systems for Extended Retention

The following Data Domain systems support the Extended Retention feature:

- ◆ DD860
- ◆ DD990
- ◆ DD4200
- ◆ DD4500
- ◆ DD7200

To upgrade these systems for Extended Retention, follow these steps:



**Once DD Extended Retention is enabled, it cannot be disabled without destroying the file system.**

**Procedure**

1. Ensure that the correct license is applied. Go to System Settings and check the Feature list for the Archiver license.
2. Users who want to upgrade an existing Data Domain system to a system with the Extended Retention software option can cause file system disruptions if they do not completely prepare the system for the Extended Retention services. It is not sufficient to disable encryption prior to installing the Extended Retention license. If encryption was ever enabled on the Data Domain system that will be converted to an Extended Retention system, contact your contracted EMC Data Domain support provider for detailed instructions on preparing an existing Data Domain system to operate with the Extended Retention software option.
3. Within the DD System Manager, go to **Data Management > File System**.
4. From the **More Tasks** menu, select **Enable Archive**. The Enable Archive menu option is available only if the file system has not already been configured for Extended Retention.
  - a. If the file system is already enabled (as a non-Extended Retention file system), you are prompted to disable it. Click **Disable** to do so.
  - b. If prompted to confirm that you want to convert the file system for use by DD Extended Retention, click **OK**.
 

After a file system is converted into an Extended Retention file system, the file system page is refreshed to include information about both tiers, and there is a new tab labeled Archive Units.
5. Go to **Data Management > File System > Summary**. Verify the Disk Space Usage summary and Compression statistics for the active and retention tiers.
6. Add storage (Extended Retention shelves) to the retention tier (see [Expanding an Active or Archive Tier on page 339](#)).
7. Set up data movement policies for MTrees to be migrated to the retention tier (see [Data Movement Policy on page 337](#)).
8. Set up the data movement schedule (see [Modifying the Data Movement Schedule on page 338](#)).

## Changing Retention Tier Compression

Within the DD System Manager, open the Data Management > File System tab and select the Summary tab.

### Modifying Retention Tier Compression

**Procedure**

1. Go to **Data Management > File System** and select the **Configuration** tab.
2. Click the **Options Edit** button.
3. Select one of the compression options from the Retention Tier Local Compression Type menu, and click **OK**.

The default is gz, which is a zip-style compression that uses the least amount of space for data storage (10% to 20% less than lz on average; however, some data sets achieve much higher compression).

See [Change Local Compression on page 144](#) for a description of other compression options.

## Upgrades and Recovery

The sections that follow describe how to perform software and hardware upgrades and how to recover data and reuse empty retention units.

### Upgrading to DD OS 5.4

When upgrading a Data Domain system with the DD Extended Retention software option to DD OS 5.4, update existing data movement schedules to bi-weekly to take advantage of the space reclamation feature.

As mentioned, you need to schedule cleaning to run after data movement completes. Do not schedule cleaning separately.

### Upgrading a Data Domain System Controller with the DD Extended Retention Software Option Enabled

The upgrade policy is the same as for a standard Data Domain controller. Upgrading from up to two major prior releases is supported.

---

#### Note

Consult with your contracted service provider and refer to the instructions in the appropriate *System Controller Upgrade Guide*.

---

If the active tier is not available, the upgrade process upgrades the system chassis and places it into the state where it is ready to create or accept a file system. Otherwise, the process proceeds to upgrade the active tier and all available retention units, and puts the system into the state that the previous upgrade has not been verified to be complete. This state is cleared by the file system after the file system is enabled and has verified that all retention units belonging to the file system have been upgraded. A subsequent upgrade is not permitted until this state is cleared.

If a retention unit becomes available after the upgrade process has finished, the unit is automatically upgraded when it is plugged into the system, or at the next system start.

## Replacing Data Domain Systems

You can replace a Data Domain system that has DD Extended Retention software with another Data Domain system that has DD Extended Retention software. For example, you could replace a DD860 with DD Extended Retention software with a DD990 with DD Extended Retention software.

The following conditions must be met for replacement:

- ◆ The new system controller must not already own an active tier or retention unit.
- ◆ The active tier and all retention units to be owned by the new system controller must be at the same version. Any active tier or retention unit that is at a different version must not be connected to the new system controller until the replacement process is completed.

A system controller replacement affects DD Extended Retention as follows:

- ◆ If the new system controller is at a more recent version than the active tier and retention units, the active tier and retention units are upgraded to that system controller's version. Otherwise, the new system controller is upgraded to the version in the active tier and retention units.

- ◆ The active tier and retention units that are connected to the new system controller become owned by the new system controller. All visible active tier and retention units are attached to the new system controller.
- ◆ If there is an active tier, the registry in the active tier is installed in the new system controller. Otherwise, the registry in the retention unit with the most recently updated registry is installed in the new system controller.

## Replication Recovery

The replication recovery procedure depends on the replication type. You can use any of these methods to recover data.

- ◆ Collection replication. For more information, see [Using DD Extended Retention with Collection Replication on page 344](#).
- ◆ MTree replication, for more information, see [MTree Replication on page 300](#).
- ◆ DD Boost managed file replication. For more information, see the *EMC Data Domain Boost for OpenStorage Administration Guide*.

### Using DD Extended Retention with Collection Replication

---

#### Note

If you need to recover only a portion of the system, such as one retention unit, from the collection replica, contact Data Domain Support.

---

The new source must be configured as a DD Extended Retention system with the same number (or more) retention units as are located at the destination. The file system must not be enabled on the new source until the retention units have been added to it and replication recovery has been initiated.

#### Procedure

1. Install the replication license on the new source.

```
# license add replication-license-key
```

2. Reset the authentication key on the destination.

```
# replication reauth col://hostB
```

3. Reconfigure replication on both the new source and destination.

```
# replication add source col://hostC destination col://hostB
```

4. Initiate recovery on the new source. The file system must not have been enabled on the new source before this step.

```
# replication recover col://hostB
```

5. Check the replication status.

```
# replication status
```

## Recovering a System with the DD Extended Retention Software Option Enabled

If the active tier and a subset of the retention units are lost and there is no replica available, Data Domain Support may be able to reconstitute any remaining sealed retention units into a new controller.

A system with the DD Extended Retention software option is designed to remain available to service read and write requests when one or more retention units are lost. The file system may not detect that a retention unit is lost until the file system restarts or tries to access data stored in the retention unit. The latter circumstance may trigger a file system



restart. After the file system has detected that a retention unit is lost, it returns an error in response to requests for data stored in that unit.

If the lost data cannot be recovered from a replica, Data Domain Support might be able to clean up the system by deleting the lost Retention Unit and any files that reside fully or partially in it.

## Recovering from System Failures

A system with DD Extended Retention is equipped with tools to address failures in different parts of the system. In a failure situation, perform recovery actions in the order listed below.

To recover from system failures:

### Procedure

1. Restore connection between the system controller and storage. If the system controller is lost, replace it with a new system controller.
2. If there is loss of data and a replica is available, try to recover the data from the replica.

If a replica is not available, limit any loss of data by leveraging the fault isolation features of the DD Extended Retention through Data Domain Support.

## Reusing a Retention Unit

If all of the files on a retention unit are no longer needed, deleting them makes the unit available for reuse. You can generate a file location report to make sure that the retention unit is indeed empty, delete the retention unit, then add it as a new retention unit.

