

Dell EMC Data Domain[®] Operating System

Version 6.1

Administration Guide

302-003-761

REV. 05

Copyright © 2010-2019 Dell Inc. or its subsidiaries All rights reserved.

Published February 2019

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

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

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

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

CONTENTS

	Preface	15
Chapter 1	Data Domain System Features and Integration	19
	Revision history.....	20
	Data Domain system overview.....	21
	Data Domain system features.....	22
	Data integrity.....	22
	Data deduplication.....	23
	Restore operations.....	23
	Data Domain Replicator.....	23
	Multipath and load balancing.....	23
	High Availability.....	24
	Random I/O handling.....	25
	System administrator access.....	26
	Licensed features.....	26
	Storage environment integration.....	28
Chapter 2	Getting Started	31
	DD System Manager overview.....	32
	Logging in and out of DD System Manager.....	32
	Logging in using a certificate.....	33
	The DD System Manager interface.....	34
	Page elements.....	34
	Banner.....	34
	Navigation panel.....	35
	Information panel.....	35
	Footer.....	35
	Help buttons.....	36
	End User License Agreement.....	36
	Configuring the system with the configuration wizard.....	36
	License page.....	37
	Network.....	37
	File System.....	39
	System Settings.....	43
	DD Boost protocol.....	45
	CIFS protocol.....	46
	NFS protocol.....	47
	DD VTL protocol.....	47
	Data Domain Command Line Interface.....	49
	Logging into the CLI.....	49
	CLI online help guidelines.....	50
Chapter 3	Managing Data Domain Systems	53
	System management overview.....	54
	HA system management overview.....	54
	HA system planned maintenance.....	55
	Rebooting a system.....	55
	Powering a system on or off	55

Power a system on.....	56
System upgrade management.....	57
Viewing upgrade packages on the system.....	58
Obtaining and verifying upgrade packages.....	58
Upgrade considerations for HA systems.....	59
Upgrading a Data Domain system.....	59
Removing an upgrade package.....	61
Managing electronic licenses.....	61
HA system license management.....	61
System storage management.....	61
Viewing system storage information.....	62
Physically locating an enclosure.....	66
Physically locating a disk.....	67
Configuring storage.....	67
DD3300 capacity expansion.....	68
Fail and unfail disks.....	69
Network connection management.....	69
HA system network connection management.....	69
Network interface management.....	70
General network settings management.....	85
Network route management.....	88
System passphrase management.....	91
Setting the system passphrase.....	92
Changing the system passphrase.....	92
System access management.....	93
Role-based access control.....	93
Access management for IP protocols.....	95
Local user account management.....	102
Directory user and group management.....	109
Change system authentication method.....	116
Configuring mail server settings.....	117
Managing time and date settings.....	117
Managing system properties.....	118
SNMP management.....	119
Viewing SNMP status and configuration.....	119
Enabling and disabling SNMP.....	121
Downloading the SNMP MIB.....	121
Configuring SNMP properties.....	121
SNMP V3 user management.....	122
SNMP V2C community management.....	123
SNMP trap host management.....	126
Autosupport report management.....	127
HA system autosupport and support bundle manageability.....	128
Enabling and disabling autosupport reporting to Data Domain.....	128
Reviewing generated autosupport reports.....	128
Configuring the autosupport mailing list.....	129
Support bundle management.....	130
Generating a support bundle.....	130
Viewing the support bundles list.....	130
Coredump management.....	131
Alert notification management.....	131
HA system alert notification management.....	132
Viewing the notification group list.....	132
Creating a notification group.....	134
Managing the subscriber list for a group.....	134
Modifying a notification group.....	135

	Deleting a notification group.....	136
	Resetting the notification group configuration.....	136
	Configuring the daily summary schedule and distribution list.....	136
	Enabling and disabling alert notification to Data Domain.....	138
	Testing the alerts email feature.....	138
	Support delivery management.....	139
	Selecting standard email delivery to Data Domain.....	139
	Selecting and configuring Secure Remote Services delivery.....	139
	Testing ConnectEMC operation.....	140
	Log file management.....	140
	Viewing log files in DD System Manager.....	141
	Displaying a log file in the CLI.....	142
	Learning more about log messages.....	142
	Saving a copy of log files.....	143
	Log message transmission to remote systems.....	144
	Remote system power management with IPMI.....	145
	IPMI and SOL limitations.....	146
	Adding and deleting IPMI users with DD System Manager.....	146
	Changing an IPMI user password.....	147
	Configuring an IPMI port.....	147
	Preparing for remote power management and console monitoring with the CLI.....	148
	Managing power with DD System Manager.....	150
	Managing power with the CLI.....	151
Chapter 4	Monitoring Data Domain Systems	153
	Viewing individual system status and identity information.....	154
	Dashboard Alerts area.....	154
	Dashboard File System area.....	155
	Dashboard Services area.....	155
	Dashboard HA Readiness area.....	156
	Dashboard Hardware area.....	156
	Maintenance System area.....	156
	Health Alerts panel.....	157
	Viewing and clearing current alerts.....	157
	Current Alerts tab.....	158
	Viewing the alerts history.....	158
	Alerts History tab.....	159
	Viewing hardware component status.....	159
	Fan status.....	160
	Temperature status.....	160
	Management panel status.....	161
	SSD status (DD6300 only).....	161
	Power supply status.....	162
	PCI slot status.....	162
	NVRAM status.....	162
	Viewing system statistics.....	163
	Performance statistics graphs.....	163
	Viewing active users.....	164
	History report management.....	164
	Types of reports.....	165
	Viewing the Task Log.....	169
	Viewing the system High Availability status.....	170
	High Availability status.....	170

Chapter 5	File System	173
	File system overview.....	174
	How the file system stores data.....	174
	How the file system reports space usage.....	174
	How the file system uses compression	174
	How the file system implements data integrity.....	176
	How the file system reclaims storage space with file system cleaning.....	176
	Supported interfaces	177
	Supported backup software.....	177
	Data streams sent to a Data Domain system	177
	File system limitations.....	180
	Monitoring file system usage.....	181
	Accessing the file system view.....	181
	Managing file system operations.....	189
	Performing basic operations.....	189
	Performing cleaning.....	191
	Performing sanitization.....	192
	Modifying basic settings.....	194
	Fast copy operations.....	196
	Performing a fast copy operation.....	196
Chapter 6	MTrees	199
	MTrees overview.....	200
	MTree limits.....	200
	Quotas.....	201
	About the MTree panel.....	201
	About the summary view.....	202
	About the space usage view (MTrees).....	206
	About the daily written view (MTrees).....	206
	Monitoring MTree usage.....	207
	Understanding physical capacity measurement.....	207
	Managing MTree operations.....	210
	Creating an MTree.....	210
	Configure and enable/disable MTree quotas.....	212
	Deleting an MTree.....	213
	Undeleting an MTree.....	213
	Renaming an MTree.....	213
Chapter 7	Snapshots	215
	Snapshots overview.....	216
	Monitoring snapshots and their schedules.....	217
	About the snapshots view.....	217
	Managing snapshots.....	218
	Creating a snapshot.....	218
	Modifying a snapshot expiration date.....	219
	Renaming a snapshot.....	219
	Expiring a snapshot.....	219
	Managing snapshot schedules.....	220
	Creating a snapshot schedule.....	220
	Modifying a snapshot schedule.....	221
	Deleting a snapshot schedule.....	222
	Recover data from a snapshot.....	222

Chapter 8	CIFS	223
	CIFS overview.....	224
	Configuring SMB signing.....	224
	Performing CIFS setup.....	225
	HA systems and CIFS.....	225
	Preparing clients for access to Data Domain systems.....	225
	Enabling CIFS services.....	225
	Naming the CIFS server.....	226
	Setting authentication parameters.....	226
	Disabling CIFS services.....	227
	Working with shares.....	227
	Creating shares on the Data Domain system.....	227
	Modifying a share on a Data Domain system.....	229
	Creating a share from an existing share.....	229
	Disabling a share on a Data Domain system.....	230
	Enabling a share on a Data Domain system.....	230
	Deleting a share on a Data Domain system.....	230
	Performing MMC administration.....	230
	Connecting to a Data Domain system from a CIFS client.....	230
	Displaying CIFS information	232
	Managing access control.....	232
	Accessing shares from a Windows client.....	233
	Providing domain users administrative access.....	233
	Allowing administrative access to a Data Domain system for domain users.....	233
	Restricting administrative access from Windows.....	234
	File access.....	234
	Monitoring CIFS operation.....	237
	Displaying CIFS status.....	237
	Display CIFS configuration.....	238
	Displaying CIFS statistics.....	240
	Performing CIFS troubleshooting.....	240
	Displaying clients current activity.....	240
	Setting the maximum open files on a connection.....	241
	Data Domain system clock.....	241
	Synchronizing from a Windows domain controller.....	242
	Synchronize from an NTP server.....	242
Chapter 9	NFS	243
	NFS overview.....	244
	HA systems and NFS.....	245
	Managing NFS client access to the Data Domain system.....	245
	Enabling NFS services.....	245
	Disabling NFS services.....	245
	Creating an export.....	245
	Modifying an export.....	247
	Creating an export from an existing export.....	248
	Deleting an export.....	248
	Displaying NFS information.....	249
	Viewing NFS status.....	249
	Viewing NFS exports.....	249
	Viewing active NFS clients.....	249
	Integrating a DDR into a Kerberos domain.....	250
	Add and delete KDC servers after initial configuration.....	251

Chapter 10	NFSv4	255
	Introduction to NFSv4.....	256
	NFSv4 compared to NFSv3 on Data Domain systems.....	256
	NFSv4 ports.....	257
	ID Mapping Overview.....	257
	External formats.....	257
	Standard identifier formats.....	257
	ACE extended identifiers.....	258
	Alternative formats.....	258
	Internal Identifier Formats.....	258
	When ID mapping occurs.....	259
	Input mapping.....	259
	Output mapping.....	259
	Credential mapping.....	260
	NFSv4 and CIFS/SMB Interoperability.....	260
	CIFS/SMB Active Directory Integration.....	261
	Default DACL for NFSv4.....	261
	System Default SIDs.....	261
	Common identifiers in NFSv4 ACLs and SIDs.....	261
	NFS Referrals.....	261
	Referral Locations.....	262
	Referral location names.....	262
	Referrals and Scaleout Systems.....	262
	NFSv4 and High Availability.....	263
	NFSv4 Global Namespaces.....	263
	NFSv4 global namespaces and NFSv3 submounts.....	263
	NFSv4 Configuration.....	264
	Enabling the NFSv4 Server.....	264
	Setting the default server to include NFSv4.....	265
	Updating existing exports.....	265
	Kerberos and NFSv4.....	265
	Configuring Kerberos with a Linux-Based KDC.....	266
	Configuring the Data Domain System to Use Kerberos Authentication.....	267
	Configuring Clients.....	267
	Enabling Active Directory.....	268
	Configuring Active Directory.....	269
	Configuring clients on Active Directory.....	269
	LDAP and NFSv4.....	270
	Configure LDAP servers.....	270
	Configure the LDAP base suffix.....	271
	Configure LDAP client authentication.....	271
	Enable LDAP.....	272
	Enable secure LDAP.....	273
	Configure LDAP server certificate verification with imported CA certificates.....	273
	Manage CA certificates for LDAP.....	274
Chapter 11	Storage Migration	275
	Storage migration overview.....	276
	Migration planning considerations.....	276
	DS60 shelf considerations.....	278
	Viewing migration status.....	278
	Evaluating migration readiness.....	279
	Migrating storage using DD System Manager.....	279

	Storage migration dialog descriptions.....	280
	Select a Task dialog.....	280
	Select Existing Enclosures dialog.....	280
	Select New Enclosures dialog.....	281
	Review Migration Plan dialog.....	281
	Verify Migration Preconditions dialog.....	281
	Migration progress dialogs.....	282
	Migrating storage using the CLI.....	283
	CLI storage migration example.....	284
Chapter 12	Metadata on Flash	289
	Overview of Metadata on Flash (MDoF)	290
	MDoF licensing and capacity.....	290
	SSD cache tier.....	291
	MDoF SSD cache tier - system management	291
	Managing the SSD cache tier.....	292
	SSD alerts.....	295
Chapter 13	SCSI Target	297
	SCSI Target overview.....	298
	Fibre Channel view.....	299
	Enabling NPIV.....	299
	Disabling NPIV.....	302
	Resources tab.....	302
	Access Groups tab.....	309
	Differences in FC link monitoring among DD OS versions.....	309
Chapter 14	Working with DD Boost	311
	About Data Domain Boost.....	312
	Managing DD Boost with DD System Manager.....	313
	Specifying DD Boost user names.....	313
	Changing DD Boost user passwords.....	314
	Removing a DD Boost user name.....	314
	Enabling DD Boost.....	314
	Configuring Kerberos.....	314
	Disabling DD Boost.....	315
	Viewing DD Boost storage units.....	315
	Creating a storage unit.....	316
	Viewing storage unit information.....	318
	Modifying a storage unit.....	320
	Renaming a storage unit.....	321
	Deleting a storage unit.....	322
	Undeleting a storage unit.....	322
	Selecting DD Boost options.....	322
	Managing certificates for DD Boost.....	324
	Managing DD Boost client access and encryption.....	326
	About interface groups.....	327
	Interfaces.....	329
	Clients.....	329
	Creating interface groups.....	330
	Enabling and disabling interface groups.....	331
	Modifying an interface group's name and interfaces.....	331
	Deleting an interface group.....	332
	Adding a client to an interface group.....	332

	Modifying a client's name or interface group.....	332
	Deleting a client from the interface group.....	333
	Using interface groups for Managed File Replication (MFR).....	333
	Destroying DD Boost.....	335
	Configuring DD Boost-over-Fibre Channel.....	335
	Enabling DD Boost users.....	336
	Configuring DD Boost.....	337
	Verifying connectivity and creating access groups.....	338
	Using DD Boost on HA systems.....	340
	About the DD Boost tabs.....	340
	Settings.....	341
	Active Connections.....	341
	IP Network.....	342
	Fibre Channel.....	342
	Storage Units.....	342
Chapter 15	DD Virtual Tape Library	345
	DD Virtual Tape Library overview.....	346
	Planning a DD VTL.....	346
	DD VTL limits.....	347
	Number of drives supported by a DD VTL.....	350
	Tape barcodes.....	351
	LTO tape drive compatibility.....	352
	Setting up a DD VTL.....	352
	HA systems and DD VTL.....	352
	DD VTL tape out to cloud.....	353
	Managing a DD VTL.....	353
	Enabling DD VTL.....	354
	Disabling DD VTL.....	355
	DD VTL option defaults.....	355
	Configuring DD VTL default options.....	356
	Working with libraries.....	357
	Creating libraries.....	357
	Deleting libraries.....	359
	Searching for tapes.....	359
	Working with a selected library.....	360
	Creating tapes.....	361
	Deleting tapes.....	362
	Importing tapes.....	363
	Exporting tapes.....	364
	Moving tapes between devices within a library.....	365
	Adding slots.....	366
	Deleting slots.....	367
	Adding CAPs.....	367
	Deleting CAPs.....	367
	Viewing changer information.....	368
	Working with drives.....	368
	Creating drives.....	369
	Deleting drives.....	370
	Working with a selected drive.....	371
	Working with tapes.....	371
	Changing a tape's write or retention lock state.....	372
	Working with the vault.....	373
	Working with the cloud-based vault.....	373
	Prepare the VTL pool for data movement.....	374

	Remove tapes from the backup application inventory.....	376
	Select tape volumes for data movement.....	376
	Restore data held in the cloud.....	378
	Manually recall a tape volume from cloud storage.....	378
	Working with access groups.....	380
	Creating an access group.....	380
	Deleting an access group.....	384
	Working with a selected access group.....	384
	Selecting endpoints for a device.....	385
	Configuring the NDMP device TapeServer group.....	385
	Working with resources.....	386
	Working with initiators.....	387
	Working with endpoints.....	388
	Working with a selected endpoint.....	389
	Working with pools.....	391
	Creating pools.....	392
	Deleting pools.....	393
	Working with a selected pool.....	393
	Converting a directory pool to an MTree pool	395
	Moving tapes between pools.....	396
	Copying tapes between pools.....	397
	Renaming pools.....	398
Chapter 16	DD Replicator	399
	DD Replicator overview.....	400
	Prerequisites for replication configuration.....	401
	Replication version compatibility.....	403
	Replication types.....	407
	Managed file replication	408
	Directory replication.....	408
	MTree replication.....	409
	Collection replication	411
	Using DD Encryption with DD Replicator.....	412
	Replication topologies.....	413
	One-to-one replication.....	414
	Bi-directional replication.....	415
	One-to-many replication.....	415
	Many-to-one replication.....	416
	Cascaded replication.....	417
	Managing replication.....	418
	Replication status.....	418
	Summary view.....	418
	DD Boost view.....	428
	Topology view.....	430
	Performance view.....	430
	Advanced Settings view.....	430
	Monitoring replication	433
	Checking replication pair status.....	433
	Viewing estimated completion time for backup jobs.....	434
	Checking replication context performance.....	434
	Tracking status of a replication process.....	434
	Replication with HA.....	434
	Replicating a system with quotas to one without.....	435
	Replication Scaling Context	435
	Directory-to-MTree replication migration.....	436

	Performing migration from directory replication to MTree replication.....	436
	Viewing directory-to-MTree migration progress.....	437
	Checking the status of directory-to-MTree replication migration....	438
	Aborting D2M replication	438
	Troubleshooting D2M.....	439
	Additional D2M troubleshooting.....	440
	Using collection replication for disaster recovery with SMT.....	440
Chapter 17	DD Secure Multitenancy	443
	Data Domain Secure Multi-tenancy overview.....	444
	SMT architecture basics.....	444
	Terminology used in Secure Multi-Tenancy (SMT).....	444
	Control path and network isolation.....	445
	Understanding RBAC in SMT.....	446
	Provisioning a Tenant Unit.....	447
	Enabling Tenant Self-Service mode.....	451
	Data access by protocol.....	451
	Multi-User DD Boost and Storage Units in SMT.....	452
	Configuring access for CIFS.....	452
	Configuring NFS access.....	452
	Configuring access for DD VTL.....	453
	Using DD VTL NDMP TapeServer	453
	Data management operations.....	453
	Collecting performance statistics.....	453
	Modifying quotas.....	454
	SMT and replication.....	454
	SMT Tenant alerts.....	455
	Managing snapshots.....	456
	Performing a file system Fast Copy.....	456
Chapter 18	DD Cloud Tier	457
	DD Cloud Tier overview.....	458
	Supported platforms.....	458
	DD Cloud Tier performance.....	460
	Configuring Cloud Tier.....	460
	Configuring storage for DD Cloud Tier.....	460
	Configuring cloud units.....	461
	Firewall and proxy settings.....	461
	Importing CA certificates.....	462
	Adding a cloud unit for Elastic Cloud Storage (ECS).....	463
	Adding a cloud unit for Virtustream.....	464
	Adding a cloud unit for Amazon Web Services S3.....	464
	Adding a cloud unit for Azure.....	466
	Adding an S3 Flexible provider cloud unit.....	467
	Modifying a cloud unit or cloud profile.....	467
	Deleting a cloud unit.....	469
	Data movement.....	470
	Adding data movement policies to MTrees.....	470
	Moving data manually.....	470
	Moving data automatically.....	471
	Recalling a file from the Cloud Tier.....	471
	Using the CLI to recall a file from the cloud tier.....	472

	Direct restore from the cloud tier.....	474
	Using the Command Line Interface (CLI) to configure DD Cloud Tier.....	474
	Configuring encryption for DD cloud units.....	478
	Information needed in the event of system loss.....	478
	Using DD Replicator with Cloud Tier.....	479
	Using DD Virtual Tape Library (VTL) with Cloud Tier.....	479
	Displaying capacity consumption charts for DD Cloud Tier.....	479
	DD Cloud Tier logs.....	480
	Using the Command Line Interface (CLI) to remove DD Cloud Tier.....	480
Chapter 19	DD Extended Retention	483
	DD Extended Retention overview.....	484
	Supported protocols in DD Extended Retention.....	486
	High Availability and Extended Retention.....	486
	Using DD Replicator with DD Extended Retention.....	486
	Collection replication with DD Extended Retention.....	486
	Directory replication with DD Extended Retention.....	487
	MTree replication with DD Extended Retention.....	487
	Managed file replication with DD Extended Retention.....	487
	Hardware and licensing for DD Extended Retention.....	488
	Hardware supported for DD Extended Retention.....	488
	Licensing for DD Extended Retention.....	491
	Adding shelf capacity licenses for DD Extended Retention.....	491
	Configuring storage for DD Extended Retention.....	492
	Customer-provided infrastructure for DD Extended Retention....	492
	Managing DD Extended Retention.....	492
	Enabling DD systems for DD Extended Retention.....	493
	Creating a two-tiered file system for DD Extended Retention.....	494
	File system panel for DD Extended Retention.....	495
	File system tabs for DD Extended Retention.....	497
	Upgrades and recovery with DD Extended Retention.....	502
	Upgrading to DD OS 5.7 with DD Extended Retention.....	502
	Upgrading hardware with DD Extended Retention.....	503
	Recovering a DD Extended Retention-enabled system.....	503
Chapter 20	DD Retention Lock	505
	DD Retention Lock overview.....	506
	DD Retention Lock protocol.....	507
	DD Retention Lock flow.....	507
	Supported data access protocols.....	508
	Enabling DD Retention Lock on an MTree.....	509
	Enabling DD Retention Lock Governance on an MTree.....	509
	Enabling DD Retention Lock Compliance on an MTree.....	510
	Client-Side Retention Lock file control.....	512
	Setting Retention Locking on a file.....	513
	Extending Retention Locking on a file.....	515
	Identifying a Retention-Locked file.....	516
	Specifying a directory and touching only those files.....	516
	Reading a list of files and touching only those files.....	516
	Deleting or expiring a file.....	517
	Using ctime or mtime on Retention-Locked files.....	517
	System behavior with DD Retention Lock.....	517
	DD Retention Lock governance.....	518
	DD Retention Lock compliance.....	519

Chapter 21	DD Encryption	531
	DD encryption overview.....	532
	Configuring encryption.....	533
	About key management.....	533
	Rectifying lost or corrupted keys.....	534
	Key manager support.....	535
	Working with the RSA DPM Key Manager.....	535
	Working with the Embedded Key Manager.....	538
	Working with KeySecure Key Manager.....	539
	Using DD System Manager to set up and manage the KeySecure Key Manager.....	539
	Using the Data Domain CLI to manage the KeySecure Key Manager	541
	How the cleaning operation works.....	545
	Key manager setup.....	545
	RSA DPM Key Manager encryption setup.....	545
	Setting up KMIP key manager.....	548
	Changing key managers after setup.....	550
	Managing certificates for RSA Key Manager.....	550
	Checking settings for encryption of data at rest.....	551
	Enabling and disabling encryption of data at rest.....	552
	Enabling encryption of data at rest.....	552
	Disabling encryption of data at rest.....	552
	Locking and unlocking the file system.....	553
	Locking the file system.....	553
	Unlocking the file system.....	554
	Changing the encryption algorithm.....	554

Preface

As part of an effort to improve its product lines, Data Domain 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, software updates, software compatibility guides, and information about Data Domain products, licensing, and service.

Contact your 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 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 Data Domain® systems with an emphasis on procedures using the 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 *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, *Data Domain DD9300 System Installation Guide*
- *Data Domain Hardware Features and Specifications Guide*
- *Data Domain Operating System USB Installation Guide*
- *Data Domain Operating System DVD Installation Guide*
- *Data Domain Operating System Release Notes*
- *Data Domain Operating System Initial Configuration Guide*
- *Data Domain Security Configuration Guide*
- *Data Domain Operating System High Availability White Paper*
- *Data Domain Operating System Command Reference Guide*

- *Data Domain Operating System MIB Quick Reference*
- *Data Domain Operating System Offline Diagnostics Suite User's Guide*
- 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*
- *Data Domain, System Controller Upgrade Guide*
- *Data Domain Expansion Shelf, Hardware Guide* (for shelf model ES30/FS15, or DS60)
- *Data Domain Boost for Partner Integration Administration Guide*
- *Data Domain Boost for OpenStorage Administration Guide*
- *Data Domain Boost for Oracle Recovery Manager Administration Guide*
- *Statement of Volatility for the Data Domain DD2500 System*
- *Statement of Volatility for the Data Domain DD4200, DD4500, or DD7200 System*
- *Statement of Volatility for the Data Domain DD6300, DD6800, or DD9300 System*
- *Statement of Volatility for the Data Domain DD9500 or DD9800 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

Data Domain uses the following conventions for special notices:

NOTICE

A notice identifies content that warns of a potential business or data loss.

Note

A note identifies information that is incidental, but not essential, to the topic. Notes can provide an explanation, a comment, reinforcement of a point in the text, or just a related point.

Typographical conventions

Data Domain uses the following type style conventions in this document:

Table 1 Typography

Bold	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"> • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Highlights a variable name that must be replaced with a variable value

Table 1 Typography (continued)

<code>Monospace bold</code>	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

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

Product information

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

Technical support

Go to Online Support and click Service Center. You will see several options for contacting Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your 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.

CHAPTER 1

Data Domain System Features and Integration

This chapter includes:

- [Revision history](#)20
- [Data Domain system overview](#) 21
- [Data Domain system features](#)..... 22
- [Storage environment integration](#).....28

Revision history

The revision history lists the major changes to this document to support DD OS Release 6.1.

Table 2 Document revision history

Revision	Date	Description
05 (6.1.2)	February 2019	This revision includes corrections to the steam counts for DD2200 systems with 8 GB of memory.
04 (6.1.2)	July 2018	<p>This revision includes information about these new features:</p> <ul style="list-style-type: none"> • Coredump file splitting • Support for using multiple Secure Remote Services gateway • Open LDAP support for NFSv4 ID mapping • DD Cloud Tier support for Microsoft Azure cool storage and AWS S3 infrequent access storage classes • Large object size support for DD Cloud Tier • Certificate-based login for Data Domain System Manager • KMIP enhancements
03 (6.1.1)	February 2018	<p>This revision includes information about these topics:</p> <ul style="list-style-type: none"> • Automatic multi-Streaming (AMS) for MTree replication • The <code>crepl-gc-gw-optim</code> option to improve throughput for collection replication • File system cleaning with collection replication
02 (6.1.1)	January 2018	<p>This revision includes information about these new features:</p> <ul style="list-style-type: none"> • Cloud Tier: <ul style="list-style-type: none"> ▪ Azure Government Cloud support ▪ S3 Flexible Cloud Provider support • DD Boost: <ul style="list-style-type: none"> ▪ Configure DD Boost authentication mode and encryption strength at the system level, or client level • DD VTL: <ul style="list-style-type: none"> ▪ View and configure VTL tape out to cloud settings ▪ Manually migrate VTL tapes to DD Cloud Tier

Table 2 Document revision history (continued)

Revision	Date	Description
		<ul style="list-style-type: none"> ▪ Manually recall VTL tapes from DD Cloud Tier • SNMP: <ul style="list-style-type: none"> ▪ View and configure SNMP engine IDs for SNMPv3 • System upgrades: <ul style="list-style-type: none"> ▪ View upgrade package checksums
01 (6.1)	June 2017	<p>This revision includes information about these new features:</p> <ul style="list-style-type: none"> • Cloud Tier: <ul style="list-style-type: none"> ▪ VTL tape out to cloud ▪ Infrequent access storage ▪ Direct restore and improved recall experience ▪ DD SM file recall enhancements ▪ Infrequent access support for Virtustream • NFS: <ul style="list-style-type: none"> ▪ Version 4 support added to DD OS ▪ Exports redesign ▪ Manageability improvements ▪ Server scaling improvements • Security: <ul style="list-style-type: none"> ▪ Key Management Interoperability Protocol (KMIP) support

Data Domain system overview

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

All systems run the Data Domain Operating System (DD OS), which provides both a command-line interface (CLI) for performing all system operations, and the 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 enclosures that add storage space.

Data Domain system features

Data Domain system features ensure data integrity, reliable restoration, efficient resource usage, and ease of management. Licensed features allow you to scale the system feature set to match your needs and budget.

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 use eight disks. Each parity stripe uses 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.
- When writing to SSD cache, the DD OS:
 - Creates an SL checksum for every record stored in the cache to detect corruption to cache data. This checksum is validated for every cache read.
 - Treats corruption to cache data as a cache miss and does not result in data loss. Therefore cache clients cannot store the latest copy of the data without some other backup mechanism such as NVRAM or HDD.
 - Removes the need for inline verification of cache writes, as cache clients can detect and handle misdirected or lost writes. This also saves I/O bandwidth.
 - Removes the need for SSD scrubbing of the the file system is, as the data in the cache keeps changing frequently and is already scrubbed by SAS Background Media Scan (BMS).

Data deduplication

DD OS data deduplication identifies redundant data during each backup and stores unique data just once.

The storage of unique data is invisible to backup software and 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 deduplication 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 deduplication.

Depending on backup volume, size, retention period, and rate of change, the amount of deduplication can vary. The best deduplication 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 with more than one Data Domain system must 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, maximum deduplication is achieved when the incremental backups and future full backups for sales data also go to Data Domain system A.

Restore operations

File restore operations create little or no contention with backup or other restore operations.

When backing up to disks on a Data Domain system, incremental backups are always reliable and can be easily accessed. With tape backups, a restore operation may rely on multiple tapes holding incremental backups. Also, the more incremental backups a site stores on multiple tapes, the more time-consuming and risky the restore process. One bad tape can kill the restore.

Using a Data Domain system, you can perform full backups more frequently without the penalty of storing redundant data. 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.

Data Domain Replicator

The Data Domain Replicator sets up and manages the replication of backup data between two Data Domain systems.

A DD Replicator pair consists of a source and a destination system and replicates a complete data set or directory from the source system to the destination 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. After replication is started, the source system automatically sends any new backup data to the destination system.

Multipath and load balancing

In a Fibre Channel multipath configuration, multiple paths are established between a Data Domain system and a backup server or backup destination array. When multiple

paths are present, the system automatically balances the backup load between the available paths.

At least two HBA ports are required to create a multipath configuration. When connected to a backup server, each of the HBA ports on the multipath is connected to a separate port on the backup server.

High Availability

The High Availability (HA) feature lets you configure two Data Domain systems as an Active-Standby pair, providing redundancy in the event of a system failure. HA keeps the active and standby systems in sync, so that if the active node were to fail due to hardware or software issues, the standby node can take over services and continue where the failing node left off.

The HA feature:

- Supports failover of backup, restore, replication and management services in a two-node system. Automatic failover requires no user intervention.
- Provides a fully redundant design with no single point of failure within the system when configured as recommended.
- Provides an Active-Standby system with no loss of performance on failover.
- Provides failover within 10 minutes for most operations. CIFS, DD VTL, and NDMP must be restarted manually.

Note

Recovery of DD Boost applications may take longer than 10 minutes, because Boost application recovery cannot begin until the DD server failover is complete. In addition, Boost application recovery cannot start until the application invokes the Boost library. Similarly, NFS may require additional time to recover.

-
- Supports ease of management and configuration through DD OS CLIs.
 - Provides alerts for malfunctioning hardware.
 - Preserves single-node performance and scalability within an HA configuration in both normal and degraded mode.
 - Supports the same feature set as stand-alone DD systems.

Note

DD Extended Retention and vDisk are not supported.

-
- Supports systems with all SAS drives. This includes legacy systems upgraded to systems with all SAS drives.

Note

The Hardware Overview and Installation Guides for the Data Domain systems that support HA describes how to install a new HA system. The *Data Domain Single Node to HA Upgrade* describes how to upgrade an existing system to an HA pair.

-
- Does not impact the ability to scale the product.
 - Supports nondisruptive software updates.

HA is supported on the following Data Domain systems:

- DD6800

- DD9300
- DD9500
- DD9800

HA architecture

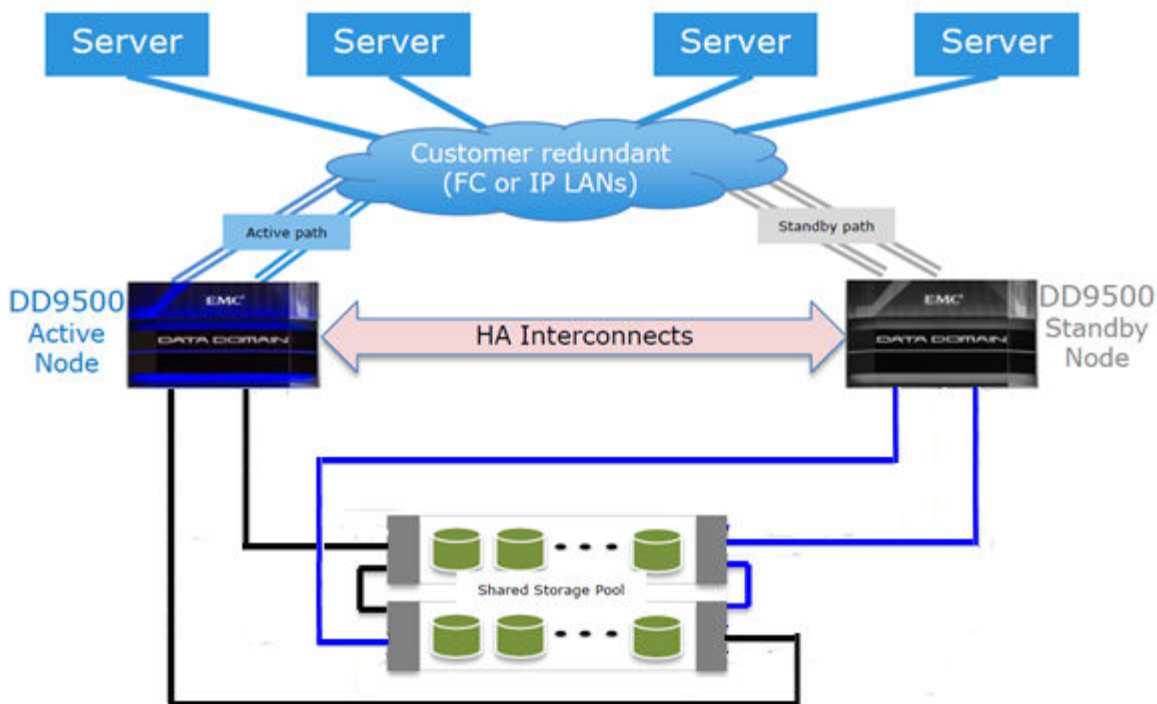
HA functionality is available for both IP and FC connections. Both nodes must have access to the same IP networks, FC SANs, and hosts in order to achieve high availability for the environment.

Over IP networks, HA uses a floating IP address to provide data access to the Data Domain HA pair regardless of which physical node is the active node.

Over FC SANs, HA uses NPIV to move the FC WWNs between nodes, allowing the FC initiators to re-establish connections after a failover.

Figure 1 on page 25 shows the HA architecture.

Figure 1 HA architecture



Random I/O handling

The random I/O optimizations included in DD OS provide improved performance for applications and use cases that generate larger amounts of random read and write operations than sequential read and write operations.

DD OS is optimized to handle workloads that consists of random read and write operations, such as virtual machine instant access and instant restore, and incremental forever backups generated by applications such as Avamar. These optimizations:

- Improve random read and random write latencies.
- Improve user IOPS with smaller read sizes.

- Support concurrent I/O operations within a single stream.
- Provide peak read and write throughput with smaller streams.

Note

The maximum random I/O stream count is limited to the maximum restore stream count of a Data Domain system.

The random I/O enhancements allow the Data Domain system to support instant access/instant restore functionality for backup applications such as Avamar and NetWorker.

System administrator access

System administrators can access the system for configuration and management using a command line interface or a graphical user interface.

- **DD OS CLI**—A command-line interface that is available through a serial console or through Ethernet connections using SSH or Telnet. CLI commands enable initial system configuration, changes to individual system settings, and display of system operation status.
 - **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.
-

Note

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

Licensed features

Feature licenses allow you to purchase only those features you intend to use. Some examples of features that require licenses are DD Extended Retention, DD Boost, and storage capacity increases.

Consult with your sales representative for information on purchasing licensed features.

Table 3 Features requiring licenses

Feature Name	License Name in Software	Description
Data Domain ArchiveStore	ARCHIVESTORE	Licenses Data Domain systems for archive use, such as file and email archiving, file tiering, and content and database archiving.
Data Domain Boost	DDBOOST	Enables the use of a Data Domain system with the following applications: Avamar, NetWorker, Oracle RMAN, Quest vRanger, Symantec Veritas NetBackup (NBU), and Backup Exec. The managed file replication (MFR) feature of DD Boost also requires the DD Replicator license.
Data Domain Capacity on Demand	CONTROLLER-COD	Enables an on-demand capacity increase for 4 TB DD2200 systems to 7.5 TB or 13.18 TB. An

Table 3 Features requiring licenses (continued)

Feature Name	License Name in Software	Description
		increase to 13.18 TB also requires the EXPANDED-STORAGE license.
Data Domain Cloud Tier	CLOUDTIER-CAPACITY	Enables a Data Domain system to move data from the active tier to low-cost, high-capacity object storage in the public, private, or hybrid cloud for long-term retention.
Data Domain Encryption	ENCRYPTION	Allows data on system drives or external storage to be encrypted while being saved and locked when moving the system to another location.
Data Domain Expansion Storage	EXPANDED-STORAGE	Allows Data Domain system storage to be expanded beyond the level provided in the base system.
Data Domain Extended Retention (formerly DD Archiver)	EXTENDED-RETENTION	Licenses the DD Extended Retention storage feature.
Data Domain I/OS (for IBM i operating environments)	I/OS	An I/OS license is required when DD VTL is used to backup systems in the IBM i operating environment. Apply this license before adding virtual tape drives to libraries.
Data Domain Replicator	REPLICATION	Adds DD Replicator for replication of data from one Data Domain system to another. A license is required on each system.
Data Domain Retention Lock Compliance Edition	RETENTION-LOCK-COMPLIANCE	Meets the strictest data retention requirements from regulatory standards such as SEC17a-4.
Data Domain Retention Lock Governance Edition	RETENTION-LOCK-GOVERNANCE	Protects selected files from modification and deletion before a specified retention period expires.
Data Domain Shelf Capacity-Active Tier	CAPACITY-ACTIVE	Enables a Data Domain system to expand the active tier storage capacity to an additional enclosure or a disk pack within an enclosure.
Data Domain Shelf Capacity-Archive Tier	CAPACITY-ARCHIVE	Enables a Data Domain system to expand the archive tier storage capacity to an additional enclosure or a disk pack within an enclosure.
Data Domain Storage Migration	STORAGE-MIGRATION-FOR-DATADOMAIN-SYSTEMS	Enables migration of data from one enclosure to another to support replacement of older, lower-capacity enclosures.
Data Domain Virtual Tape Library (DD VTL)	VTL	Enables the use of a Data Domain system as a virtual tape library over a Fibre Channel network. This license also enables the NDMP Tape Server feature, which previously required a separate license.

Table 3 Features requiring licenses (continued)

Feature Name	License Name in Software	Description
High Availability	HA-ACTIVE-PASSIVE	Enables the High Availability feature in an Active-Standby configuration. You only need to purchase one HA license; the license runs on the active node and is mirrored to the standby node.

Storage environment integration

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, DD Boost, or DD VTL protocols.
- Search for *compatibility documents* at <https://support.emc.com> for information on the applications that work with the different configurations.
- 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 that is accessed through a Fibre Channel connection. The DD 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: Avamar, 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 2 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/DD VTL/DD Boost
10. Data Verification
11. File system
12. Global deduplication and compression
13. RAID

As shown in [Figure 2](#) on page 29, 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 DD VTL and DD Boost).

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 the following 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:

- [DD System Manager overview](#) 32
- [Logging in and out of DD System Manager](#) 32
- [The DD System Manager interface](#) 34
- [Configuring the system with the configuration wizard](#) 36
- [Data Domain Command Line Interface](#) 49
- [Logging into the CLI](#) 49
- [CLI online help guidelines](#) 50

DD System Manager overview

DD System Manager is a browser-based graphical user interface, available through Ethernet connections, for managing a single system from any location. DD System Manager provides a single, consolidated management interface that allows for configuration and monitoring of many system features and system settings.

Note

Data Domain Management Center allows you to manage multiple systems from a single browser window.

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

Logging in and out of DD System Manager

Use a browser to log in to 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 **Secure Login**.

Secure login with HTTPS requires a digital certificate to validate the identity of the DD OS system and to support bidirectional encryption between DD System Manager and a browser. DD OS includes a self-signed certificate, and DD OS allows you to import your own certificate.

The default settings of most browsers do not automatically accept a self-signed certificate. This does not prevent you from using the self-signed certificate; it just means that you must respond to a warning message each time you perform

a secure log in, or you must install the certificate in your browser. For instructions on how to install the certificate in your browser, see your browser documentation.

3. Enter your assigned username and password.

Note

The initial username is *sysadmin* and the initial password is the system serial number. For information on setting up a new system, see the *Data Domain Operating System Initial Configuration Guide*.

4. Click **Log In**.

If this is the first time you have logged in, the Home view appears in the Information panel.

Note

If you enter an incorrect password 4 consecutive times, the system locks out the specified username for 120 seconds. The login count and lockout period are configurable and might be different on your system.

Note

If this is the first time you are logging in, you might be required to change your password. If the system administrator has configured your username to require a password change, you must change the password before gaining access to DD System Manager.

5. To log out, click the log out button in the DD System Manager banner.

When you log out, the system displays the log in page with a message that your log out is complete.

Logging in using a certificate

As an alternative to logging in using a username and password, you can log in to DD System Manager with a certificate issued by a Certificate Authority (CA).

To log in using a certificate, you must have authorization privileges on the Data Domain system, and the Data Domain system must trust the CA certificate. Your username must be specified in the common-name field in the certificate.

Procedure

1. Ensure that you have a user account on the Data Domain system.

You can be either a local user or a name service user (NIS/AD). For a name service user, your group-to-role mapping must be configured on the Data Domain system.

2. Use the following CLI command to import the public key from the CA that issued the certificate: `adminaccess certificate import ca application login-auth.`

3. Load the certificate in PKCS12 format in your browser.

Once the CA certificate is trusted by the Data Domain system, a **Log in with certificate** link is visible on the HTTPS login screen.

- Click **Log in with certificate** and choose the certificate from the list of certificates prompted by the browser.

Results

The Data Domain system validates the user certificate against the trust store. Based on authorization privileges associated with your account, a System Manager session is created for you.

The DD System Manager interface

The DD System Manager interface provides common elements on most pages that enable you to navigate through the configuration and display options and display context sensitive help.

Page elements

The primary page elements are the banner, the navigation panel, the information panels, and footer.

Figure 3 DD System Manager page components

The screenshot displays the DD System Manager interface. The banner at the top (1) shows the Dell EMC logo, the title 'Data Domain System Manager', and buttons for Refresh, Log Out, and Help. The navigation panel on the left (2) lists various system components. The main content area features an Alerts section with a table of alerts, and three information panels: File System, Services, and Hardware. The footer (4) contains system information such as the version, OS, model, user, and role.

Count	Type	Most recent Alerts
2	Hardware	No connection is detected on the Fibre Channel Port.
0	Replication	No Alerts
0	File System	No Alerts
0	Others	No Alerts

File System: Status: Running, Used: 0.2 (Total: 8516.9 GB), Data Written (Last 24 hrs): Logical Written: 0.0 GB, Physical Written: 0.0 GB, Compression Factor: 0.0x

Services: Replication: Errors: 0 Warning: 0 Normal: 1, DD-VTL: Running, CIFS: Enabled, NFS: Enabled, DD Boost: Enabled, vDisk: Disabled

Hardware: Enclosures: 2, Storage: 17

DD System Manager: rtp-ddr19.datadomain.com OS: 6.12.1-594656 Model: DC890 User: sysadmin Role: admin

- Banner
- Navigation panel
- Information panels
- Footer

Banner

The DD System Manager banner displays the program name and buttons for **Refresh**, **Log Out**, and **Help**.

Navigation panel

The Navigation panel displays the highest level menu selections that you can use to identify the system component or task that you want to manage.

The Navigation panel displays the top two levels of the navigation system. Click any top level title to display the second level titles. Tabs and menus in the Information panel provide additional navigation controls.

Information panel

The Information panel displays information and controls related to the selected item in the Navigation panel. The information panel is where you find system status information and configure a system.

Depending on the feature or task selected in the Navigation panel, the Information panel may display a tab bar, topic areas, table view controls, and the More Tasks menu.

Tab bar

Tabs provide access to different aspects of the topic selected in the Navigation panel.

Topic areas

Topic areas divide the Information panel into sections that represent different aspects of the topic selected in the Navigation panel or parent tab.

For high-availability (HA) systems, the HA Readiness tab on the System Manager dashboard indicates whether the HA system is ready to fail over from the active node to the standby node. You can click on **HA Readiness** to navigate to the **High Availability** section under **HEALTH**.

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 provide a More Tasks menu at the top right of the view that contains commands related to the current view.

Footer

The DD System Manager footer displays important information about the management session.

The banner lists the following information.

- System hostname.
- DD OS version
- Selected system model number.
- User name and role for the current logged in user.

Help buttons

Help buttons display a ? and appear in the banner, in the title of many areas of the Information panel, and in many dialogs. Click the help button to display a help window related to the current feature you are using.

The help window provides a contents button and navigation button above the help. Click the contents button to display the guide contents and a search button that you can use to search the help. Use the directional arrow buttons to page through the help topics in sequential order.

End User License Agreement

To view the End User License Agreement (EULA), select **Maintenance > System > View EULA**.

Configuring the system with the configuration wizard

There are two wizards, a DD System Manager configuration wizard and a Command Line Interface (CLI) configuration wizard. The configuration wizards guide you through a simplified configuration of your system to get your system operating quickly.

After you complete the basic configuration with a wizard, you can use additional configuration controls in DD System Manager and the CLI to further configure your system.

Note

The following procedure describes how to start and run the DD System Manager configuration wizard after the initial configuration of your system. For instructions on running the configuration wizards at system startup, see the *Data Domain Operating System Initial Configuration Guide*.

Note

If you want to configure your system for high availability (HA), you must perform this operation using the CLI Configuration Wizard. For more information, see the *Data Domain DD9500/DD9800 Hardware Overview and Installation Guide* and the *Data Domain Operating System Initial Configuration Guide*.

Procedure

1. Select **Maintenance > System > Configure System**.
2. Use the controls at the bottom of the Configuration Wizard dialog to select which features you want to configure and to advance through the wizard. To display help for a feature, click the help icon (question mark) in the lower left corner of the dialog.

License page

The License page displays all installed licenses. Click **Yes** to add, modify, or delete a license, or click **No** to skip license installation.

License Configuration

The **Licenses Configuration** section allows you add, modify or delete licenses from a license file. Data Domain Operating System 6.0 and later supports ELMS licensing, which allows you to include multiple features in a single license file upload.

When using the Configuration Wizard on a system without any licenses configured on it, select the license type from the drop-down, and click the ... button. Browse to the directory where the license file resides, and select it for upload to the system.

Table 4 License Configuration page values

Item	Description
Add Licenses	Select this option to add licenses from a license file.
Replace Licenses	If licenses are already configured the Add Licenses selection changes to Replace Licenses . Select this option to replace the licenses already added.
Delete Licenses	Select this option to delete licenses already configured on the system.

Network

The **Network** section allows you to configure the network settings. Click **Yes** to configure the network settings, or click **No** to skip network configuration.

Network General page

The General page allows you to configure network settings that define how the system participates in an IP network.

To configure these network settings outside of the configuration wizard, select **Hardware > Ethernet**.

Table 5 General page settings

Item	Description
Obtain Settings using DHCP	Select this option to specify that the system collect network settings from a Dynamic Host Control Protocol (DHCP) server. When you configure the network interfaces, at least one of the interfaces must be configured to use DHCP.
Manually Configure	Select this option to use the network settings defined in the Settings area of this page.
Host Name	Specifies the network hostname for this system.

Table 5 General page settings (continued)

Item	Description
	<p>Note</p> <p>If you choose to obtain the network settings through DHCP, you can manually configure the hostname at Hardware > Ethernet > Settings or with the <code>net set hostname</code> command. You must manually configure the host name when using DHCP over IPv6.</p>
Domain Name	Specifies the network domain to which this system belongs.
Default IPv4 Gateway	Specifies the IPv4 address of the gateway to which the system will forward network requests when there is no route entry for the destination system.
Default IPv6 Gateway	Specifies the IPv6 address of the gateway to which the system will forward network requests when there is no route entry for the destination system.

Network Interfaces page

The Interfaces page allows you to configure network settings that define how each interface participates in an IP network.

To Configure these network settings outside of the configuration wizard, select **Hardware > Ethernet > Interfaces**.

Table 6 Interfaces page settings

Item	Description
Interface	Lists the interfaces available on your system.
Enabled	Shows whether each interface is enabled (checkbox selected) or disabled (not selected). Click the checkbox to toggle the interface between the enabled and disabled states.
DHCP	Shows the current Dynamic Host Control Protocol (DHCP) configuration for each interface. Select v4 for IPv4 DHCP connections, v6 for IPv6 connections, or no to disable DHCP.
IP Address	Specifies an IPv4 or IPv6 address for this system. To configure the IP address, you must set DHCP to No .
	<p>Note</p> <p>DD140, DD160, DD610, DD620, and DD630 systems do not support IPv6 on interface eth0a (eth0 on systems that use legacy port names) or on any VLANs created on that interface.</p>
Netmask	Specifies the network mask for this system. To configure the network mask, you must set DHCP to No .
Link	Displays whether the Ethernet link is active (Yes) or not (No).

Network DNS page

The DNS page allows you to configure how the system obtains IP addresses for DNS servers in a Domain Name System (DNS).

To Configure these network settings outside of the configuration wizard, select **Hardware > Ethernet > Settings**.

Table 7 DNS page settings

Item	Description
Obtain DNS using DHCP.	Select this option to specify that the system collect DNS IP addresses from a Dynamic Host Control Protocol (DHCP) server. When you configure the network interfaces, at least one of the interfaces must be configured to use DHCP.
Manually configure DNS list.	Select this option when you want to manually enter DNS server IP addresses.
Add (+) button	Click this button to display a dialog in which you can add a DNS IP address to the DNS IP Address list. You must select Manually configure DNS list before you can add or delete DNS IP addresses.
Delete (X) button	Click this button to delete a DNS IP address from the DNS IP Address list. You must select the IP address to delete before this button is enabled. You must also select Manually configure DNS list before you can add or delete DNS IP addresses.
IP Address Checkboxes	Select a checkbox for a DNS IP address that you want to delete. Select the DNS IP Address checkbox when you want to delete all IP addresses. You must select Manually configure DNS list before you can add or delete DNS IP addresses.

File System

The **File System** section allows you to configure Active and Cloud Tier storage. Each has a separate wizard page. You can also create the File System within this section. The configuration pages cannot be accessed if the file system is already created.

Anytime you display the **File System** section when the File System has not been created, the system displays an error message. Continue with the procedure to create the file system.

Configure storage tier pages

The configure storage tier pages allow you to configure storage for each licensed tier on the system, Active Tier, Archive Tier, and DD Cloud Tier. Each tier has a separate wizard page. The storage tier configuration pages cannot be accessed if the file system is already created.

Configure Active Tier

The Configure Active Tier section allows you to configure the Active Storage Tier devices. The Active Tier is where back up data resides. To add storage to the Active

Tier, select one or more devices and add them to the tier. You can add storage devices up to the capacity licenses installed.

The DD3300 system requires 4 TB devices for the Active Tier.

Table 8 Addable Storage

Item	Description
ID (Device in DD VE)	The disk identifier, which can be any of the following. <ul style="list-style-type: none"> The enclosure and disk number (in the form Enclosure Slot, or Enclosure Pack for DS60 shelves) A device number for a logical device such as those used by DD VTL and vDisk A LUN
Disks	The disks that comprise the disk pack or LUN. This does not apply to DD VE instances.
Model	The type of disk shelf. This does not apply to DD VE instances.
Disk Count	The number of disks in the disk pack or LUN. This does not apply to DD VE instances.
Disk Size (Size in DD VE)	The data storage capacity of the disk when used in a Data Domain system. ^a
License Needed	The licensed capacity required to add the storage to the tier.
Failed Disks	Failed disks in the disk pack or LUN. This does not apply to DD VE instances.
Type	SCSI. This only applies to DD VE instances.

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.

Table 9 Active Tier values

Item	Description
ID (Device in DD VE)	The disk identifier, which can be any of the following. <ul style="list-style-type: none"> The enclosure and disk number (in the form Enclosure Slot, or Enclosure Pack for DS60 shelves). This does not apply to DD VE instances. A device number for a logical device such as those used by DD VTL and vDisk A LUN
Disks	The disks that comprise the disk pack or LUN. This does not apply to DD VE instances.
Model	The type of disk shelf. This does not apply to DD VE instances.
Disk Count	The number of disks in the disk pack or LUN. This does not apply to DD VE instances.
Disk Size (Size in DD VE)	The data storage capacity of the disk when used in a Data Domain system. ^a

Table 9 Active Tier values (continued)

Item	Description
License Used	The licensed capacity consumed by the storage.
Failed Disks	Failed disks in the disk pack or LUN. This does not apply to DD VE instances.
Configured	New or existing storage. This does not apply to DD VE instances.
Type	SCSI. This only applies to DD VE instances.

- 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.

Configure Archive Tier

The Configure Archive Tier section allows you to configure the Archive Storage Tier devices. The Archive Tier is where data archived with the DD Extended Retention feature resides. To add storage to the Archive Tier, select one or more devices and add them to the tier. You can add storage devices up to the capacity licenses installed. Archive Tier storage is not available on the DD3300 system, or on DD VE instances.

Table 10 Addable Storage

Item	Description
ID	The disk identifier, which can be any of the following. <ul style="list-style-type: none"> The enclosure and disk number (in the form Enclosure Slot, or Enclosure Pack for DS60 shelves) A device number for a logical device such as those used by DD VTL and vDisk A LUN
Disks	The disks that comprise the disk pack or LUN.
Model	The type of disk shelf.
Disk Count	The number of disks in the disk pack or LUN.
Disk Size (Size in DD VE)	The data storage capacity of the disk when used in a Data Domain system. ^a
License Needed	The licensed capacity required to add the storage to the tier.
Failed Disks	Failed disks in the disk pack or LUN.

- 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.

Table 11 Archive Tier values

Item	Description
ID	The disk identifier, which can be any of the following. <ul style="list-style-type: none"> The enclosure and disk number (in the form Enclosure Slot, or Enclosure Pack for DS60 shelves). This does not apply to DD VE instances.

Table 11 Archive Tier values (continued)

Item	Description
	<ul style="list-style-type: none"> • A device number for a logical device such as those used by DD VTL and vDisk • A LUN
Disks	The disks that comprise the disk pack or LUN.
Model	The type of disk shelf.
Disk Count	The number of disks in the disk pack or LUN.
Disk Size (Size in DD VE)	The data storage capacity of the disk when used in a Data Domain system. ^a
License Used	The licensed capacity consumed by the storage.
Failed Disks	Failed disks in the disk pack or LUN.
Configured	New or existing storage.

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.

Configure Cloud Tier

The Configure Cloud Tier section allows you to configure the Cloud Storage Tier devices. To add storage to the Cloud Tier, select one or more devices and add them to the tier. You can add storage devices up to the capacity licenses installed.

The DD3300 system requires 1 TB devices for DD Cloud Tier.

Table 12 Addable Storage

Item	Description
ID (Device in DD VE)	<p>The disk identifier, which can be any of the following.</p> <ul style="list-style-type: none"> • The enclosure and disk number (in the form Enclosure Slot, or Enclosure Pack for DS60 shelves) • A device number for a logical device such as those used by DD VTL and vDisk • A LUN
Disks	The disks that comprise the disk pack or LUN. This does not apply to DD VE instances.
Model	The type of disk shelf. This does not apply to DD VE instances.
Disk Count	The number of disks in the disk pack or LUN. This does not apply to DD VE instances.
Disk Size (Size in DD VE)	The data storage capacity of the disk when used in a Data Domain system. ^a
License Needed	The licensed capacity required to add the storage to the tier.
Failed Disks	Failed disks in the disk pack or LUN. This does not apply to DD VE instances.
Type	SCSI. This only applies to DD VE instances.

Table 12 Addable Storage (continued)

- 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.

Table 13 Cloud Tier values

Item	Description
ID (Device in DD VE)	The disk identifier, which can be any of the following. <ul style="list-style-type: none"> The enclosure and disk number (in the form Enclosure Slot, or Enclosure Pack for DS60 shelves). This does not apply to DD VE instances. A device number for a logical device such as those used by DD VTL and vDisk A LUN
Disks	The disks that comprise the disk pack or LUN. This does not apply to DD VE instances.
Model	The type of disk shelf. This does not apply to DD VE instances.
Disk Count	The number of disks in the disk pack or LUN. This does not apply to DD VE instances.
Disk Size (Size in DD VE)	The data storage capacity of the disk when used in a Data Domain system. ^a
License Used	The licensed capacity consumed by the storage.
Failed Disks	Failed disks in the disk pack or LUN. This does not apply to DD VE instances.
Configured	New or existing storage. This does not apply to DD VE instances.
Type	SCSI. This only applies to DD VE instances.

- 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.

Create File System page

The Create File System page displays the allowed size of each storage tier in the file system, and allows you to automatically enable the file system after it is created.

System Settings

The **System Settings** section allows you to configure system passwords, and email settings. Click **Yes** to configure the system settings, or click **No** to skip system settings configuration.

System Settings Administrator page

The Administrator page allows you to configure the administrator password and how the system communicates with the administrator.

Table 14 Administrator page settings

Item	Description
User Name	The default administrator name is <i>sysadmin</i> . The sysadmin user cannot be renamed or deleted.
Old Password	Type the old password for sysadmin.
New Password	Type the new password for sysadmin.
Verify New Password	Retype the new password for sysadmin.
Admin Email	Specify the email address to which DD System Manager sends alert and autosupport email messages.
Send Alert Notification Emails to this address	Check to configure DD System Manager to send alert notifications to the Admin email address as alert events occur.
Send Daily Alert Summary Emails to this address	Check to configure DD System Manager to send alert summaries to the Admin email address at the end of each day.
Send Autosupport Emails to this address	Check to configure DD System Manager to send the Admin user autosupport emails, which are daily reports that document system activity and status.

System Settings Email/Location page

The Email/Location page allows you to configure the mail server name, control what system information is sent to Data Domain, and specify a location name to identify your system.

Table 15 Email/Location page settings

Item	Description
Mail Server	Specify the name of the mail server that manages emails to and from the system.
Send Alert Notification Emails to Data Domain	Check to configure DD System Manager to send alert notification emails to Data Domain.
Send Vendor Support Notification Emails to Data Domain	Check to configure DD System Manager to send vendor support notification emails to Data Domain.
Location	Use this optional attribute as needed to record the location of your system. If you specify a location, this information is stored as the SNMP system location.

DD Boost protocol

The **DD Boost Settings** section allows you to configure the DD Boost protocol settings. Click **Yes** to configure the DD Boost Protocol settings, or click **No** to skip DD Boost configuration.

DD Boost Protocol Storage Unit page

The Storage Unit page allows you to configure DD Boost storage units.

To configure these settings outside of the configuration wizard, select **Protocols > DD Boost > Storage Units > + (plus sign)** to add a storage unit, **the pencil** to modify a storage unit, or **X** to delete a storage unit.

Table 16 Storage Unit page settings

Item	Description
Storage Unit	The name of your DD Boost Storage Unit. You may optionally change this name.
User	<p>For the default DD Boost user, either select an existing user, or select Create a new Local User, and enter their User name, Password, and Management Role. This role can be one of the following:</p> <ul style="list-style-type: none"> • <i>Admin role:</i> Lets you configure and monitor the entire Data Domain system. • <i>User role:</i> Lets you monitor Data Domain systems and change your own password. • <i>Security role:</i> In addition to user role privileges, lets you set up security-officer configurations and manage other security-officer operators. • <i>Backup-operator role:</i> In addition to user role privileges, lets you create snapshots, import and export tapes to, or move tapes within a DD VTL. • <i>None role:</i> Intended only for DD Boost authentication, so you cannot monitor or configure a Data Domain system. None is also the parent role for the SMT tenant-admin and tenant-user roles. None is also the preferred user type for DD Boost storage owners. Creating a new local user here only allows that user to have the "none" role.

DD Boost Protocol Fibre Channel page

The Fibre Channel page allows you to configure DD Boost Access Groups over Fibre Channel.

To configure these settings outside of the configuration wizard, select **Protocols > DD Boost > Fibre Channel > + (plus sign)** to add an access group, **the pencil** to modify an access group, or **X** to delete an access group.

Table 17 Fibre Channel page settings

Item	Description
Configure DD Boost over Fibre Channel	Select the checkbox if you want to configure DD Boost over Fibre Channel.
Group Name (1-128 Chars)	Create an Access Group. Enter a unique name. Duplicate access groups are not supported.
Initiators	Select one or more initiators. Optionally, replace the initiator name by entering a new one. An initiator is a backup client that connects to the system to read and write data using the FC (Fibre Channel) protocol. A specific initiator can support DD Boost over FC or DD VTL, but not both.
Devices	The devices to be used are listed. They are available on all endpoints. An endpoint is the logical target on the Data Domain system to which the initiator connects.

CIFS protocol

The **CIFS Protocol** settings section allows you to configure the CIFS protocol settings. Click **Yes** to configure the CIFS protocol settings, or click **No** to skip CIFS configuration.

Data Domain systems use the term MTree to describe directories. When you configure a directory path, DD OS creates an MTree where the data will reside.

CIFS Protocol Authentication page

The Authentication page enables you to configure Active Directory and Workgroup for your system.

To configure these settings outside of the configuration wizard, select **Administration > Access > Authentication**.

Table 18 Authentication page settings

Item	Description
Active Directory/Kerberos Authentication	Expand this panel to enable, disable, and configure Active Directory Kerberos authentication.
Workgroup Authentication	Expand this panel to configure Workgroup authentication.

CIFS Protocol Share page

The Share page enables you to configure a CIFS protocol share name and a directory path for your system.

To configure these settings outside of the configuration wizard, select **Protocols > CIFS > Shares > Create**.

Table 19 Share page settings

Item	Description
Share Name	Enter a share name for the system.

Table 19 Share page settings (continued)

Item	Description
Directory Path	Enter a directory path for the system.
Add (+) button	Click + to enter a system client.
Pencil icon	Modify a client.
Delete (X) button	Click X to delete a selected client.

NFS protocol

The **NFS Protocol** settings section allows you to configure the NFS protocol settings. Click **Yes** to configure the NFS protocol settings, or click **No** to skip NFS configuration.

Data Domain systems use the term MTree to describe directories. When you configure a directory path, DD OS creates an MTree where the data will reside.

NFS Protocol Export page

The Export page enables you to configure an NFS protocol export directory path, network clients, and NFSv4 referrals.

To configure these settings outside of the configuration wizard, select **Protocols > NFS > Create**.

Table 20 Export page settings

Item	Description
Directory Path	Enter a pathname for the export.
Add (+) button	Click + to enter a system client or NFSv4 referral.
Pencil icon	Modify a client or NFSv4 referral.
Delete (X) button	Click X to delete a selected client or NFSv4 referral.

DD VTL protocol

The **DD VTL Protocol** settings section allows you to configure the Data Domain Virtual Tape Library settings. Click **Yes** to configure the DD VTL settings, or click **No** to skip DD VTL configuration.

VTL Protocol Library page

The Library page allows you to configure the DD VTL protocol settings for a library.

To configure these settings outside of the configuration wizard, select **PROTOCOLS > VTL > Virtual Tape Libraries > VTL Service > Libraries > More Tasks > Library > Create**

Table 21 Library page settings

Item	Description
Library Name	Enter a name of from 1 to 32 alphanumeric characters.

Table 21 Library page settings (continued)

Item	Description
Number of Drives	Number of supported tape drives.
Drive Model	Select the desired model from the drop-down list: <ul style="list-style-type: none"> • IBM-LTO-1 • IBM-LTO-2 • IBM-LTO-3 • IBM-LTO-4 • IBM-LTO-5 (default) • HP-LTO-3 • HP-LTO-4
Number of Slots	Enter the number of slots per library: <ul style="list-style-type: none"> • Up to 32,000 slots per library • Up to 64,000 slots per system • This should be equal to, or greater than, the number of drives.
Number of CAPs	(Optional) Enter the number of cartridge access ports (CAPs): <ul style="list-style-type: none"> • Up to 100 CAPs per library • Up to 1000 CAPs per system
Changer Model Name	Select the desired model from the drop-down list: <ul style="list-style-type: none"> • L180 (default) • RESTORER-L180 • TS3500 • I2000 • I6000 • DDVTL
Starting Barcode	Enter the desired barcode for the first tape, in the format A990000LA.
Tape Capacity	(Optional) Enter the tape capacity. If not specified, the capacity is derived from the last character of the barcode.

VTL Protocol Access Group page

The Access Group page allows you to configure DD VTL protocol settings for an access group.

To configure these settings outside of the configuration wizard, select **PROTOCOLS > VTL > Access Groups > Groups > More Tasks > Group > Create**.

Table 22 Access Group page settings

Item	Description
Group Name	Enter a unique name of from 1 - 128 characters. Duplicate access groups are not supported.
Initiators	Select one or more initiators. Optionally, replace the initiator name by entering a new one. An initiator is a backup client that connects to a system to read and write data using the Fibre Channel (FC) protocol. A specific initiator can support DD Boost over FC or DD VTL, but not both.
Devices	The devices (drives and changer) to be used are listed. These are available on all endpoints. An endpoint is the logical target on the Data Domain system to which the initiator connects.

Data Domain Command Line Interface

The Command Line Interface (CLI) is a text driven interface that can be used instead of or in addition to DD System Manager. Most management tasks can be performed in DD System Manager or with the CLI. In some cases, the CLI offers configuration options and reports that are not yet supported in DD System Manager.

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 do the following.

- 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`.
- Show the next available option, if you do not enter any characters before pressing the Tab key.
- Show partial matched tokens or complete a unique entry, if you enter characters before pressing the Tab key.

The *Data Domain Operating System Command Reference Guide* provides information for each of the CLI commands. Online help is available and provides the complete syntax for each command.

Logging into the CLI

You can access the CLI using a direct connection to the system or using an Ethernet connection through SSH or Telenet.

Before you begin

To use the CLI, you must establish a local or remote connection to the system using one of the following methods.

- If you are connecting through a serial console port on the system, connect a terminal console to the port and use the communication settings: 9600 baud, 8 data bits, no parity, and 1 stop bit.
- If the system provides keyboard and monitor ports, connect a keyboard and monitor to those ports.
- If you are connecting through Ethernet, connect a computer with SSH or Telnet client software to an Ethernet network that can communicate with the system.

Procedure

1. If you are using an SSH or Telnet connection to access the CLI, start the SSH or Telnet client and specify the IP address or host name of the system.

For information on initiating the connection, see the documentation for the client software. The system prompts you for your username.

2. When prompted, enter your system username.
3. When prompted, enter your system password.

The following example shows SSH login to a system named *mysystem* using SSH client software.

```
# ssh -l sysadmin mysystem.mydomain.com
Data Domain OS 5.6.0.0-19899
Password:
```

CLI online help guidelines

The CLI displays two types of help, syntax-only help and command-description help that includes the command syntax. Both types of help offer features that allow you reduce the time it takes to find the information you need.

The following guidelines describe how to use syntax-only help.

- To list the top-level CLI commands, enter a question mark (?), or type the command `help` at the prompt.
- To list all forms of a top-level command, enter the command with no options at the prompt or enter *command* ?.
- To list all commands that use a specific keyword, enter `help keyword` or ? *keyword*.
For example, ? `password` displays all Data Domain system commands that use the `password` argument.

The following guidelines describe how to use command-description help.

- To list the top-level CLI commands, enter a question mark (?), or type the command `help` at the prompt.
- To list all forms of a top-level command with an introduction, enter `help command` or ? *command*.
- The end of each help description is marked `END`. Press Enter to return to the CLI prompt.
- When the complete help description does not fit in the display, the colon prompt (:) appears at the bottom of the display. The following guidelines describe what you can do when this prompt appears.
 - To move through the help display, use the up and down arrow keys.
 - To quit the current help display and return to the CLI prompt, press `q`.
 - To display help for navigating the help display, press `h`.

- To search for text in the help display, enter a slash character (/) followed by a pattern to use as search criteria and press Enter. Matches are highlighted.

CHAPTER 3

Managing Data Domain Systems

This chapter includes:

• System management overview	54
• Rebooting a system	55
• Powering a system on or off	55
• System upgrade management	57
• Managing electronic licenses	61
• System storage management	61
• Network connection management	69
• System passphrase management	91
• System access management	93
• Configuring mail server settings	117
• Managing time and date settings	117
• Managing system properties	118
• SNMP management	119
• Autosupport report management	127
• Support bundle management	130
• Coredump management	131
• Alert notification management	131
• Support delivery management	139
• Log file management	140
• Remote system power management with IPMI	145

System management overview

DD System Manager allows you to manage the system on which DD System Manager is installed.

- 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 6.0, DD System Manager supports the addition of systems for replication for DD OS Version 5.6 to 5.7 plus the next two releases.

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 following table recommends the maximum number of user sessions supported by DD System Manager:

Table 23 Maximum number of users supported by DD System Manager

System Model	Maximum Active Users	Maximum Logged In Users
4 GB models ^a	5	10
8 GB models ^b	10	15
16 GB and greater models ^c	10	20

a. Includes DD140 and DD2200 (4TB)

b. Includes DD610 and DD630

c. Includes DD670, DD860, DD890, DD990, DD2200 (>7.5TB), DD4200, DD4500, DD6300, DD6800, DD7200, DD9300, DD9500, and DD9800

Note

Initial HA system set-up cannot be done from the DD System Manager, but the status of an already-configured HA system can be viewed from DD System Manager.

HA system management overview

The HA relationship between the two nodes, one active and one standby, is setup through DDSH CLIs.

Initial set-up can be run on either of the two nodes but only one at a time. It is a precondition of HA that the system interconnect and identical hardware is setup on both nodes first.

Note

Both DDRs are required to have identical hardware which will be validated during setup and system boot-up.

If set-up is from a fresh install of systems, the `ha create` command needs to be run on the node with the license installed. If setup is from an existing system and a new fresh install system (upgrade), then it should be run from the existing system.

HA system planned maintenance

The HA architecture provides a rolling upgrade, which reduces maintenance downtime for a DD OS upgrade.

With a rolling upgrade, the HA nodes are upgraded one by one, coordinately and automatically. The standby node is rebooted and upgraded first. The newly upgraded node then takes over the active role through an HA failover. After the failover, the second node is rebooted and assumes the role of the standby node after the upgrade.

System upgrade operations that require data conversion cannot start until both systems are upgraded to the same level and HA state is fully restored.

Rebooting a system

Reboot a system when a configuration change, such as changing the time zone, requires that you reboot the system.

Procedure

1. Select **Maintenance > System > Reboot System**.
2. Click **OK** to confirm.

Powering a system on or off

When powering a system off and on, it is important that you follow the proper procedure to preserve the file system and configuration integrity.

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.

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.

For HA systems, a connection to both nodes is required.

Complete the following steps to power off a Data Domain system.

Procedure

1. Verify that I/O on the system is stopped.

Run the following commands:

- `cifs show active`
- `nfs show active`
- `system show stats view sysstat interval 2`
- `system show perf`

2. For HA systems, verify the health of the HA configuration.

Run the following command:

```
ha status
```

```
HA System Name: apollo-ha3a.emc.com
HA System Status: highly available
Node Name           Node ID   Role      HA State
-----
apollo-ha3a-p0.emc.com    0       active   online
apollo-ha3a-p1.emc.com    1       standby  online
-----
```

Note

This output sample is from a healthy system. If the system is being shut down to replace a failed component, the HA System Status will be degraded, and one or both nodes will show offline for the HA State.

3. Run the `alerts show current` command. For HA pairs, run the command on the active node first, and then the standby node.
4. For HA systems, run the `ha offline` command if the system is in a highly available state with both nodes online. Skip this step if the HA status is degraded.
5. Run the `system poweroff` command. For HA pairs, run the command on the active node first, and then the standby node.

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

6. Remove the power cords from the power supplies on the controller or controllers.
7. Verify the blue power LED is off on the controller or controllers to confirm that the system is powered down.

Power a system on

Restore power to the Data Domain system when the system downtime is complete.

Procedure

1. Power on any expansion shelves before powering on the Data Domain controller. Wait approximately three minutes after all expansion shelves are turned on.

Note

A 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 there is a power button on the controller, press the power button (as shown in the *Installation and Setup Guide* for your Data Domain system). For HA systems, power on the active node first, and then the standby node.
3. For HA systems, verify the health of the HA configuration.

Run the following command:

```
ha status
```

```
HA System Name: apollo-ha3a.emc.com
HA System Status: highly available
Node Name           Node ID   Role      HA State
-----
apollo-ha3a-p0.emc.com    0       active   online
-----
```



```
apollo-ha3a-p1.emc.com      1      standby  offline
-----
```

4. For HA systems, if one of the nodes displays as offline, run the `ha online` command on that node to restore the HA configuration.
5. Run the `alerts show current` command. For HA pairs, run the command on the active node first, and then the standby node.

System upgrade management

To upgrade a DD OS system, you must verify that there is sufficient room for the new software on the target system, transfer the software to the system to be upgraded, and then start the upgrade. For an HA system, transfer the software to the active node and start the upgrade from the active node.

For HA systems, use the floating IP address to access DD System Manager to perform software upgrades.

CAUTION

DD OS 6.0 uses Secure Remote Support version 3 (ESRSv3). Upgrading a system running DD OS 5.X to DD OS 6.0 removes the existing ConnectEMC configuration from the system. After the upgrade is complete, reconfigure ConnectEMC manually.

If the system uses MD5-signed certificates, regenerate the certificates with a stronger hash algorithm during the upgrade process.

Minimally disruptive upgrade

The minimally disruptive upgrade (MDU) feature lets you upgrade specific software components or apply bug fixes without needing to perform a system reboot. Only those services that depend on the component being upgraded are disrupted, so the MDU feature can prevent significant downtime during certain software upgrades.

Not all software components qualify for a minimally disruptive upgrade; such components must be upgraded as part of a regular DD OS system software upgrade. A DD OS software upgrade uses a large RPM (upgrade bundle), which performs upgrade actions for all of the components of DD OS. MDU uses smaller component bundles, which upgrade specific software components individually.

RPM signature verification

RPM signature verification validates Data Domain RPMs that you download for upgrade. If the RPM has not been tampered with, the digital signature is valid and you can use the RPM as usual. If the RPM has been tampered with, the corruption invalidates the digital signature, and the RPM is rejected by DD OS. An appropriate error message is displayed.

Note

When upgrading from 5.6.0.x to 6.0, first upgrade the 5.6.0.x system to 5.6.1.x (or later) before upgrading to 6.0.

Support software

DD OS 6.1 introduces a type of software package called support software. Support software is provided by Data Domain Support Engineering to address specific issues. By default, the Data Domain system does not allow support software to be installed on the system. Contact Support for more information about support software.

Viewing upgrade packages on the system

DD System Manager allows you to view and manage up to five upgrade packages on a system. Before you can upgrade a system, you must download an upgrade package from the Online Support site to a local computer, and then upload it to the target system.

Procedure

1. Select **Maintenance > System**.
2. Optionally, select an upgrade package and click **View Checksum** to display the MD5 and SHA256 checksums of the upgrade package.

Results

For every package stored on the system, DD System Manager displays the filename, file size, and last modified date in the list titled: Upgrade Packages Available on Data Domain System.

Obtaining and verifying upgrade packages

You can use DD System Manager to locate upgrade package files on the Data Domain Support Web site and upload copies of those files to a system.

Note

You can use FTP or NFS to copy an upgrade package to a system. DD System Manager is limited to managing 5 system upgrade packages, but there are no restrictions, other than space limitations, when you manage the files directly in the `/ddvar/releases` directory. FTP is disabled by default. To use NFS, `/ddvar` needs to be exported and mounted from an external host).

Procedure

1. Select **Maintenance > System**.
2. To obtain an upgrade package, click the **EMC Online Support** link, click Downloads, and use the search function to locate the package recommended for your system by Support personnel. Save the upgrade package to the local computer.
3. Verify that there are no more than four packages listed in the Upgrade Packages Available on Data Domain System list.

DD System Manager can manage up to five upgrade packages. If five packages appear in the list, remove at least one package before uploading the new package.
4. Click **Upload Upgrade Package** to initiate the transfer of the upgrade package to the system.
5. In the Upload Upgrade Package dialog, click **Browse** to open the Choose File to Upload dialog. Navigate to the folder with the downloaded file, select the file, and click **Open**.
6. Click **OK**.

An upload progress dialog appears. Upon successful completion of the upload, the download file (with a `.rpm` extension) appears in the list titled: Upgrade Packages Available on Data Domain System.

7. To verify the upgrade package integrity, click **View Checksum** and compare the calculated checksum displayed in the dialog to the authoritative checksum on the Online Support site.
8. To manually initiate an upgrade precheck, select an upgrade package and click **Upgrade Precheck**.

Upgrade considerations for HA systems

HA systems require one unique pre-check before initiating the upgrade operation, and one unique post-check after the upgrade is complete.

The HA system must be in a highly available state, with both nodes online before performing the DD OS upgrade. Run the `ha status` command to verify the HA system state.

```
# ha status
HA System Name: apollo-ha3a.emc.com
HA System Status: highly available
Node Name           Node ID  Role      HA State
-----
apollo-ha3a-p0.emc.com 0        active    online
apollo-ha3a-p1.emc.com 1        standby   online
-----
```

DD OS automatically recognizes the HA system and performs the upgrade procedure on both nodes.

After the upgrade procedure is complete, run the `ha status` command again to verify that the system is in a highly available state, and both nodes are online.

Upgrading a Data Domain system

When an upgrade package file is present on a system, you can use DD System Manager to perform an upgrade using that upgrade package.

Before you begin

Read the DD OS Release Notes for the complete upgrade instructions and coverage of all the issues that can impact the upgrade.

The procedure that follows describes how to initiate an upgrade using DD System Manager. Log out of any Data Domain CLI sessions on the system where the upgrade is to be performed before using DD System Manager to upgrade the system.

Note

Upgrade package 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. Log into DD System Manager on 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 6.0, upgrades are permitted from Releases 5.6 and 5.7.

Note

As recommended in the Release Notes, reboot the Data Domain system before upgrading to verify that the hardware is in a clean state. If any issues are discovered during the reboot, resolve those issues before starting the upgrade. For an MDU upgrade, a reboot may not be needed.

2. Select **Data Management > File System**, and verify that the file system is enabled and running.
 3. Select **Maintenance > System**.
 4. From the Upgrade Packages Available on Data Domain System list, select the package to use for the upgrade.
-

Note

You must select an upgrade package for a newer version of DD OS. DD OS does not support downgrades to previous versions.

5. Click **Perform System Upgrade**.
The System Upgrade dialog appears and displays information about the upgrade and a list of users who are currently logged in to the system to be upgraded.
 6. Verify the version of the upgrade package, and click **OK** to continue with the upgrade.
The System Upgrade dialog displays the upgrade status and the time remaining.
When upgrading the system, you must wait for the upgrade to complete before using DD System Manager to manage the system. If the system restarts, the upgrade might continue after the restart, and DD System Manager displays the upgrade status after login. If possible, keep the System Upgrade progress dialog open until the upgrade completes or the system powers off. When upgrading DD OS Release 5.5 or later to a newer version, and if the system upgrade does not require a power off, a Login link appears when the upgrade is complete.
-

Note

To view the status of an upgrade using the CLI, enter the `system upgrade status` command. Log messages for the upgrade are stored in `/ddvar/log/debug/platform/upgrade-error.log` and `/ddvar/log/debug/platform/upgrade-info.log`.

7. If the system powers down, you must remove AC power from the system to clear the prior configuration. Unplug all of the power cables for 30 seconds and then plug them back in. The system powers on and reboots.
8. If the system does not automatically power on and there is a power button on the front panel, press the button.

After you finish

For environments that use self-signed SHA-256 certificates, the certificates must be regenerated manually after the upgrade process is complete, an trust must be re-established with external systems that connect to the Data Domain system.

1. Run the `adminaccess certificate generate self-signed-cert regenerate-ca` command to regenerate the self-signed CA and host certificates.

Regenerating the certificates breaks existing trust relationships with external systems.

2. Run the `adminaccess trust add host hostname type mutual` command to reestablish mutual trust between the Data Domain system and the external system.

Removing an upgrade package

A maximum of five upgrade packages can be uploaded to a system with DD System Manager. If the system you are upgrading contains five upgrade packages, you must remove at least one package before you can upgrade the system.

Procedure

1. Select **Maintenance > System**.
2. From the list titled Upgrade Packages Available on Data Domain System, select the package to remove. One package can be removed at a time.
3. Click **Remove Upgrade Package**.

Managing electronic licenses

Add and delete electronic licenses from the Data Domain system. Refer to the applicable *Data Domain Operating System Release Notes* for the most up-to-date information on product features, software updates, software compatibility guides, and information about products, licensing, and service.

HA system license management

HA is a licensed feature, and the system licensing key is registered by following the same steps to add any other license to the DD system.

A system will be configured as Active-Standby, where one node is designated "standby." Only one set of licenses will be required rather than needing individual licenses for each node. During failover, the licenses on one node will failover to the other node.

System storage management

System storage management features allow you to view the status and configuration of your storage space, flash a disk LED to facilitate disk identification, and change the storage configuration.

Note

All storage connected or used by the two-node Active-Standby HA system can be viewed as a single system.

Viewing system storage information

The storage status area shows the current status of the storage, such as Operational or Non-Operational, and the storage migration status. Below the Status area are tabs that organize how the storage inventory is presented.

Procedure

1. To display the storage status, select **Hardware** > **Storage**.
2. If an alerts link appears after the storage status, click the link to view the storage alerts.
3. If the Storage Migration Status is Not licensed, you can click **Add License** to add the license for this feature.

Overview tab

The Overview tab displays information for all disks in the Data Domain system organized by type. The categories that display are dependent on the type of storage configuration in use.

The Overview tab lists the discovered storage in one or more of the following sections.

- **Active Tier**
Disks in the Active Tier are currently marked as usable by the file system. Disks are listed in two tables, Disks in Use and Disks Not in Use.
- **Retention Tier**
If the optional Data Domain Extended Retention (formerly DD Archiver) license is installed, this section shows the disks that are configured for DD Extended Retention storage. Disks are listed in two tables, Disks in Use and Disks Not in Use. For more information, see the *Data Domain Extended Retention Administration Guide*.
- **Cache Tier**
SSDs in the Cache Tier are used for caching metadata. The SSDs are not usable by the file system. Disks are listed in two tables, Disks in Use and Disks Not in Use.
- **Cloud Tier**
Disks in the Cloud Tier are used to store the metadata for data that resides in cloud storage. The disks are not usable by the file system. Disks are listed in two tables, Disks in Use and Disks Not in Use.
- **Addable Storage**
For systems with optional enclosures, this section shows the disks and enclosures that can be added to the system.
- **Failed/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.
- **Migration History**
Shows the history of migrations.

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.

Click a section plus (+) button to display detailed information, or click the minus (-) button to hide the detailed information.

Table 24 Disks In Use column label descriptions

Item	Description
Disk Group	The name of the disk group that was created by the file system (for example, dg1).
State	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).
Size	The size of the disk group (for example, 25.47 TiB).

Table 25 Disks Not In Use column label descriptions

Item	Description
Disk	The disk identifier, which can be any of the following. <ul style="list-style-type: none"> The enclosure and disk number (in the form Enclosure Slot) A device number for a logical device such as those used by DD VTL and vDisk A LUN
Slot	The enclosure where the disk is located.
Pack	The disk pack, 1-4, within the enclosure where the disk is located. This value will only be 2-4 for DS60 expansion shelves.
State	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. ^a
Type	The disk connectivity and type (For example, SAS).

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.

Enclosures tab

The Enclosures tab displays a table summarizing the details of the enclosures connected to the system.

The Enclosures tab provides the following details.

Table 26 Enclosures tab column label descriptions

Item	Description
Enclosure	The enclosure number. Enclosure 1 is the head unit.
Serial Number	The enclosure serial number.

Table 26 Enclosures tab column label descriptions (continued)

Item	Description
Disks	The disks contained in the enclosure, in the format <i><Enclosure-number>.1- <Enclosure-number>. <N></i> .
Model	The enclosure model. For enclosure 1, the model is Head Unit.
Disk Count	The number of disks in the enclosure.
Size	The data storage capacity of the disk when used in a Data Domain system. ^a
Failed Disks	The failed disks in the enclosure.
Temperature Status	The temperature status of the enclosure.

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.

Disks tab

The Disks tab displays information on each of the system disks. You can filter the disks viewed to display all disks, disks in a specific tier, or disks in a specific group.

The Disk State table displays a summary status table showing the state of all system disks.

Table 27 Disks State table column label descriptions

Item	Description
Total	The total number of inventoried disks in the Data Domain system.
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 DD 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.
Migrating	The number of disks serving as the source of a storage migration.
Destination	The number of disks serving as the destination of a storage migration.
Not Installed	The number of empty disk slots that the system can detect.

The Disks table displays specific information about each disk installed in the system.

Table 28 Disks table column label descriptions

Item	Description
Disk	<p>The disk identifier, which can be:</p> <ul style="list-style-type: none"> • The enclosure and disk number (in the form <i>Enclosure.Slot</i>). • A device number for a logical device such as those used by DD VTL and vDisk.. • A LUN.
Slot	The enclosure where the disk is located.
Pack	The disk pack, 1-4, within the enclosure where the disk is located. This value will only be 2-4 for DS60 expansion shelves.
State	<p>The status of the disk, which can be one of the following.</p> <ul style="list-style-type: none"> • Absent. No disk is installed in the indicated location. • Available. An available disk is allocated to the active or retention tier, but it is not currently in use. • Copy Recovery. The disk has a high error rate but is not failed. RAID is currently copying the contents onto a spare drive and will fail the drive once the copy reconstruction is complete. • Destination. The disk is in use as the destination for storage migration. • Error. The disk has a high error rate but is not failed. The disk is in the queue for copy reconstruction. The state will change to Copy Recovery when copy reconstruction begins. • Foreign. The disk has been assigned to a tier, but the disk data indicates the disk may be owned by another system. • In-Use. The disk is being used for backup data storage. • Known. The disk is a supported disk that is ready for allocation. • Migrating. The disk is in use as the source for storage migration. • Powered Off. The disk power has been removed by Support. • Reconstruction. The disk is reconstructing in response to a <code>disk fail</code> command or by direction from RAID/SSM. • Spare. The disk is available for use as a spare. • System. System disks store DD OS and system data. No backup data is stored on system disks. • Unknown. An unknown disk is not allocated to the active or retention tier. It might have been failed administratively or by the RAID system.

Table 28 Disks table column label descriptions (continued)

Item	Description
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.
Disk Life Used	The percentage of an SSD's rated life span consumed.
Type	The disk connectivity and type (For example, SAS).

Reconstruction tab

The Reconstruction tab displays a table that provides additional information on reconstructing disks.

The following table describes the entries in the Reconstructing table.

Table 29 Reconstruction table column label descriptions

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 is complete.

When a spare disk is available, the 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.

Physically locating an enclosure

If you have trouble determining which physical enclosure corresponds to an enclosure displayed in DD System Manager, you can use the CLI beacon feature to flash the enclosure IDENT LEDs and all the disk LEDs that indicate normal operation.

Procedure

1. Establish a CLI session with the system.
2. Type `enclosure beacon enclosure`.
3. Press `ctrl-c` to stop the LED flashing.

Physically locating a disk

If you have trouble determining which physical disk corresponds to a disk displayed in DD System Manager, you can use the beacon feature to flash an LED on the physical disk.

Procedure

1. Select **Hardware > Storage > Disks**.
2. 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.

3. Click **Stop** to stop the LED beaconing.

Configuring storage

Storage configuration features allow you to add and remove storage expansion enclosures from the active, retention, and cloud tiers. Storage in an expansion enclosure (which is sometimes called an expansion shelf) is not available for use until it is added to a tier.

Note

Additional storage requires the appropriate license or licenses and sufficient memory to support the new storage capacity. Error messages display if more licenses or memory is needed.

DD6300 systems support the option to use ES30 enclosures with 4 TB drives (43.6 TiB) at 50% utilization (21.8 TiB) in the active tier if the available licensed capacity is exactly 21.8 TiB. The following guidelines apply to using partial capacity shelves.

- No other enclosure types or drive sizes are supported for use at partial capacity.
- A partial shelf can only exist in the Active tier.
- Only one partial ES30 can exist in the Active tier.
- Once a partial shelf exists in a tier, no additional ES30s can be configured in that tier until the partial shelf is added at full capacity.

Note

This requires licensing enough additional capacity to use the remaining 21.8 TiB of the partial shelf.

- If the available capacity exceeds 21.8 TB, a partial shelf cannot be added.
- Deleting a 21 TiB license will not automatically convert a fully-used shelf to a partial shelf. The shelf must be removed, and added back as a partial shelf.

Procedure

1. Select **Hardware > Storage > Overview**.
2. Expand the dialog for one of the available storage tiers:

- **Active Tier**
 - **Extended Retention Tier**
 - **Cache Tier**
 - **Cloud Tier**
3. Click **Configure**.
 4. In the Configure Storage dialog, select the storage to be added from the **Addable Storage** list.
 5. In the **Configure** list, select either **Active Tier** or **Retention Tier**.

The maximum amount of storage that can be added to the active tier depends on the DD controller used.

Note

The licensed capacity bar shows the portion of licensed capacity (used and remaining) for the installed enclosures.

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**.

DD3300 capacity expansion

The DD3300 system is available in three different capacity configurations. Capacity expansions from one configuration to another are supported.

The DD3300 system is available in the following capacity configurations:

- 4 TB
- 16 TB
- 32 TB

The following upgrade considerations apply:

- A 4 TB system can be upgraded to 16 TB.
- A 16 TB system can be upgraded to 32 TB.
- There is no upgrade path from 4 TB to 32 TB.

Select **Maintenance > System** to access information about capacity expansion, and to initiate the capacity expansion process.

The capacity expansion is a one-time process. The **Capacity Expansion History** pane displays whether or not the system has already been expanded. If the system has not been expanded, click the **Capacity Expand** button to initiate the capacity expansion.

All capacity expansions require the installation of additional disks and memory in the system. Do not attempt to expand the capacity until the hardware upgrades are complete. The following table lists the hardware upgrade requirements for capacity expansion.

Capacity expansion	Additional memory	Additional HDDs	Additional SSD
4 TB to 16 TB	32 GB	6 x 4 TB HDDs	1 x 480 GB SSD
16 TB to 32 TB	16 GB	6 x 4 TB HDDs	N/A

The *Data Domain DD3300 Field Replacement and Upgrade Guide* provides detailed instructions for expanding system capacity.

Capacity Expand

Select the target capacity from the **Select Capacity** drop-down list. A capacity expansion can be prevented by insufficient memory, insufficient physical capacity (HDDs), the system has already been expanded, or the target for capacity expansion is not supported. If the capacity expansion cannot be completed, the reason will display here.

Capacity expansion history

The **Capacity Expansion History** table displays details about the capacity of the system. The table provides the capacity of the system when the software was first installed, the date of the initial software installation. If the capacity was expanded, the table also provides the expanded capacity, and the date the expansion was performed.

Fail and unfail disks

Disk fail functionality allows you to manually set a disk to a failed state to force reconstruction of the data stored on the disk. Disk unfail functionality allows you to take a disk in a failed state and return it to operation

Fail a disk

Fail a disk and force reconstruction. Select **Hardware > Storage > Disks > Fail**.

Unfail a disk

Make a disk previously marked Failed or Foreign usable to the system. Select **Hardware > Storage > Disks > Unfail**.

Network connection management

Network connection management features allow you view and configure network interfaces, general network settings, and network routes.

HA system network connection management

The HA system relies on two different types of IP addresses, fixed and floating. Each type has specific behaviors and limitations.

On an HA system, Fixed IP addresses:

- Are used for node management via the CLI
- Are attached ("fixed") to the node
- Can be static or DHCP, IPv6 SLAAC

- Configuration is done on the specific node with the optional `type fixed` argument

Note

All filesystem access should be through a floating IP.

Floating IP addresses only exist in the two-node HA system; during failover, the IP address "float" to the new active node and are:

- Only configured on the active node
- Used for filesystem access and most configuration
- Can only be static
- Configuration requires the `type floating` argument

Network interface management

Network interface management features allow you to manage the physical interfaces that connect the system to a network and create logical interfaces to support link aggregation, load balancing, and link or node failover.

Viewing interface information

The Interfaces tab allows you to manage physical and virtual interfaces, VLANs, DHCP, DDNS, and IP addresses and aliases.

Consider the following guidelines when managing IPv6 interfaces.

- The command-line interface (CLI) supports IPv6 for basic Data Domain network and replication commands, but not for backup and DD 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, as is Managed File Replication using DD Boost.
- There are some restrictions for interfaces with 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. Select **Hardware > Ethernet > Interfaces**.

The following table describes the information on the Interfaces tab.

Table 30 Interface tab label descriptions

Item	Description
Interface	The name of each interface associated with the selected system.
Enabled	Whether the interface is enabled.

Table 30 Interface tab label descriptions (continued)

Item	Description
	<ul style="list-style-type: none"> Select Yes to enable the interface and connect it to the network. Select No to disable the interface and disconnect it from the network.
DHCP	Indicates if the interface is configured manually (<code>no</code>), by a DHCP (Dynamic Host Configuration Protocol) IPv4 server (<code>v4</code>), or by a DHCP IPv6 server (<code>v6</code>).
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 Ethernet connection is active (Yes/No).
Address Type	On an HA system, the Address Type indicates Fixed, Floating, or Interconnect.
Additional Info	Additional settings for the interface. For example, the bonding mode.
IPMI interfaces configured	Displays Yes or No and indicates if IPMI health monitoring and power management is configured for the interface.

- To filter the interface list by interface name, enter a value in the **Interface Name** field and click **Update**.

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

- To filter the interface list by interface type, select a value from the **Interface Type** menu and click **Update**.

On an HA system, there is a filter dropdown to filter by IP Address Type (Fixed, Floating, or Interconnect).

- To return the interfaces table to the default listing, click **Reset**.
- Select an interface in the table to populate the Interface Details area.

Table 31 Interface Details label descriptions

Item	Description
Auto-generated Addresses	Displays the automatically generated IPv6 addresses for the selected interface.
Auto Negotiate	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.
Cable	Shows whether the interface is Copper or Fiber.

Table 31 Interface Details label descriptions (continued)

Item	Description
	<p>Note</p> <p>Some interfaces must be up before the cable status is valid.</p>
Duplex	Used in conjunction with the Speed value to set the data transfer protocol. Options are Unknown, Full, Half.
Hardware Address	The MAC address of the selected interface. For example, 00:02:b3:b0:8a:d2.
Interface Name	Name of the selected interface.
Latent Fault Detection (LFD) - HA systems only	The LFD field has a View Configuration link, displaying a pop-up that lists LFD addresses and interfaces.
Maximum Transfer Unit (MTU)	MTU value assigned to the interface.
Speed	Used in conjunction with the Duplex value to set the rate of data transfer. Options are Unknown, 10 Mb/s, 100 Mb/s, 1000 Mb/s, 10 Gb/s.
	<p>Note</p> <p>Auto-negotiated interfaces must be set up before speed, duplex, and supported speed are visible.</p>
Supported Speeds	Lists all of the speeds that the interface can use.

- To view IPMI interface configuration and management options, click **View IPMI Interfaces**.

This link displays the **Maintenance > IPMI** information.

Physical interface names and limitations

The format of physical interface names varies on different Data Domain systems and option cards, and limitations apply to some interfaces.

- 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 provide 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`.
- DD2200 systems provide four on-board 1G Base-T NIC ports: `ethMa` (top left), `ethMb` (top right), `ethMc` (bottom left), and `ethMd` (bottom right).
- DD2500 systems provide 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 provide one on-board Ethernet port, which is ethMa.
- For systems ranging between DD140 and DD990, the physical interface names for I/O 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 I/O 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 DD7200 I/O 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

Review the general interface configuration guidelines before configuring system interfaces.

- When supporting both backup and replication traffic, if possible, use 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, if possible, do not use 10 GbE interfaces for replication traffic because 10 GbE interfaces are optimized for faster traffic.
- If a Data Domain service uses a non-standard port and the user wants to upgrade to DD OS 6.0, or the user wants to change a service to use a non-standard port on a DD OS 6.0 system, add a net filter function for all the clients using that service to allow the client IP addresses to use the new port.
- On DD4200, DD4500, and DD7200 systems that use IPMI, if possible, reserve interface ethMa for IPMI traffic and system management traffic (using protocols such as HTTP, Telnet, and SSH). Backup data traffic should be directed to other interfaces.

Configuring physical interfaces

You must configure at least one physical interface before the system can connect to a network.

Procedure

1. Select **Hardware > Ethernet > Interfaces**.
2. Select an interface to configure.

Note

DD140, DD160, DD610, DD620, and DD630 systems do not support IPv6 on interface eth0a (eth0 on systems that use legacy port names) or on any VLANs created on that interface.

3. Click **Configure**.
4. In the Configure Interface dialog, determine how the interface IP address is to be set:

Note

On an HA system, the Configure Interface dialog has a field for whether or not to designate the Floating IP (Yes/No). Selecting **Yes** the `Manually Configure IP Address` radio button is auto-selected; Floating IP interfaces can only be manually configured.

- Use DHCP to assign the IP address—in the IP Settings area, select **Obtain IP Address using DHCP** and select either **DHCPv4** for IPv4 access or **DHCPv6** for IPv6 access. Setting a physical interface to use DHCP automatically enables the interface.
-

Note

If you choose to obtain the network settings through DHCP, you can manually configure the hostname at **Hardware > Ethernet > Settings** or with the `net set hostname` command. You must manually configure the host name when using DHCP over IPv6.

- Specify IP Settings manually—in the IP Settings area, select **Manually configure IP Address**. The **IP Address** and **Netmask** fields become active.
5. If you chose to manually enter the IP address, enter an IPv4 or IPv6 address. If you entered an IPv4 address, enter a netmask address.
-

Note

You can assign just one IP address to an interface with this procedure. If you assign another IP address, the new address replaces the old address. To attach an additional IP address to an interface, create an IP alias.

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 I/O modules:
 - Dual Port 10GbE SR Optical with LC connectors (using SFPs)
 - Dual Port 10GbE Direct Attach Copper (SFP+ cables)
 - Quad port 2 port 1GbE Copper (RJ45) /2 port 1GbE 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.
 - On DD2500, DD4200, DD4500, and DD7200 10GbE I/O 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.

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 a protocol that registers local IP addresses on a Domain Name System (DNS) server. In this release, DD System Manager supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command.

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**.

MTU size values

The MTU size must be set properly to optimize the performance of a network connection. An incorrect MTU size can negatively affect interface 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

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

Moving a static IP address

A specific static IP address must be assigned to only one interface on a system. A static IP address must be properly removed from an interface before it is configured on another interface.

Procedure

1. If the interface that hosts the static IP address is part of a DD Boost interface group, remove the interface from that group.
2. Select **Hardware > Ethernet > Interfaces**.
3. Remove the static IP address that you want to move.
 - a. Select the interface that is currently using the IP address you want to move.
 - b. In the Enabled column, select **No** to disable the interface.
 - c. Click **Configure**.
 - d. Set the IP Address to 0.

Note

Set the IP address to 0 when there is no other IP address to assign to the interface. The same IP address must not be assigned to multiple interfaces.

- e. Click **Next**, and click **Finish**.
4. Add the removed static IP address to another interface.
 - a. Select the interface to which you want to move the IP address.
 - b. In the Enabled column, select **No** to disable the interface.
 - c. Click **Configure**.
 - d. Set the IP Address to match the static IP address you removed.
 - e. Click **Next**, and click **Finish**.
 - f. In the Enabled column, select **Yes** to enable the updated interface.

Virtual interface configuration guidelines

Virtual interface configuration guidelines apply to failover and aggregate virtual interfaces. There are additional guidelines that apply to either failover or aggregate interfaces but not both.

- 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 is 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 support 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 for aggregated links 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**.
- For a bonded interface, the bonded interface is created with remaining slaves if the hardware for a slave interface fails. If no slaves, the bonded interface is created with no slaves. This slave hardware failure will generate managed alerts, one per failed slave.

Note

The alert for a failed slave disappears after the failed slave is removed from the system. If new hardware is installed, the alerts disappear and the bonded interface uses the new slave interface after the reboot.

- On DD4200, DD4500, and DD7200 systems, the ethMa interface does not support failover or link aggregation.

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. These guidelines are provided to help you optimize your use of link aggregation.

- 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.
- 10 Gb single-port optical Ethernet cards do not support link aggregation.
- 1 GbE and 10 GbE interfaces cannot be aggregated together.
- Copper and optical interfaces cannot be aggregated together.
- On DD4200, DD4500, and DD7200 systems, the ethMA interface does not support link aggregation.

Guidelines for configuring a virtual interface for failover

Link failover provides improved network stability and performance by identifying backup interfaces that can support network traffic when the primary interface is not operating. These guidelines are provided to help you optimize your use of link failover.

- 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.

Virtual interface creation

Create a virtual interface to support link aggregation or failover. The virtual interface serves as a container for the links to be aggregated or associated for failover.

Creating a virtual interface for link aggregation

Create a virtual interface for link aggregation to serve as a container to associate the links that participate in aggregation.

A link aggregation interface must specify a link bonding mode and may require a hash selection. 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.

Procedure

1. Select **Hardware > Ethernet > Interfaces**.
2. In the Interfaces table, disable the physical interface where the virtual interface is to be added by clicking **No** in the **Enabled** column.
3. From the **Create** menu, select **Virtual Interface**.
4. In the Create Virtual Interface dialog, 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.

5. In the **Bonding Type** list, select **Aggregate**.

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.

6. In the **Mode** list, select a bonding mode.

Specify the mode that is compatible with the requirements of the system to which the interfaces are directly attached.

- **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 uses a control protocol that communicates to the other end and coordinates which links within the bond are available for use. LACP provides a kind of heartbeat failover and must be configured at both ends of the link.

7. If you selected Balanced or LACP mode, specify a bonding hash type in the **Hash** list.

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).

8. 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.

9. Enter an IP address, or enter 0 to specify no IP address.
10. Enter a netmask address or prefix.
11. 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 Mb, 100 Mb, 1000 Mb, and 10 Gb.
 - 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.
 - The 10 GbE copper NIC default is 10 Gb. If a copper interface is set to 1000 Mb or 10 Gb line speed, duplex must be full-duplex.
12. Specify the MTU setting.
- To select the default value (1500), click **Default**.
 - To select a different setting, enter the setting in the **MTU** box. Ensure that all of your network components support the size set with this option.
13. Optionally, select Dynamic DNS Registration option.
- Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server. In this release, DD System Manager supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command.
- The DDNS must be registered to enable this option.
14. Click **Next**.
- The Configure Interface Settings summary page appears. The values listed reflect the new system and interface state.
15. Click **Finish** and **OK**.

Creating a virtual interface for link failover

Create a virtual interface for link failover to serve as a container to associate the links that will participate in failover.

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. 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.

Procedure

1. Select **Hardware > Ethernet > Interfaces**.
2. In the interfaces table, disable the physical interface to which the virtual interface is to be added by clicking **No** in the **Enabled** column.
3. From the **Create** menu, select **Virtual Interface**.
4. In the Create Virtual Interface dialog, 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.

5. In the **Bonding Type** list, select **Failover**.
6. Select an interface to add to the failover configuration, and click **Next**. Virtual aggregate interfaces can be used for failover.

The Create virtual interface *veth_name* dialog appears.

7. Enter an IP address, or enter 0 to specify no IP address.
8. Enter a netmask or prefix.
9. Specify the Speed/Duplex options.

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

- Select **Autonegotiate Speed/Duplex** to allow the network interface card to autonegotiate the line speed and duplex setting for an interface.
- Select **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 Mb, 100 Mb, 1000 Mb, and 10 Gb.
 - Half-duplex is available for 10 Mb and 100 Mb speeds only.
 - 1000 Mb and 10 Gb 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 Gb or 10 Gb line speed, the duplex must be full-duplex.

10. Specify MTU setting.
 - To select the default value (1500), click **Default**.
 - To select a different setting, enter the setting in the MTU box. Ensure that all of your network path components support the size set with this option.
11. Optionally, select Dynamic DNS Registration option.

Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server. In this release, DD System Manager supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command.

The DDNS must be registered to enable this option.

Note

This option disables DHCP for this interface.

12. Click **Next**.

The Configure Interface Settings summary page appears. The values listed reflect the new system and interface state.
13. Complete the Interface, click **Finish** and **OK**.

Modifying a virtual interface

After you create a virtual interface, you can update the settings to respond to network changes or resolve issues.

Procedure

1. Select **Hardware > Ethernet > Interfaces**.
2. 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.
3. In the **Interfaces** column, select the interface and click **Configure**.
4. In the **Configure Virtual Interface** dialog, change the settings.
5. Click **Next** and **Finish**.

Configuring a VLAN

Create a new VLAN interface from either a physical interface or a virtual interface.

The recommended total VLAN count 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. Select **Hardware > Ethernet > Interfaces**.
2. In the interfaces table, select the interface to which you want to add the VLAN.
The interface you select must be configured with an IP address before you can add a VLAN.
3. Click **Create** and select **VLAN**.
4. In the Create VLAN dialog box, specify a VLAN ID by entering a number in the **VLAN Id** box.

The range of a VLAN ID is between 1 and 4094 inclusive.

5. 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.

6. Enter a netmask or prefix.
7. Specify the MTU setting.

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.

- To select the default value (1500), click **Default**.
- To select a different setting, enter the setting in the MTU box. 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.

8. Specify Dynamic DNS Registration option.

Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server. In this release, DD System Manager

supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command.

The DDNS must be registered to enable this option.

9. Click **Next**.

The **Create VLAN** summary page appears.

10. Review the configuration settings, click **Finish**, and click **OK**.

Modifying a VLAN interface

After you create a VLAN interface, you can update the settings to respond to network changes or resolve issues.

Procedure

1. Select **Hardware > Ethernet > Interfaces**.
2. 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.
3. In the Interfaces column, select the checkbox of the interface and click **Configure**.
4. In the **Configure VLAN Interface** dialog, change the settings.
5. Click **Next** and **Finish**.

Configuring an IP alias

An IP alias assigns an additional IP address to 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.

Note

When using a Data Domain HA system, if a user is created and logs in to the standby node without logging into the active node first, the user will not have a default alias to use. Therefore, in order to use aliases on the standby node, the user should login to the active node first.

Procedure

1. Select **Hardware > Ethernet > Interfaces**.
2. Click **Create**, and select **IP Alias**.
The Create IP Alias dialog appears.
3. Specify an IP alias ID by entering a number in the **IP ALIAS Id** box.
The range is 1 to 4094 inclusive.
4. Enter an IPv4 or IPv6 address.
5. If you entered an IPv4 address, enter a netmask address.
6. Specify Dynamic DNS Registration option.

Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server. In this release, DD System Manager

supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command.

The DDNS must be registered to enable this option.

7. Click **Next**.

The Create IP Alias summary page appears.

8. Review the configuration settings, click **Finish**, and **OK**.

Modifying an IP alias interface

After you create an IP alias, you can update the settings to respond to network changes or resolve issues.

Procedure

1. Select **Hardware > Ethernet > Interfaces**.
2. 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.
3. In the **Interfaces** column, select the checkbox of the interface and click **Configure**.
4. In the Configure IP Alias dialog box, change the settings as described in the procedure for creating an IP Alias.
5. Click **Next** and **Finish**.

Registering interfaces with DDNS

Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server.

In this release, DD System Manager supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command. 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. Select **Hardware > Ethernet > Interfaces > DDNS Registration**.
2. In the DDNS Windows Mode Registration dialog, click **Add** to add an interface to the DDNS.

The Add Interface dialog box appears.

- a. Enter a name in the **Interface** field.
- b. Click **OK**.
3. 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**.

4. 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.
5. 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. Select **Hardware > Ethernet > Interfaces**.
2. Click the box next to each interface you want to destroy (Virtual or VLAN or IP Alias).
3. Click **Destroy**.
4. Click **OK** to confirm.

Viewing an interface hierarchy in the tree view

The Tree View dialog displays the association between physical and virtual interfaces.

Procedure

1. Select **Hardware > Ethernet > Interfaces > Tree View**.
2. In the Tree View dialog box, click the plus or minus boxes to expand or contract the tree view that shows the hierarchy.
3. Click **Close** to exit this view.

General network settings management

The configuration settings for hostname, domain name, search domains, host mapping, and DNS list are managed together on the Settings tab.

Viewing network settings information

The Settings tab displays the current configuration for the hostname, domain name, search domains, host mapping, and DNS.

Procedure

1. Select **Hardware > Ethernet > Settings**.

Results

The Settings tab displays the following information.

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, if possible, manually configure the host and domain names.

When configuring the hostname and domain name, consider the following guidelines.

- Do not include an underscore in the hostname; it is incompatible with some browsers.
- Replication and CIFS authentication must be reconfigured after you change the names.
- If a system was previously added without a fully qualified name (no domain name), a domain name change requires that you remove and add the affected system or update the Search Domain List to include the new domain name.

Procedure

1. Select **Hardware > Ethernet > Settings**.
2. Click **Edit** in the **Host Settings** area. The Configure Host dialog appears.
3. To manually configure the host and domain names:
 - a. Select **Manually configure host**.
 - 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.

4. 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

Use the domain search list to define which domains the system can search.

Procedure

1. Select **Hardware > Ethernet > Settings**.
2. Click **Edit** in the Search Domain List area.
3. To add a search domain using the Configure Search Domains dialog:
 - a. Click Add (+).
 - b. In the Add Search Domain dialog, 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.

4. To remove a search domain using the Configure Search Domains dialog:
 - 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.

Adding and deleting host maps

A host map links an IP address to a hostname, so that either the IP address or the hostname can be used to specify the host.

Procedure

1. Select **Hardware > Ethernet > Settings**.
2. To add a host map, do the following.
 - a. In the Hosts Mapping area, click **Add**.
 - b. In the Add Hosts dialog, enter the IP address of the host in the **IP Address** box.
 - c. Click Add (+).
 - d. In the Add Host dialog, enter a hostname, such as *id##.yourcompany.com*, in the **Host Name** box.
 - e. Click **OK** to add the new hostname to the Host Name list.

- f. Click **OK** to return to the Settings tab.
3. To delete a host map, do the following.
 - a. In the Hosts Mapping area, select the host mapping to delete.
 - b. Click Delete (**X**).

Configuring DNS IP addresses

DNS IP addresses specify the DNS servers the system can use to get IP addresses for host names that are not in the host mapping table.

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. Select **Hardware > Ethernet > Settings**.
2. Click **Edit** in the DNS List area.
3. 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.
4. 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.
5. 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.

Network route management

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 IPv4 policy-based routing, which allows only one route to a default gateway per routing table. There can be multiple route tables and multiple default gateways. A routing table is created for each address that has the same subnet as a default gateway. The routing rules send the packets with the source IP

address that matches the IP address used to create the table to that routing table. All other packets that do not have source IP addresses that match a routing table are sent to the main routing table.

Within each routing table, static routes can be added, but because source routing is used to get packets to the table, the only static routes that will work are static routes that use the interface that has the source address of each table. Otherwise it needs to be put into the main table.

Other than the IPv4 source routing done to these other routing tables, Data Domain systems use source-based routing for the main routing IPv4 and IPv6 tables, which means that outbound network packets that match the subnet of multiple interfaces are routed only over the physical interface whose IP address matches the source IP address of the packets, which is where they originated.

For IPv6, set static routes when 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.

Static routes can be added and deleted from individual routing tables by adding or deleting the table from the route specification. This provides the rules to direct packets with specific source addresses through specific route tables. If a static route is required for packets with those source addresses, the routes must be added the specific table where the IP address is routed.

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. This is not needed if the source address is IPv4 and has a default gateway associated with it. In that case, the source routing is already handled via its own routing table.

Viewing route information

The Routes tab displays the default gateways, static routes, and dynamic routes.

Procedure

1. Select **Hardware > Ethernet > Routes**.

Results

The Static Routes area lists the route specification used to configure each static route. The Dynamic Routes table lists information for each of the dynamically assigned routes.

Table 32 Dynamic Routes column label descriptions

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.

Table 32 Dynamic Routes column label descriptions (continued)

Item	Description
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.
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, if possible, manually configure the default gateway IP address.

Procedure

1. Select **Hardware > Ethernet > Routes**.
2. Click **Edit** next to the default gateway type (IPv4 or IPv6) you want to configure.
3. To manually configure the default gateway address:
 - a. Select **Manually Configure**.
 - b. Enter the gateway address in the **Gateway** box.
 - c. Click **OK**.
4. 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

Static routes define destination hosts or networks that they system can communicate with.

Procedure

1. Select **Hardware > Ethernet > Routes**.

2. Click **Create** in the Static Routes area.
3. In the **Create Routes** dialog, select the interface you want to host the static route, and click **Next**.
4. Specify the Destination.
 - 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.
5. Optionally, specify the gateway to use to connect to the destination network or host.
 - a. Select **Specify a gateway for this route**.
 - b. Enter the gateway address in the **Gateway** box.
6. Review the configuration and click **Next**.
The create routes Summary page appears.
7. Click **Finish**.
8. After the process is completed, click **OK**.
The new route specification is listed in the Route Spec list.

Deleting static routes

Delete a static route when you no longer want the system to communicate with a destination host or network.

Procedure

1. Select **Hardware > Ethernet > Routes**.
2. Select the Route Spec of the route specification to delete.
3. Click **Delete**.
4. Click **Delete** to confirm and then click **Close**.

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

System passphrase management

The system passphrase is a key that allows a Data Domain system to be transported with encryption keys on the system. The encryption keys protect the data and the system passphrase protects the encryption keys.

The system passphrase is a human-readable (understandable) key (like a smart card) which is used to generate a machine usable AES 256 encryption key. 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.

Setting the system passphrase

The system passphrase must be set before the system can support data encryption or request digital certificates.

Before you begin

No minimum system passphrase length is configured when DD OS is installed, but the CLI provides a command to set a minimum length. To determine if a minimum length is configured for the passphrase, enter the `system passphrase option show` CLI command.

Procedure

1. Select **Administration > Access > Administrator Access**.

If the system passphrase is not set, the **Set Passphrase** button appears in the Passphrase area. If a system passphrase is configured, the **Change Passphrase** button appears, and your only option is to change the passphrase.

2. Click the **Set Passphrase** button.

The Set Passphrase dialog appears.

3. Enter the system passphrase in the boxes and click **Next**.

If a minimum length is configured for the system passphrase, the passphrase you enter must contain the minimum number of characters.

Results

The system passphrase is set and the **Change Passphrase** button replaces the **Set Passphrase** button.

Changing the system passphrase

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.

Changing the passphrase requires two-user authentication to protect against data shredding.

Procedure

1. Select **Administration > Access > Administrator Access**.
2. To change the system passphrase, click **Change Passphrase**.

The Change Passphrase dialog appears.

Note

The file system must be disabled to change the passphrase. If the file system is running, you are prompted to disable it.

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 passphrase when changing the passphrase.

- The new passphrase, which must contain the minimum number of characters configured with the `system passphrase option set min-length` command.
4. Click the checkbox for **Enable file system now**.
 5. Click **OK**.

NOTICE

Be sure to take care of the passphrase. If the passphrase is lost, you can never unlock the file system and access the data; the data is irrevocably lost.

System access management

System access management features allow you to control system access to users in a local database or in a network directory. Additional controls define different access levels and control which protocols can access the system.

Role-based access control

Role-based access control (RBAC) is an authentication policy that controls which DD System Manager controls and CLI commands a user can access on a system.

For example, users who are assigned the *admin* role can configure and monitor an entire system, while users who are assigned the *user* role are limited to monitoring a system. When logged into DD System Manager, users see only the program features that they are permitted to use based on the role assigned to the user. The following roles are available for administering and managing the DD OS.

admin

An *admin* role user can configure and monitor the entire Data Domain system. Most configuration features and commands are available only to *admin* role users. However, some features and commands require the approval of a *security* role user before a task is completed.

limited-admin

The *limited-admin* role can configure and monitor the Data Domain system with some limitations. Users who are assigned this role cannot perform data deletion operations, edit the registry, or enter bash or SE mode.

user

The *user* role enables users to monitor systems and change their own password. Users who are assigned the *user* management role can view system status, but they cannot change the system configuration.

security (security officer)

A *security* role user, who may be referred to as a security officer, can manage other security officers, authorize procedures that require security officer approval, and perform all tasks supported for user-role users.

The *security* role is provided to comply with the Write Once Read-Many (WORM) regulation. This regulation requires electronically stored corporate data be kept in an unaltered, original state for purposes such as eDiscovery. Data Domain added auditing and logging capabilities to enhance this feature. As a result of compliance regulations, most command options for administering sensitive operations, such

as DD Encryption, DD Retention Lock Compliance, and archiving now require security officer approval.

In a typical scenario, an *admin* role user issues a command and, if security officer approval is required, the system displays a prompt for approval. To proceed with the original task, the security officer must enter his or her username and password on the same console at which the command was run. If the system recognizes the security officer credentials, the procedure is authorized. If not, a security alert is generated.

The following are some guidelines that apply to security-role users:

- Only the *sysadmin* user (the default user created during the DD OS installation) can create the first security officer, after which the privilege to create security officers is removed from the *sysadmin* user.
- After the first security officer is created, only security officers can create other security officers.
- Creating a security officer does not enable the authorization policy. To enable the authorization policy, a security officer must log in and enable the authorization policy.
- Separation of privilege and duty apply. *admin* role users cannot perform security officer tasks, and security officers cannot perform system configuration tasks.
- During an upgrade, if the system configuration contains security officers, a *sec-off-defaults* permission is created that includes a list of all current security officers.

backup-operator

A *backup-operator* role user can perform all tasks permitted for *user* role users, create snapshots for MTrees, import, export, and move tapes between elements in a virtual tape library, and copy tapes across pools.

A *backup-operator* role user can also add and delete SSH public keys for non-password-required log ins. (This function is used mostly for automated scripting.) He or she can add, delete, reset and view CLI command aliases, synchronize modified files, and wait for replication to complete on the destination system.

none

The *none* role is for DD Boost authentication and tenant-unit users only. A *none* role user can log in to a Data Domain system and can change his or her password, but cannot monitor, manage, or configure the primary system. When the primary system is partitioned into tenant units, either the *tenant-admin* or the *tenant-user* role is used to define a user's role with respect to a specific tenant unit. The tenant user is first assigned the *none* role to minimize access to the primary system, and then either the *tenant-admin* or the *tenant-user* role is appended to that user.

tenant-admin

A *tenant-admin* role can be appended to the other (non-tenant) roles when the Secure Multi-Tenancy (SMT) feature is enabled. A *tenant-admin* user can configure and monitor a specific tenant unit.

tenant-user

A *tenant-user* role can be appended to the other (non-tenant) roles when the SMT feature is enabled. The *tenant-user* role enables a user to monitor a specific tenant unit and change the user password. Users who are assigned the *tenant-*

user management role can view tenant unit status, but they cannot change the tenant unit configuration.

Access management for IP protocols

This feature manages system access for the FTP, FTPS, HTTP, HTTPS, SSH, SCP, and Telnet protocols.

Viewing the IP services configuration

The Administrator Access tab displays the configuration status for the IP protocols that can be used to access the system. FTP and FTPS are the only protocols that are restricted to administrators.

Procedure

1. Select **Administration > Access > Administrator Access**.

Results

The Access Management page displays the Administrator Access, Local Users, Authentication, and Active Users tabs.

Table 33 Administrator Access tab information

Item	Description
Passphrase	If no passphrase is set, the Set Passphrase button appears. If a passphrase is set, the Change Passphrase button appears.
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 Configure . 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 Configure . 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.
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. If possible, set a session timeout maximum of five minutes. Use the Advanced tab of the dialog box to set a timeout in seconds.

Managing FTP access

The File Transfer Protocol (FTP) allows administrators to access files on the Data Domain system.

You can enable either FTP or FTPS access to users who are assigned the admin management role. FTP access allows admin 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

Only users who are assigned the admin management role are permitted to access the system using FTP

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. Select **Administration > Access > Administrator Access**.
 2. Select **FTP** and click **Configure**.
 3. To manage FTP access and which hosts can connect, select the **General** tab and do the following:
 - a. To enable FTP access, select **Allow FTP Access**.
 - b. To enable all hosts to connect, select **Allow all hosts to connect**.
 - c. To restrict access to select hosts, select **Limit Access to the following systems**, and modify the **Allowed Hosts** list.
-

Note

You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address.

- To add a host, click **Add (+)**. Enter the host identification and click **OK**.
 - To modify a host ID, select the host in the **Hosts** list and click **Edit** (pencil). Change the host ID and click **OK**.
 - To remove a host ID, select the host in the **Hosts** list and click **Delete (X)**.
4. To set a session timeout, select the **Advanced** tab, and enter the timeout value in seconds.
-

Note

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

5. Click **OK**.

If FTPS is enabled, a warning message appears with a prompt to click **OK** to proceed.

Managing FTPS access

The FTP Secure (FTPS) protocol allows administrators to access files on the Data Domain system.

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. Consider the following guidelines when using FTPS.

- Only users who are assigned the admin management role are permitted to access the system using FTPS.
- When you enable FTPS access, FTP access 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.
- When you issue the `get` command, the fatal error message `SSL_read: wrong version number lftp` appears if matching versions of SSL are not installed on the Data Domain system and compiled on the LFTP client. As a workaround, attempt to re-issue the `get` command on the same file.

Procedure

1. Select **Administration > Access > Administrator Access**.
2. Select **FTPS** and click **Configure**.
3. To manage FTPS access and which hosts can connect, select the **General** tab and do the following:
 - a. To enable FTPS access, select **Allow FTPS Access**.
 - b. To enable all hosts to connect, select **Allow all hosts to connect**.
 - c. To restrict access to select hosts, select **Limit Access to the following systems**, and modify the hosts list.

Note

You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address.

- To add a host, click **Add (+)**. Enter the host identification and click **OK**.
 - To modify a host ID, select the host in the **Hosts** list and click **Edit** (pencil). Change the host ID and click **OK**.
 - To remove a host ID, select the host in the **Hosts** list and click **Delete (X)**.
4. To set a session timeout, select the **Advanced** tab and enter the timeout value in seconds.

Note

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

5. Click **OK**. If FTP is enabled, a warning message appears and prompts you to click **OK** to proceed.

Managing HTTP and HTTPS access

HTTP or HTTPS access is required to support browser access to DD System Manager.

Procedure

1. Select **Administration > Access > Administrator Access**.
2. Select **HTTP** or **HTTPS** and click **Configure**.
The Configure HTTP/HTTPS Access dialog appears and displays tabs for general configuration, advanced configuration, and certificate management.
3. To manage the access method and which hosts can connect, select the **General** tab and do the following:
 - a. Select the checkboxes for the access methods you want to allow.
 - b. To enable all hosts to connect, select **Allow all hosts to connect**.
 - c. To restrict access to select hosts, select **Limit Access to the following systems**, and modify the host list.

Note

You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address.

- To add a host, click **Add (+)**. Enter the host identification and click **OK**.
 - To modify a host ID, select the host in the **Hosts** list and click **Edit** (pencil). Change the host ID and click **OK**.
 - To remove a host ID, select the host in the **Hosts** list and click **Delete (X)**.
4. To configure system ports and session timeout values, select the **Advanced** tab, and complete the form.
 - In the **HTTP Port** box, enter the port number. Port 80 is assigned by default.
 - In the **HTTPS Port** box, enter the number. Port 443 is assigned by default.
 - In the **Session Timeout** box, enter the interval in seconds that must elapse before a 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 host certificates for HTTP and HTTPS

A host certificate allows browsers to verify the identity of the system when establishing management sessions.

Requesting a host certificate for HTTP and HTTPS

You can use DD System Manager to generate a host certificate request, which you can then forward to a Certificate Authority (CA).

Note

You must configure a system passphrase (system passphrase set) before you can generate a CSR.

Procedure

1. Select **Administration > Access > Administrator Access**.
2. In the Services area, select **HTTP** or **HTTPS** and click **Configure**.
3. Select the **Certificate** tab.
4. Click **Add**.

A dialog appears for the protocol you selected earlier in this procedure.

5. Click **Generate the CSR for this Data Domain system**.

The dialog expands to display a CSR form.

Note

DD OS supports one active CSR at a time. After a CSR is generated, the **Generate the CSR for this Data Domain system** link is replaced with the **Download the CSR for this Data Domain system** link. To delete a CSR, use the `adminaccess certificate cert-signing-request delete` CLI command.

6. Complete the CSR form and click **Generate and download a CSR**.

The CSR file is saved at the following path: `/ddvar/certificates/CertificateSigningRequest.csr`. Use **SCP**, **FTP** or **FTPS** to transfer the CSR file from the system to a computer from which you can send the CSR to a CA.

Adding a host certificate for HTTP and HTTPS

You can use DD System Manager to add a host certificate to the system.

Procedure

1. If you did not request a host certificate, request a host certificate from a certificate authority.
2. When you receive a host certificate, copy or move it to the computer from which you run DD Service Manager.
3. Select **Administration > Access > Administrator Access**.
4. In the Services area, select **HTTP** or **HTTPS** and click **Configure**.
5. Select the **Certificate** tab.
6. Click **Add**.

A dialog appears for the protocol you selected earlier in this procedure.

7. To add a host certificate enclosed in a .p12 file, do the following:

- a. Select **I want to upload the certificate as a .p12 file**.
 - b. Type the password in the **Password** box.
 - c. Click **Browse** and select the host certificate file to upload to the system.
 - d. Click **Add**.
8. To add a host certificate enclosed in a .pem file, do the following:
- a. Select **I want to upload the public key as a .pem file and use a generated private key**.
 - b. Click **Browse** and select the host certificate file to upload to the system.
 - c. Click **Add**.

Deleting a host certificate for HTTP and HTTPS

DD OS supports one host certificate for HTTP and HTTPS. If the system is currently using a host certificate and you want to use a different host certificate, you must delete the current certificate before adding the new certificate.

Procedure

1. Select **Administration > Access > Administrator Access**.
2. In the Services area, select **HTTP** or **HTTPS** and click **Configure**.
3. Select the **Certificate** tab.
4. Select the certificate you want to delete.
5. Click **Delete**, and click **OK**.

Managing SSH and SCP access

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

Procedure

1. Select **Administration > Access > Administrator Access**.
2. Select **SSH** or **SCP** and click **Configure**.
3. To manage the access method and which hosts can connect, select the **General** tab.
 - a. Select the checkboxes for the access methods you want to allow.
 - b. To enable all hosts to connect, select **Allow all hosts to connect**.
 - c. To restrict access to select hosts, select **Limit Access to the following systems**, and modify the host list.

Note

You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address.

- To add a host, click Add (+). Enter the host identification and click **OK**.
- To modify a host ID, select the host in the **Hosts** list and click Edit (pencil). Change the host ID and click **OK**.
- To remove a host ID, select the host in the **Hosts** list and click Delete (X).

4. To configure system ports and session timeout values, click the **Advanced** tab.
 - In the **SSH/SCP** Port text entry box, enter the port number. Port 22 is assigned by default.
 - In the **Session Timeout** 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.

5. Click **OK**.

Managing Telnet access

Telnet is an insecure protocol that enables network access to the system CLI.

Note

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

Procedure

1. Select **Administration > Access > Administrator Access**.
2. Select **Telnet** and click **Configure**.
3. To manage Telnet access and which hosts can connect, select the **General** tab.
 - a. To enable Telnet access, select **Allow Telnet Access**.
 - b. To enable all hosts to connect, select **Allow all hosts to connect**.
 - c. To restrict access to select hosts, select **Limit Access to the following systems**, and modify the host list.

Note

You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address.

- To add a host, click Add (+). Enter the host identification and click **OK**.
 - To modify a host ID, select the host in the **Hosts** list and click Edit (pencil). Change the host ID and click **OK**.
 - To remove a host ID, select the host in the **Hosts** list and click Delete (X).
4. To set a session timeout, select the **Advanced** tab and enter the timeout value in seconds.
-

Note

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

5. Click **OK**.

Local user account management

A local user is a user account (user name and password) that is configured on the Data Domain system instead of being defined in a Windows Active Directory, Windows Workgroup, or NIS directory.

UID conflicts: local user and NIS user accounts

When you set up a Data Domain system in an NIS environment, be aware of potential UID conflicts between local and NIS user accounts.

Local user accounts on a Data Domain system start with a UID of 500. To avoid conflicts, consider the size of potential local accounts when you define allowable UID ranges for NIS users.

Viewing local user information

Local users are user accounts that are defined on the system, rather than in Active Directory, a Workgroup, or UNIX. You can display the local user's username, management role, login status, and target disable date. You can also display the user's password controls and the tenant units the user can access.

Note

The user-authentication module uses Greenwich Mean Time (GMT). To ensure that user accounts and passwords expire correctly, configure settings to use the GMT that corresponds to the target local time.

Procedure

1. Select **Administration > Access > Local Users** .

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

Table 34 Local user list column label descriptions

Item	Description
Name	The user ID, as added to the system.
Management Role	The role displayed is admin, user, security, backup-operator, or none. In this table, Tenant user roles are displayed as <i>none</i> . To see an assigned tenant role, select the user and view the role in the Detailed Information area.
Status	<ul style="list-style-type: none"> • Active—User access to the account is permitted. • Disabled—User access to the account is denied because the account is administratively disabled, the current date is beyond the account expiration date, or a locked account's password requires renewal. • Locked—User access is denied because the password expired.
Disable Date	The date the account is set to be disabled.
Last Login From	The location where the user last logged in.

Table 34 Local user list column label descriptions (continued)

Item	Description
Last Login Time	The time the user last logged in.

Note

User accounts configured with the admin or security officer roles can view all users. Users with other roles can view only their own user accounts.

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

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

Table 35 Detailed User Information, Row Label Descriptions

Item	Description
Tenant-User	The list of tenant units the user can access as a tenant-user role user.
Tenant-Admin	The list of tenant units the user can access as a tenant-admin role user.
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 90.
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 by selecting **More Tasks > Change Login Options**.

Creating local users

Create local users when you want to manage access on the local system instead of through an external directory. Data Domain systems support a maximum of 500 local user accounts.

Procedure

1. Select **Administration > Access > Local Users**.

The Local Users view appears.

2. Click **Create** to create a new user.

The Create User dialog appears.

3. Enter user information in the General Tab.

Table 36 Create User dialog, general controls

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.
Management Role	The role assigned to the user, which can be admin, user, security, backup-operator, or none. .
Note	
Only the sysadmin user (the default user created during the DD OS installation) can create the first security-role user. After the first security-role user is created, only security-role users can create other security-role users.	
Force Password Change	Select this checkbox to require that the user change the password during the first login when logging in to DD System Manager or to the CLI with SSH or Telnet.

The default value for the minimum length of a password is 6 characters. The default value for the 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

Sysadmin is the default admin-role user and cannot be deleted or modified.

4. To manage password and account expiration, select the Advanced tab and use the controls described in the following table.

Table 37 Create User dialog, advanced controls

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 90.
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.

Table 37 Create User dialog, advanced controls (continued)

Item	Description
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.

5. Click **OK**.

Note

Note: The default password policy can change if an admin-role user changes them (**More Tasks > Change Login Options**). The default values are the initial default password policy values.

Modifying a local user profile

After you create a user, you can use DD System Manager to modify the user configuration.

Procedure

1. Select **Administration > Access > Local Users**.
The Local Users view appears.
2. Click a user name from the list.
3. Click **Modify** to make changes to a user account.
The Modify User dialog box appears.
4. Update the information on the General tab.

Note

If SMT is enabled and a role change is requested from none to any other role, the change is accepted only if the user is not assigned to a tenant-unit as a management-user, is not a DD Boost user with its default-tenant-unit set, and is not the owner of a storage-unit that is assigned to a tenant-unit.

Note

To change the role for a DD Boost user that does not own any storage units, unassign it as a DD Boost user, change the user role, and re-assign it as a DD Boost user again.

Table 38 Modify User dialog, general controls

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

5. Update the information on the Advanced tab.

Table 39 Modify User dialog, advanced controls

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 90.
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.

6. 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.

The sysadmin user cannot be deleted. Admin users cannot delete security officers. Only security officers can delete, enable, and disable other security officers.

Procedure

1. Select **Administration > Access > Local Users**.
The Local Users view appears.
2. Click one or more user names from the list.
3. Click **Delete** to delete the user accounts.
The Delete User dialog box appears.
4. Click **OK** and **Close**.

Enabling and disabling local users

Admin users can enable or disable all users except the sysadmin user and users with the security role. The sysadmin user cannot be disabled. Only Security officers can enable or disable other security officers.

Procedure

1. Select **Administration > Access > Local Users**.
The Local Users view appears.
2. Click one or more user names from the list.
3. Click either **Enable** or **Disable** to enable or disable user accounts.
The Enable or Disable User dialog box appears.
4. Click **OK** and **Close**.

Enabling security authorization

You can use the Data Domain system command-line interface (CLI) to enable and disable the security authorization policy.

For information on the commands used in this procedure, see the *Data Domain Operating System Command Reference Guide*.

Note

The DD Retention Lock Compliance license must be installed. You are not permitted to disable the authorization policy on DD Retention Lock Compliance systems.

Procedure

1. Log into the CLI using a security officer username and password.
2. To enable the security officer authorization policy, enter: `# authorization policy set security-officer enabled`

Changing user passwords

After you create a user, you can use DD System Manager to change the user's password. Individual users can also change their own passwords.

Procedure

1. Click **Administration > Access > Local Users**.
The Local Users view appears.
2. Click a user name from the list.
3. Click **Change Password** to change the user password.
The Change Password dialog box appears.
If prompted, enter your old password.
4. Enter the new password into the **New Password** box.
5. Enter the new password again into **Verify New Password** box.
6. Click **OK**.

Modifying the password policy and login controls

The password policy and login controls define login requirements for all users. Administrators can specify how often a password must be changed, what is required to create a valid password, and how the system responds to invalid login attempts.

Procedure

1. Select **Administration > Access**.
2. Select **More Tasks > Change Login Options**.
The Change Login Options dialog appears.
3. Specify the new configuration in the boxes for each option. To select the default value, click **Default** next to the appropriate option.
4. Click **OK** to save the password settings.

Change Login Options dialog

Use this dialog to set the password policy and specify the maximum login attempts and lockout period.

Table 40 Change Login Options dialog controls

Item	Description
Minimum Days Between Change	The minimum number of days between password changes that you allow a user. This value must be less than the Maximum Days Between Change value minus the Warn Days Before Expire value. The default setting is 0.
Maximum Days Between Change	The maximum number of days between password changes that you allow a user. The minimum value is 1. The default value is 90.
Warn Days Before Expire	The number of days to warn the users before their password expires. This value must be less than the Maximum Days Between Change value minus the Minimum Days Between Change value. The default setting is 7.
Disable Days After Expire	The system disables a user account after password expiration according to the number of days specified with this option. Valid entries are <i>never</i> or number greater than or equal to 0. The default setting is never.
Minimum Length of Password	The minimum password length required. Default is 6.
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"> • Lowercase letters (a-z) • Uppercase letters (A-Z) • Numbers (0-9) • Special Characters (\$, %, #, +, and so on)
Lowercase Character Requirement	Enable or disable the requirement for at least one lowercase character. The default setting is disabled.
Uppercase Character Requirement	Enable or disable the requirement for at least one uppercase character. The default setting is disabled.
One Digit Requirement	Enable or disable the requirement for at least one numerical character. The default setting is disabled.
Special Character Requirement	Enable or disable the requirement for at least one special character. The default setting is disabled.
Max Consecutive Character Requirement	Enable or disable the requirement for a maximum of three repeated characters. The default setting is disabled.
Prevent use of Last N Passwords	Specify the number of remembered passwords. The range is 0 to 24, and the default settings is 1.

Table 40 Change Login Options dialog controls (continued)

Item	Description
	<p>Note</p> <p>If this setting is reduced, the remembered password list remains unchanged until the next time the password is changed. For example, if this setting is changed from 4 to 3, the last four passwords are remembered until the next time the password is changed.</p>
Maximum login attempts	Specifies the maximum number of login attempts before a mandatory lock is applied to a user account. This limit applies to all user accounts, including sysadmin. A locked user cannot log in while the account is locked. The range is 4 to 10, and the default value is 4.
Unlock timeout (seconds)	Specifies how long a user account is locked after the maximum number of login attempts. When the configured unlock timeout is reached, a user can attempt login. The range is 120 to 600 seconds, and the default period is 120 seconds.

Directory user and group management

You can use DD System Manager to manage access to the system for users and groups in Windows Active Directory, Windows Workgroup, and NIS. Kerberos authentication is an option for CIFS and NFS clients.

Viewing Active Directory and Kerberos information

The Active Directory Kerberos configuration determines the methods CIFS and NFS clients use to authenticate. The Active Directory/Kerberos Authentication panel displays this configuration.

Procedure

1. Select **Administration > Access > Authentication**.
2. Expand the Active Directory/Kerberos Authentication panel.

Table 41 Active Directory/ Kerberos Authentication label descriptions

Item	Description
Mode	The type of authentication mode. In Windows/Active Directory mode, CIFS clients use Active Directory and Kerberos authentication, and NFS clients use Kerberos authentication. In Unix mode, CIFS clients use Workgroup authentication (without Kerberos), and NFS clients use Kerberos authentication. In Disabled mode, Kerberos authentication is disabled and CIFS clients use Workgroup authentication.
Realm	The realm name of the Workgroup or Active Directory.
DDNS	Whether or not the Dynamic Domain Name System is enabled.

Table 41 Active Directory/ Kerberos Authentication label descriptions (continued)

Item	Description
Domain Controllers	The name of the domain controller for the Workgroup or Active Directory.
Organizational Unit	The name of the organizations unit for the Workgroup or Active Directory.
CIFS Server Name	The name of the CIFS server in use (Windows mode only).
WINS Server	The name of the WINS server in use (Windows mode only).
Short Domain Name	An abbreviated name for the domain.
NTP	Enabled/Disabled (UNIX mode only)
NIS	Enabled/Disabled (UNIX mode only)
Key Distribution Centers	Hostname(s) or IP(s) of KDC in use (UNIX mode only)
Active Directory Administrative Access	Enabled/Disabled: Click to Enable or disable administrative access for Active Directory (Windows) groups.

Table 42 Active Directory administrative groups and roles

Item	Description
Windows Group	The name of the Windows group.
Management Role	The role of the group (admin, user, and so on)

Configuring Active Directory and Kerberos authentication

Configuring Active Directory authentication makes the Data Domain system part of a Windows Active Directory realm. CIFS clients and NFS clients use Kerberos authentication.

Procedure

1. Select **Administration > Access > Authentication**.
The Authentication view appears.
2. Expand the Active Directory/Kerberos Authentication panel.
3. Click **Configure...** next to Mode to start the configuration wizard.
The Active Directory/Kerberos Authentication dialog appears.
4. Select **Windows/Active Directory** and click **Next**.
5. Enter the full realm name for the system (for example: domain1.local), the user name, and password for the Data Domain system. Then click **Next**.

Note

Use the complete realm name. Ensure that the user is assigned sufficient privileges to join the system to the domain. The user name and password must be compatible with Microsoft requirements for the Active Directory domain. This user must also be assigned permission to create accounts in this domain.

6. Select the default CIFS server name, or select **Manual** and enter a CIFS server name.

7. To select domain controllers, select **Automatically assign**, or select **Manual** and enter up to three domain controller names.
You can enter fully qualified domain names, hostnames, or IP (IPv4 or IPv6) addresses.
8. To select an organizational unit, select **Use default Computers**, or select **Manual** and enter an organization unit name.

Note

The account is moved to the new organizational unit.

9. Click **Next**.
The Summary page for the configuration appears.
10. Click **Finish**.
The system displays the configuration information in the Authentication view.
11. To enable administrative access, click **Enable** to the right of **Active Directory Administrative Access**.

Authentication mode selections

The authentication mode selection determines how CIFS and NFS clients authenticate using supported combinations of Active Directory, Workgroup, and Kerberos authentication.

DD OS supports the following authentication options.

- **Disabled:** Kerberos authentication is disabled for CIFS and NFS clients. CIFS clients use Workgroup authentication.
- **Windows/Active Directory:** Kerberos authentication is enabled for CIFS and NFS clients. CIFS clients use Active Directory authentication.
- **Unix:** Kerberos authentication is enabled for only NFS clients. CIFS clients use Workgroup authentication.

Managing administrative groups for Active Directory

You can use the Active Directory/Kerberos Authentication panel to create, modify, and delete Active Directory (Windows) groups and assign management roles (admin, backup-operator, and so on) to those groups.

To prepare for managing groups, select **Administration** > **Access** > **Authentication**, expand the Active Directory/Kerberos Authentication panel, and click the Active Directory Administrative Access **Enable** button.

Creating administrative groups for Active Directory

Create an administrative group when you want to assign a management role to all the users configured in an Active Directory group.

Before you begin

Enable Active Directory Administrative Access on the Active Directory/Kerberos Authentication panel in the **Administration** > **Access** > **Authentication** page.

Procedure

1. Click **Create...**
2. Enter the domain and group name separated by a backslash. For example: domainname\groupname.

3. Select the management role for the group from the drop-down menu.
4. Click **OK**.

Modifying administrative groups for Active Directory

Modify an administrative group when you want to change the administrative group name or management role configured for an Active Directory group.

Before you begin

Enable Active Directory Administrative Access on the Active Directory/Kerberos Authentication panel in the **Administration > Access > Authentication** page.

Procedure

1. Select a group to modify under the **Active Directory Administrative Access** heading.
2. Click **Modify...**
3. Modify the domain and group name. These names are separated by a backslash. For example: domainname\groupname.
4. Modify the management role for the group by selecting a different role from the drop-down menu.

Deleting administrative groups for Active Directory

Delete an administrative group when you want to terminate system access for all the users configured in an Active Directory group.

Before you begin

Enable Active Directory Administrative Access on the Active Directory/Kerberos Authentication panel in the **Administration > Access > Authentication** page.

Procedure

1. Select a group to delete under the **Active Directory Administrative Access** heading.
2. Click **Delete**.

Configuring UNIX Kerberos authentication

Configuring UNIX Kerberos authentication enables NFS clients to use Kerberos authentication. CIFS clients use Workgroup authentication.

Before you begin

NIS must be running for UNIX-mode Kerberos authentication to function. For instructions about enabling Kerberos, see the section regarding enabling NIS services.

Configuring Kerberos for UNIX enables NFS clients to use Kerberos authentication. CIFS clients use Workgroup authentication.

Procedure

1. Select **Administration > Access > Authentication**.
The Authentication view appears.
2. Expand the Active Directory/Kerberos Authentication panel.
3. Click **Configure...** next to Mode to start the configuration wizard.
The Active Directory/Kerberos Authentication dialog appears.

4. Select **Unix** and click **Next**.
5. Enter the realm name (for example: domain1.local), and up to three host names or IP addresses (IPv4 or IPv6) for key distribution centers (KDCs).
6. Optionally, click **Browse** to upload a keytab file, and click **Next**.

The Summary page for the configuration appears.

Note

Keytab files are generated on the authentication servers (KDCs) and contain a shared secret between the KDC server and the DDR.

NOTICE

A keytab file must be uploaded and imported for Kerberos authentication to operate correctly.

7. Click **Finish**.

The system displays the configuration information in the Active Directory/Kerberos Authentication panel.

Disabling Kerberos authentication

Disabling Kerberos authentication prevents CIFS and NFS clients from using Kerberos authentication. CIFS clients use Workgroup authentication.

Procedure

1. Select **Administration > Access Management > Authentication**.

The Authentication view appears.

2. Expand the Active Directory/Kerberos Authentication panel.
3. Click **Configure...** next to Mode to start the configuration wizard.

The Active Directory/Kerberos Authentication dialog appears.

4. Select **Disabled** and click **Next**.

The system displays a summary page with changes appearing in bold text.

5. Click **Finish**.

The system displays Disabled next to Mode in the Active Directory/Kerberos Authentication panel.

Viewing Workgroup authentication information

Use the Workgroup Authentication panel to view Workgroup configuration information.

Procedure

1. Select **Administration > Access > Authentication**.
2. Expand the Workgroup Authentication panel.

Table 43 Workgroup Authentication label descriptions

Item	Description
Mode	The type of authentication mode (Workgroup or Active Directory).

Table 43 Workgroup Authentication label descriptions (continued)

Item	Description
Workgroup name	The specified workgroup
CIFS Server Name	The name of the CIFS server in use.
WINS Server	The name of the WINS server in use.

Configuring workgroup authentication parameters

Workgroup authentication parameters allow you to configure a Workgroup name and CIFS server name.

Procedure

1. Select **Administration > Access > Authentication**.

The Authentication view appears.

2. Expand the Workgroup Authentication panel.
3. Click **Configure**.

The Workgroup Authentication dialog appears.

4. For Workgroup Name, select **Manual** and enter a workgroup name to join, or use the default.

The Workgroup mode joins a Data Domain system to a workgroup domain.

5. For CIFS Server Name, select **Manual** and enter a server name (the DDR), or use the default.
6. Click **OK**.

Viewing NIS authentication information

The NIS Authentication panel displays the NIS configuration parameters and whether NIS authentication is enabled or disabled.

Procedure

1. Select **Administration > Access > Authentication**.

The Authentication view appears.

2. Expand the NIS Authentication panel.

Results

Table 44 NIS Authentication panel items

Item	Description
NIS Status	Enabled or Disabled.
Domain Name	The name of the domain for this service.
Server	Authentication server(s).
NIS Group	The name of the NIS group.

Table 44 NIS Authentication panel items (continued)

Item	Description
Management Role	The role of the group (admin, user, and so on).

Enabling and disabling NIS authentication

Use the NIS Authentication panel to enable and disable NIS authentication.

Procedure

1. Select **Maintenance > Access > Authentication**.
The Authentication view appears.
2. Expand the NIS Authentication panel.
3. Click **Enable** next to NIS Status to enable or **Disable** to disable NIS Authentication.
The Enable or Disable NIS dialog box appears.
4. Click **OK**.

Configuring the NIS domain name

Use the NIS Authentication panel to configure the NIS domain name.

Procedure

1. Select **Administration > Access > Authentication**.
The Authentication view appears.
2. Expand the NIS Authentication panel.
3. Click **Edit** next to Domain Name to edit the NIS domain name.
The Configure NIS Domain Name dialog box appears.
4. Enter the domain name in the **Domain Name** box.
5. Click **OK**.

Specifying NIS authentication servers

Use the NIS Authentication panel to specify NIS authentication servers.

Procedure

1. Select **Administration > Access > Authentication**.
The Authentication view appears.
2. Expand the NIS Authentication panel.
3. Below Domain Name, select one of the following:
 - **Obtain NIS Servers from DHCP** The system automatically obtains NIS servers using DHCP
 - **Manually Configure** Use the following procedures to manually configure NIS servers.

- To add an authentication server, click Add (+) in the server table, enter the server name, and click **OK**.
 - To modify an authentication server, select the authentication server name and click the edit icon (pencil). Change the server name, and click **OK**.
 - To remove an authentication server name, select a server, click the X icon, and click **OK**.
4. Click **OK**.

Configuring NIS groups

Use the NIS Authentication panel to configure NIS groups.

Procedure

1. Select **Administration > Access > Authentication**.
The Authentication view appears.
2. Expand the NIS Authentication panel.
3. Configure the NIS groups in the NIS Group table.
 - To add a 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**.
4. Click **OK**.

Change system authentication method

The Data Domain system supports password-based authentication, or certificate-based authentication. Password-based authentication is the default method.

Before you begin

Certificate-based authentication requires SSH keys and CA certificates are imported to allow users to authenticate with the system when password-based authentication is disabled.

Complete the following steps to change the system authentication method from password-based authentication to certificate-based authentication.

Procedure

1. Select **Administration > Access**.
The Access Management view appears.
2. Click **Manage CA Certificates**.
3. Click **Add** to create a new certificate.
4. Add the certificate.
 - Select **I want to upload the certificate as a .pem file** and click **Choose File** to select the certificate file and upload it to the system.

- Select **I want to copy and paste the certificate text** to copy and paste the certificate text into the text field.
5. Click **Add**.
 6. Select **More Tasks > Change Login Options**.
 7. In the **Password Based Login** drop-down menu, select **Disable**.

Note

The drop-down menu is disabled if the required SSH keys and CA certificates are not configured on the system

8. Click **OK**.
If a security policy is configured, the system prompts for security officer credentials. Provide the credentials and click **OK**.

Reset the system authentication method to password-based authentication.

Complete the following steps to change the system authentication method from certificate-based authentication to password-based authentication.

Procedure

1. Select **Administration > Access**.
The Access Management view appears.
2. Select **More Tasks > Change Login Options**.
3. In the **Password Based Login** drop-down menu, select **Enable**.
4. Click **OK**.
If a security policy is configured, the system prompts for security officer credentials. Provide the credentials and click **OK**.

Configuring mail server settings

The Mail Server tab allows you to specify the mail server to which DD OS sends email reports.

Procedure

1. Select **Administration > Settings > 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**.

Managing time and date settings

The Time and Date Settings tab allows you to view and configure the system time and date or configure the Network Time Protocol to set the time and date.

Procedure

1. To view the current time and date configuration, select **Administration > Settings > Time and Date Settings**.

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

2. To change the configuration, select **More Tasks > Configure Time Settings**.
The Configure Time Settings dialog appears.
3. In the **Time Zone** dropdown list, select the time zone where the Data Domain system resides.
4. To manually set the time and date, select **None**, type the date in the **Date** box, and select the time in the **Time** dropdown lists.
5. To use NTP to synchronize the time, select NTP and 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.

6. Click **OK**.
7. 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.

Managing system properties

The System Properties tab allows you to view and configure system properties that identify the managed system location, administrator email address, and host name.

Procedure

1. To view the current configuration, select **Administration > Settings > System Properties**.
The System Properties tab displays the system location, the administrator email address, and the administrator hostname.
2. To change the configuration, select **More Tasks > Set System Properties**.
The Set System Properties dialog box appears.
3. In the **Location** box, enter information about where the Data Domain system is located.
4. In the **Admin Email** box, enter the email address of the system administrator.
5. In the **Admin Host** box, enter the name of the administration server.

6. Click **OK**.

SNMP management

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. DD OS also supports 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.

The Data Domain system SNMP agent accepts queries for Data Domain-specific information from management systems using SNMP v1, v2c, and v3. SNMP V3 provides a greater degree of security than v2c and v1 by replacing cleartext community strings (used for authentication) with user-based authentication using either MD5 or SHA1. Also, SNMP v3 user authentication packets can be encrypted and their integrity verified with either DES or AES.

Data Domain systems can send SNMP traps (which are alert messages) using SNMP v2c and SNMP v3. Because SNMP v1 traps are not supported, if possible, use SNMP v2c or v3.

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

- The *Data Domain Operating System Initial Configuration Guide* describes how to set up the Data Domain system to use SNMP monitoring.
- The *Data Domain Operating System MIB Quick Reference* describes the full set of MIB parameters included in the Data Domain MIB branch.

Viewing SNMP status and configuration

The SNMP tab displays the current SNMP status and configuration.

Procedure

1. Select **Administration > Settings > SNMP**.

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

SNMP tab labels

The SNMP tab labels identify the overall SNMP status, SNMP property values, and the configurations for SNMPv3 and SNMPv2.

Status

The Status area displays the operational status of the SNMP agent on the system, which is either Enabled or Disabled.

SNMP Properties

Table 45 SNMP Properties descriptions

Item	Description
SNMP System Location	The location of the Data Domain system being monitored.
SNMP System Contact	The person designated as the person to contact for the Data Domain system administration.
SNMP System Notes	(Optional) Additional SNMP configuration data.
SNMP Engine ID	A unique hexadecimal identifier for the Data Domain system.

SNMP V3 Configuration

Table 46 SNMP Users column descriptions

Item	Description
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, which can be Read-only or Read-write.
Authentication Protocols	The Authentication Protocol used to validate the SNMP user, which can be MD5, SHA1, or None.
Privacy Protocol	The encryption protocol used during the SNMP user authentication, which can be AES, DES, or None.

Table 47 Trap Hosts column descriptions

Item	Description
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

Table 48 Communities column descriptions

Item	Description
Community	The name of the community. For example, public, private, or localCommunity.
Access	The access permission assigned, which can be Read-only or Read-write.
Hosts	The hosts in this community.

Table 49 Trap Hosts column descriptions

Item	Description
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

Use the SNMP tab to enable or disable SNMP.

Procedure

1. Select **Administration > Settings > SNMP**.
2. In the Status area, click **Enable** or **Disable**.

Downloading the SNMP MIB

Use the SNMP tab to download the SNMP MIB.

Procedure

1. Select **Administration > Settings > SNMP**.
2. Click **Download MIB file**.
3. In the Opening DATA_DOMAIN.mib dialog box, select **Open**.
4. 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.

5. Save the MIB or exit the browser.

Configuring SNMP properties

Use the SNMP tab to configure the text entries for system location and system contact.

Procedure

1. Select **Administration > Settings > SNMP**.
2. In the SNMP Properties area, click **Configure**.
The SNMP Configuration dialog box appears.
3. In the text fields, specify the following information: and/or an
 - **SNMP System Location:** A description of where the Data Domain system is located.

- **SNMP System Contact:** The email address of the system administrator for the Data Domain system.
- **SNMP System Notes:** (Optional) Additional SNMP configuration information.
- **SNMP Engine ID:** A unique identifier for the SNMP entity. The engine ID must be 5-34 hexadecimal characters (SNMPv3 only).

Note

The system displays an error if the SNMP engine ID does not meet the length requirements, or uses invalid characters.

4. Click **OK**.

SNMP V3 user management

Use the SNMP tab to create, modify, and delete SNMPv3 users and trap hosts.

Creating SNMP V3 users

When you create SNMPv3 users, you define a username, specify either read-only or read-write access, and select an authentication protocol.

Procedure

1. Select **Administration > Settings > 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 for whom you want to grant access to the Data Domain system agent. 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

You can modify the access level (read-only or read-write) and authentication protocol for existing SNMPv3 users.

Procedure

1. Select **Administration > Settings > 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

Use the SNMP tab to delete existing SNMPv3 users.

Procedure

1. Select **Administration > Settings > 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.

SNMP V2C community management

Define SNMP v2c communities (which serve as passwords) to control management system access to the Data Domain system. To restrict access to specific hosts that use the specified community, assign the hosts to the community.

Note

The SNMP V2c 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. SNMP V3 provides authentication and encryption features to prevent interception.

Note

SNMP community definitions do not enable the transmission of SNMP traps to a management station. You must define trap hosts to enable trap submission to management stations.

Creating SNMP V2C communities

Create communities to restrict access to the DDR system or for use in sending traps to a trap host. You must create a community and assign it to a host before you can select that community for use with the trap host.

Procedure

1. Select **Administration > Settings > 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 for whom you want to grant access to the Data Domain system agent.
 4. Select either read-only or read-write access for this community.
 5. If you want to associate the community to one or more hosts, add the hosts as follows:

- 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 and lists the selected hosts.

Modifying SNMP V2C Communities

Procedure

1. Select **Administration > Settings > SNMP**.
2. In the Communities area, select the checkbox for the community and click **Modify**.

The Modify SNMP V2C Community dialog box appears.

3. To change the access mode for this community, select either **read-only** or **read-write** access.

Note

The Access buttons for the selected community are disabled when a trap host on the same system is configured as part of that community. To modify the access setting, delete the trap host and add it back after the community is modified.

4. To add one or more hosts to this community, do the following:
 - 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. To delete one or more hosts from the host list, do the following:

Note

DD System Manager does not allow you to delete a host when a trap host on the same system is configured as part of that community. To delete a trap host from a community, delete the trap host and add it back after the community is modified.

Note

The Access buttons for the selected community are not disabled when the trap host uses an IPv6 address and the system is managed by an earlier DD OS version that does not support IPv6. If possible, always select a management system that uses the same or a newer DD OS version than the systems it manages.

- a. Select the checkbox for each host or click the Host check box in the table head to select all listed hosts.
- b. Click the delete button (X).
6. To edit a host name, do the following:
 - a. Select the checkbox for the host.
 - b. Click the edit button (pencil icon).
 - c. Edit the host name.
 - d. Click **OK**.
7. Click **OK**.
The modified community entry appears in the Communities table.

Deleting SNMP V2C communities

Use the SNMP tab to delete existing SNMPv2 communities.

Procedure

1. Select **Administration > Settings > 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.

SNMP trap host management

Trap host definitions enable Data Domain systems to send alert messages in SNMP trap messages to an SNMP management station.

Creating SNMP V3 and V2C trap hosts

Trap host definitions identify remote hosts that receive SNMP trap messages from the system.

Before you begin

If you plan to assign an existing SNMP v2c community to a trap host, you must first use the Communities area to assign the trap host to the community.

Procedure

1. Select **Administration > Settings > 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 appears.
3. In the **Host** box, enter the IP address or domain name of the SNMP Host to receive traps.
4. In the **Port** box, 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.

Note

The Community list displays only those communities to which the trap host is already assigned.

6. To create a new community, do the following:
 - a. Select **Create New Community** in the Community drop-down menu.
 - b. Enter the name for the new community in the **Community** box.
 - c. Select the Access type.
 - d. Click the add (+) button.
 - e. Enter the trap host name.
 - f. Click **OK**.
 - g. Click **OK**.
7. Click **OK**.

Modifying SNMP V3 and V2C trap hosts

You can modify the port number and community selection for existing trap host configurations.

Procedure

1. Select **Administration > Settings > 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.

3. To modify the port number, enter a new port number in the **Port** box (port 162 is a common port).
4. Select the user (SNMP V3) or the community (SNMP V2C) from the drop-down menu.

Note

The Community list displays only those communities to which the trap host is already assigned.

5. To create a new community, do the following:
 - a. Select **Create New Community** in the Community drop-down menu.
 - b. Enter the name for the new community in the **Community** box.
 - c. Select the Access type.
 - d. Click the add (+) button.
 - e. Enter the trap host name.
 - f. Click **OK**.
 - g. Click **OK**.
6. Click **OK**.

Removing SNMP V3 and V2C trap hosts

Use the SNMP tab to delete existing trap host configurations.

Procedure

1. Select **Administration > Settings > 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.

Autosupport report management

The Autosupport feature generates a report called an ASUP. The ASUP shows system identification information, consolidated output from a number of Data Domain system

commands, and entries from various log files. Extensive and detailed internal statistics appear at the end of the report. This report is designed to aid Data Domain Support in debugging system problems.

An ASUP is generated every time the file system is started, which is usually once per day. However, the file system can be started more than once in a day.

You can configure email addresses to receive the daily ASUP reports, and you can enable or disable sending of these reports to Data Domain. The default time for sending the daily ASUP is 06.00 a.m, and it is configurable. When sending ASUPs to Data Domain, you have the option to select the legacy unsecure method or the ConnectEMC method, which encrypts the information before transmission.

HA system autosupport and support bundle manageability

Configuration is done on the active node and mirrored to the standby node; therefore, the same configuration is on both nodes, but there is not a consolidated ASUP and support bundle.

Autosupport and support bundle on the active node also includes filesystem, replication, protocol, and full HA information in addition to local node information. Autosupport and support bundle on the standby node only have local node information plus some HA information (configuration and status), but no filesystem/replication/protocol information. The autosupports and support bundles from both the nodes will be needed to debug issues related to HA system status (filesystem, replication, protocols, and HA configuration).

Enabling and disabling autosupport reporting to Data Domain

You can enable or disable autosupport reporting to Data Domain without affecting whether or not alerts are sent to Data Domain.

Procedure

1. To view the autosupport reporting status, select **Maintenance > Support > Autosupport**.

The autosupport reporting status is highlighted next to the Scheduled autosupport label in the Support area. Depending on the current configuration, either an **Enable** or a **Disable** button appears in the Scheduled autosupport row.

2. To enable autosupport reporting to Data Domain, click **Enable** in the Scheduled autosupport row.
3. To disable autosupport reporting to Domain, click **Disable** in the Scheduled autosupport row.

Reviewing generated autosupport reports

Review autosupport reports to view system statistics and configuration information captured in the past. The system stores a maximum of 14 autosupport reports.

Procedure

1. 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.

CLI equivalent


```
# autosupport show history
```

2. 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.

Configuring the autosupport mailing list

Autosupport mailing list subscribers receive autosupport messages through email. Use the Autosupport tab to add, modify, and delete subscribers.

Autosupport emails are sent through the configured mail server 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.

CLI equivalent

```
# autosupport add asup-detailed emails djones@company.com
# autosupport add alert-summary emails djones@company.com
```

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).

CLI equivalent

```
# autosupport del asup-detailed emails djones@company.com
# autosupport del alert-summary emails djones@company.com
```

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.

- Click **OK** to close the Configure Autosupport Subscribers dialog box.

The revised autosupport email list appears in the Autosupport Mailing List area.

Support bundle management

A support bundle is a file that contains system configuration and operation information. It is a good practice to generate a support bundle before a software upgrade or a system topology change (such as a controller upgrade).

Data Domain Support often requests a support bundle when providing assistance.

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.

Procedure

- Select **Maintenance > Support > Support Bundles**.
- Click **Generate Support Bundle**.

Note

The system supports a maximum of five support bundles. If you attempt to generate an sixth support bundle, the system automatically deletes the oldest support bundle. You can also delete support bundles using the CLI command `support bundle delete`.

Also, if you generate a support bundle on a upgraded system that contains a support bundle named using the old format, `support-bundle.tar.gz`, that file is renamed to use the newer name format.

- Email the file to customer support at support@emc.com.

Note

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

Viewing the support bundles list

Use the Support Bundles tab to view the support bundle files on the system.

Procedure

- 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 `hostname-support-bundle-datestamp.tar.gz`. An example filename is `localhost-support-bundle-1127103633.tar.gz`, which indicates that the support bundle was created on the localhost system on November 27th at 10:36:33.

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

Coredump management

When DD OS crashes due to a coredump, a core file describing the problem is created in the `/ddvar/core` directory. This file may be large, and difficult to copy off the Data Domain system.

If the core file cannot be copied off the Data Domain system because it is too large, run the `support coredump split <filename> by <n> {MiB|GiB}` command, where:

- `<filename>` is the name of the core file in the `/ddvar/core` directory
- `<n>` is the number of smaller chunks to break the core file into

Note

A single core file can be broken down into a maximum of 20 chunks. The command will fail with an error if the specified size would result in more than 20 chunks.

For example, splitting a 42.1 MB core file named `cpmdb.core.19297.1517443767` into 10 MB chunks would result in five chunks.

```
# support coredump split cpmdb.core.19297.1517443767 10 MiB
cpmdb.core.19297.1517443767 will be split into 5 chunks.
Splitting...
```

The md5 and split chunks of `cpmdb.core.19297.1517443767`:

File	Size	Time Created
-----	-----	-----
cpmdb.core.19297.1517443767_5_01	10.0 MiB	Mon Feb 5 11:50:57 2018
cpmdb.core.19297.1517443767_5_02	10.0 MiB	Mon Feb 5 11:50:57 2018
cpmdb.core.19297.1517443767_5_03	10.0 MiB	Mon Feb 5 11:50:57 2018
cpmdb.core.19297.1517443767_5_04	10.0 MiB	Mon Feb 5 11:50:57 2018
cpmdb.core.19297.1517443767_5_05	2.1 MiB	Mon Feb 5 11:50:57 2018
cpmdb.core.19297.1517443767.md5	0 MiB	Mon Feb 5 11:50:58 2018
-----	-----	-----

Download the files as soon as possible. Otherwise they will be automatically delete in 48 hours.

Alert notification management

The alert feature generates event and summary reports that can be distributed to configurable email lists and to Data Domain.

Event reports are sent immediately and provide detailed information on a system event. The distribution lists for event alerts are called *notification groups*. You can configure a notification group to include one or more email addresses, and you can configure the types and severity level of the event reports sent to those addresses. For example, you might configure one notification group for individuals who need to know about critical events and another group for those who monitor less critical events. Another option is to configure groups for different technologies. For example, you might configure one notification group to receive email messages about all network events and another group to receive messages about storage issues.

Summary reports are sent daily and provide a summary of the events that occurred during the last 24 hours. Summary reports do not include all the information that is provided in event reports. The default generation time for the daily report is 08.00 a.m, and it can be changed. Summary reports are sent using a dedicated email list that is separate from the event notification groups.

You can enable or disable alert distribution to Data Domain. When sending reports to Data Domain, you have the option to select the legacy unsecure method or Secure Remote Services for secure transmissions.

HA system alert notification management

The alert feature on an HA system generates event and summary report like a non-HA system but how the HA system manages these alerts is different due to the two node system set-up.

Initial alert configuration is completed on the active node and mirrored to the stand-by (i.e., same configuration on both nodes). Local and AM-Alerts are emailed according to the notification settings and include information indicating they are from an HA system and from which node, the active or standby, that generated the alerts.

If there are active alerts on the file system, replication, or protocols when a failover occurs, these active alerts continue to show on the new active node after failover if the alert conditions have not cleared up.

Historical alerts on the filesystem, replication, and protocols stay with the node where they originated rather than failing over together with the filesystem on a failover. This means the CLIs on the active node will not present a complete/continuous view of historical alerts for filesystem, replication, and protocols

During a failover, local historical alerts stay with the node from which they were generated; however, the historical alerts for the filesystem, replication, and protocols (generally called "logical alerts") fail over together with the filesystem.

Note

The **Health > High Availability** panel displays only alerts that are HA-related. Those alerts can be filtered by major HA component, such as HA Manager, Node, Interconnect, Storage, and SAS connection.

Viewing the notification group list

A notification group defines a set of alert types (classes) and a group of email addresses (for subscribers). Whenever the system generates an alert type selected in a notification list, that alert is sent to the list subscribers.

Procedure

1. Select **Health > Alerts > Notification**.

CLI equivalent

```
# alerts notify-list show
```

2. To limit (filter) the entries in the Group Name list, type a group name in the Group Name box or a subscriber email in the Alert Email box, and click **Update**.
-

Note

Click **Reset** to display all configured groups.

3. To display detailed information for a group, select the group in the Group Name list.

Notification tab

The Notification tab allows you to configure groups of email address that receive system alerts for the alert types and severity levels you select.

Table 50 Group Name list, column label descriptions

Item	Description
Group Name	The configured name for the group.
Classes	The number of alert classes that are reported to the group.
Subscribers	The number of subscribers who are configured to receive notifications through email.

Table 51 Detailed Information, label descriptions

Item	Description
Class	A service or subsystem that can forward alerts. The listed classes are those for which the notification group receives alerts.
Severity	The severity level that triggers an email to the notification group. All alerts at the specified severity level and above are sent to the notification group.
Subscribers	The subscribers area displays a list of all email addresses configured for the notification group.

Table 52 Notification tab controls

Control	Description
Add button	Click the Add button to begin creating a notification group.
Class Attributes Configure button	Click this Configure button to change the classes and severity levels that generate alerts for the selected notification group.
Delete button	Click the Delete button to delete the selected notification group.
Filter By: Alert Email box	Enter text in this box to limit the group name list entries to groups that include an email address that contains the specified text.
Filter By: Group Name box	Enter text in this box to limit the group name list entries to group names that contain the specified text.
Modify button	Click the Modify button to modify the configuration for the selected notification group.
Reset button	Click this button to remove any entries in the Filter By boxes and display all group names.

Table 52 Notification tab controls (continued)

Control	Description
Subscribers Configure button	Click this Configure button to change the email list for the selected notification group.
Update button	Click this button to update the group name list after you enter text in a filter box.

Creating a notification group

Use the Notification tab to add notification groups and select the severity level for each group.

Procedure

1. Select **Health > Alerts > Notification**.
2. Click **Add**.
The Add Group dialog box appears.
3. Type the group name in the **Group Name** box.
4. Select the checkbox of one or more alert classes of which to be notified.
5. To change the default severity level (Warning) for a class, select another level in the associated list box.

The severity levels are listed in ascending severity level. *Emergency* is the highest severity level.

6. Click **OK**.

CLI equivalent

```
# alerts notify-list create eng_grp class hardwareFailure
```

Managing the subscriber list for a group

Use the Notification tab to add, modify, or delete email addresses from a notification group subscriber list.

Procedure

1. Select **Health > Alerts > Notification**.
2. Select the checkbox of a group in the Notifications group list, and do one of the following.
 - Click **Modify** and select **Subscribers**.
 - Click **Configure** in the Subscribers list.
3. To add a subscriber to the group, do the following.
 - a. Click the + icon.
The Email Address dialog box appears.
 - b. Enter the email address of a subscriber.
 - c. Click **OK**.

CLI equivalent

```
# alerts notify-list add eng_lab emails
mlee@urcompany.com,bob@urcompany.com
```

4. To modify an email address, do the following.
 - a. Click the checkbox of the email address in the **Subscriber Email** list.
 - b. Click the pencil icon.
 - c. Edit the email address in the Email Address dialog box.
 - d. Click **OK**.
5. To delete an email address, click the checkbox of the email address in the **Subscriber Email** list and click the **X** icon.

CLI equivalent

```
# alerts notify-list del eng_lab emails bob@urcompany.com
```

6. Click **Finish** or **OK**.

Modifying a notification group

Use the Notification table to modify the attribute classes in an existing group.

Procedure

1. Select **Health > Alerts > Notification**.
2. Select the checkbox of the group to modify in the group list.
3. 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**.

CLI equivalent

```
# alerts notify-list add eng_lab class cloud severity warning
# alerts notify-list del eng_lab class cloud severity notice
```

4. 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**.

CLI equivalent

```
# alerts notify-list add eng_lab emails
mlee@urcompany.com,bob@urcompany.com
# alerts notify-list del eng_lab emails bob@urcompany.com
```

5. Click **OK**.

Deleting a notification group

Use the Notification tab to delete one or more existing notification groups.

Procedure

1. Select **Health > Alerts > Notification**.
2. Select one or more checkboxes of groups in the Notifications group list, and click **Delete**.

The Delete Group dialog box appears.

3. Verify the deletion and click **OK**.

CLI equivalent

```
# alerts notify-list destroy eng_grp
```

Resetting the notification group configuration

Use the Notification tab to remove all notification groups added and to remove any changes made to the Default group.

Procedure

1. Select **Health > Alerts > Notification**.
2. Select **More Tasks > Reset Notification Groups**.
3. In the Reset Notification Groups dialog box, click **Yes** in the verification dialog.

CLI equivalent

```
# alerts notify-list reset
```

Configuring the daily summary schedule and distribution list

Every day, 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 you might want to address soon.

A fan failure is an example of a noncritical issue that you might want to address as soon as is reasonably possible. When Support receives the failure notification, they contact you to arrange for component replacement.

Procedure

1. Select **Health > Alerts > Daily Alert Summary**.
2. If the default deliver time of 8 AM is not acceptable, do the following.
 - a. Click **Schedule**.

The Schedule Alert Summary dialog box appears.

- b. Use the list boxes to select the hour, minute, and either AM or PM for the summary report.
- c. Click **OK**.

CLI equivalent

```
# autosupport set schedule alert-summary daily 1400
```

- 3. To configure the daily alert subscriber list, do the following.

- a. Click **Configure**.

The Daily Alert Summary Mailing List dialog box appears.

- b. Modify the daily alert subscriber list as follows.

- To add a subscriber, click the + icon, type the email address, and click **OK**.

CLI equivalent

```
# autosupport add alert-summary emails djones@company.com
```

- 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**.

CLI equivalent

```
# autosupport del alert-summary emails djones@company.com
```

- c. Click **Finish**.

Daily Alert Summary tab

The Daily Alert Summary tab allows you to configure an email list of those who want to receive a summary of all system alerts once each day. The people on this list do not receive individual alerts unless they are also added to a notification group.

Table 53 Daily Alert Summary, label descriptions

Item	Description
Delivery Time	The delivery time shows the configured time for daily emails.
Email List	This list displays the email addresses of those who receive the daily emails.

Table 54 Daily Alert Summary tab controls

Control	Description
Configure button	Click the Configure button to edit the subscriber email list.
Schedule button	Click the Schedule button to configure the time that the daily report is sent.

Enabling and disabling alert notification to Data Domain

You can enable or disable alert notification to Data Domain without affecting whether or not autosupport reports are sent to Data Domain.

Procedure

1. To view the alert reporting status, select **Maintenance > Support > Autosupport**.

The alert notification status is highlighted in green next to the Real-time alert label in the Support area. Depending on the current configuration, either an **Enable** or a **Disable** button appears in the Real-time alert row.

2. To enable alert reporting to Data Domain, click **Enable** in the Real-time alert row.
3. To disable alert reporting to Data Domain, click **Disable** in the Real-time alert row.

Testing the alerts email feature

Use the Notification tab to send a test email to select notification groups or email addresses. This feature allows you to determine if the system is configured correctly to send alert messages.

Procedure

1. To control whether or not a test alert is sent to Data Domain, do the following.
 - a. Select **Maintenance > Support > Autosupport**.
 - b. In the **Alert Support** area, click **Enable** or **Disable** to control whether or not the test email is sent .

You cannot change the email address.

2. Select **Health > Alerts > Notification**.
3. Select **More Tasks > Send Test Alert**.
The Send Test Alert dialog box appears.
4. In the **Notification Groups** list, select groups to receive the test email and click **Next**.
5. Optionally, add additional email addresses to receive the email.
6. Click **Send Now** and **OK**.

CLI equivalent

```
# alerts notify-list test jsmith@yourcompany.com
```

7. If you disabled sending of the test alert to Data Domain and you want to enable this feature now, do the following.
 - a. Select **Maintenance > Support > Autosupport**.
 - b. In the **Alert Support** area, click **Enable** .

Results

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`

Support delivery management

Delivery management defines how alerts and autosupport reports are sent to Data Domain. By default, alerts and autosupport reports are sent to Data Domain Customer Support using the standard (unsecure) email. The ConnectEMC method sends messages in a secure format through the Secure Remote Services Virtual Edition (VE) gateway.

When the ConnectEMC method is used with a Secure Remote Services gateway, one benefit is that one gateway can forward messages from multiple systems, and this allows you to configure network security for only the Secure Remote Services gateway instead of for multiple systems. Also, a usage intelligence report is generated and sent if electronic licenses are adopted.

When configuring a Secure Remote Services gateway, the Data Domain system supports registering multiple gateways to provide redundancy.

Selecting standard email delivery to Data Domain

When you select the standard (non-secure) email delivery method, this method applies to both alert and autosupport reporting.

Procedure

1. Select **Maintenance > Support > Autosupport**.
2. Click **Configure** in the Channel row in the Support area.
The Configure EMC Support Delivery dialog appears. The delivery method is displayed after the Channel label in the Support area.
3. In the **Channel** list box, select **Email to datadomain.com**.
4. Click **OK**.

CLI equivalent

```
# support notification method set email
```

Selecting and configuring Secure Remote Services delivery

Secure Remote Services Virtual Edition (VE) Gateway provides automated connect home and remote support activities through an IP-based solution that is enhanced by a comprehensive security system.

An on-premise Secure Remote Services version 3 gateway provides the ability to monitor both on-premise Data Domain systems and DD VE instances, and cloud-based DD VE instances.

Procedure

1. Select **Maintenance > Support > Autosupport**.
2. Click **Configure** in the Channel row in the Support area.

The Configure EMC Support Delivery dialog box appears. The delivery method is displayed after the Channel label in the Support area.

3. In the **Channel** list box, select **EMC Secure Remote Services**.
4. Type the gateway hostname and select the local IP address for the Data Domain system.
5. Click **OK**.
6. Type the service link username and password.
7. Click **Register**.

Secure Remote Services details are displayed in the Autosupport panel.

CLI equivalent

```
# support connectemc device register ipaddr esrs-gateway [host-list] [ha-peer ipaddr]
```

Note

The `ha-peer` parameter is required when configuring Secure Remote Services on Data Domain HA pairs to register both nodes.

Testing ConnectEMC operation

A CLI command allows you to test ConnectEMC operation by sending a test message to Support through the Secure Remote Services gateway.

Procedure

1. To test ConnectEMC operation, use the CLI.

```
#support connectemc test
Sending test message through ConnectEMC...
Test message successfully sent through ConnectEMC.
```

Log file management

The Data Domain system maintains a set of log files, which can be bundled and sent to Support to assist in troubleshooting any system issues that may arise. Log files cannot be modified or deleted by any user with DD System Manager, but they can be copied from the log directory and managed off of the system.

Note

Log messages on an HA system are preserved on the node where the log file originated.

Log files are rotated weekly. Every Sunday at 0:45 a.m., the system automatically opens new log files for the existing logs and renames the previous files with appended numbers. For example, after the first week of operation, the previous `week messages` file is renamed `messages . 1`, and new messages are stored in a new `messages` file. Each numbered file is rolled to the next number each week. For example, after the second week, the file `messages . 1` is rolled to `messages . 2`. If a `messages . 2` file already existed, it rolls to `messages . 3`. At the end of the retention period (shown in the table below, the expired log is deleted. For example, an existing `messages . 9` file is deleted when `messages . 8` rolls to `messages . 9`.

Except as noted in this topic, the log files are stored in `/ddvar/log`.

Note

Files in the `/ddvar` directory can be deleted using Linux commands if the Linux user is assigned *write* permission for that directory.

The set of log files on each system is determined by the features configured on the system and the events that occur. The following table describes the log files that the system can generate.

Table 55 System log files

Log File	Description	Retention Period
audit.log	Messages about user log-in events.	15 weeks
cifs.log	Log messages from the CIFS subsystem are logged only in <code>debug/cifs/cifs.log</code> . Size limit of 50 MiB.	10 weeks
messages	Messages about general system events, including commands executed.	9 weeks
secure.log	Messages regarding user events such as successful and failed logins, user additions and deletions, and password changes. Only Admin role users can view this file.	9 weeks
space.log	<p>Messages about disk space usage by system components, and messages from the clean process. A space use message is generated every hour. Each time the clean process runs, it creates approximately 100 messages. All messages are in comma-separated-value format with tags you can use to separate the disk space messages from the clean process messages. You can use third-party software to analyze either set of messages. The log file uses the following tags.</p> <ul style="list-style-type: none"> • CLEAN for data lines from clean operations. • CLEAN_HEADER for lines that contain headers for the clean operations data lines. • SPACE for disk space data lines. • SPACE_HEADER for lines that contain headers for the disk space data lines. 	A single file is kept permanently. There is no log file rotation for this log.

Viewing log files in DD System Manager

Use the Logs tab to view and open the system log files in DD System Manager.

Procedure

1. Select **Maintenance > Logs**.

The Logs list displays log file names and the size and generation date for each log file.

2. 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.

Displaying a log file in the CLI

Use the `log view` command to view a log file in the CLI.

Procedure

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

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.emc.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.

Learning more about log messages

Look up error messages in the Error Message Catalog for your DD OS version.

In the log file is text similar to the following.

```
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 as follows.

```
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. Go to the Online Support website at <https://support.emc.com>, enter *Error Message Catalog* in the search box, and click the search button.
2. In the results list, locate the catalog for your system and click on the link.

3. Use your browser search tool to search for a unique text string in the message.

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.
```

4. 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

Save log file copies to another device when you want to archive those files.

Use NFS, CIFS mount, or FTP to copy the files to another machine. If using CIFS or NFS, mount `/ddvar` to your desktop and copy the files from the mount point. The following procedure describes how to use FTP to move 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 includes the IP address of 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 as shown in the following 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`.

Log message transmission to remote systems

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).

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

Viewing the log file transmission configuration

Use the `log host show` CLI command to view whether log file transmission is enabled and which hosts receive log files.

Procedure

1. To display the configuration, enter the `log host show` command.

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

Enabling and disabling log message transmission

You must use CLI commands to enable or disable log message transmission.

Procedure

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

Adding or removing a receiver host

You must use CLI commands to add or remove a receiver host.

Procedure

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

The following command adds the system named *log-server* to the hosts that receive log messages.

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

The following command removes the system named *log-server* from the hosts that receive log messages.

```
log host del log-server
```

The following command disables the sending of logs and clears the list of destination hostnames..

```
log host reset
```

Remote system power management with IPMI

Select DD systems support remote power management using the Intelligent Platform Management Interface (IPMI), and they support remote monitoring of the boot sequence using Serial over LAN (SOL).

IPMI power management takes place between an IPMI initiator and an IPMI remote host. The IPMI initiator is the host that controls power on the remote host. To support remote power management from an initiator, the remote host must be configured with an IPMI username and password. The initiator must provide this username and password when attempting to manage power on a remote host.

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 an initiator 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.

You can use both DD System Manager and the CLI to configure IPMI users on a remote system. After you configure IPMI on a remote system, you can use IPMI initiator features on another system to log in and manage power.

Note

If a system cannot support IPMI due to hardware or software limitations, DD System Manager displays a notification message when attempting to navigate to a configuration page.

SOL is used to view the boot sequence after a power cycle on a remote system. SOL enables text console data that is normally sent to a serial port or to a directly attached console to be sent over a LAN and displayed by a management host.

The DD OS CLI allows you to configure a remote system for SOL and view the remote console output. This feature is supported only in the CLI.

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.

IPMI and SOL limitations

IPMI and SOL support is limited on some Data Domain systems.

- IPMI is supported on all systems supported by this release except the following systems: DD140, DD610, and DD630.
 - IPMI user support varies as follows.
 - Model DD990: Maximum user IDs = 15. Three default users (NULL, anonymous, root). Maximum user IDs available = 12.
 - Models DD640, DD4200, DD4500, DD7200, and DD9500: Maximum user IDs = 10. Two default users (NULL, root). Maximum user IDs available = 8.
 - SOL is supported on the following systems: DD160, DD620, DD640, DD670, DD860, DD890, DD990, DD2200, DD2500 (requires DD OS 5.4.0.6 or later), DD4200, DD4500, DD7200, and DD9500.
-

Note

User root is not supported for IPMI connections on DD160 systems.

Adding and deleting IPMI users with DD System Manager

Each system contains its own list of configured IPMI users, which is used to control access to local power management features. Another system operating as an IPMI initiator can manage remote system power only after providing a valid username and password.

To give an IPMI user the authority to manage power on multiple remote systems, you must add that user to each of the remote systems.

Note

The IPMI user list for each remote 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.

Procedure

1. Select **Maintenance > IPMI**.
2. To add a user, complete the following steps.
 - a. Above the IPMI Users table, click **Add**.
 - b. 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).

- c. Click **Create**.

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

3. To delete a user, complete the following steps.
 - a. In the IPMI Users list, select a user and click **Delete**.
 - b. In the Delete User dialog box, click **OK** to verify user deletion.

Changing an IPMI user password

Change the IPMI user password to prevent use of the old password for power management.

Procedure

1. Select **Maintenance > IPMI**.
2. In the IPMI Users table, select a user, and click **Change Password**.
3. In the Change Password dialog box, type the password in the appropriate text box and reenter the password in the **Verify Password** box.
4. Click **Update**.

Configuring an IPMI port

When you configure an IPMI port for a system, you select the port from a network ports list and specify the IP configuration parameters for that port. The selection of IPMI ports displayed is determined by the Data Domain system model.

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 > Ethernet > 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 rule described earlier. On these systems, IPMI port, `bmc0a`, corresponds to shared port `ethMa` in the network interface list. If possible, reserve the shared port `ethMa` 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, if possible, do not use the link aggregation feature on the shared interface because link state changes can interfere with IPMI connectivity.

Procedure

1. 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 who can access the managed system.

Table 56 Network Ports list column descriptions

Item	Description
Port	The logical name for a port that supports IPMI communications.
Enabled	Whether the port is enabled for IPMI (Yes or No).
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 57 IPMI Users list column descriptions

Item	Description
User Name	The name of a user with authority to power manage the remote system.

- In the **Network Ports** table, select a port to configure.

Note

If the IPMI port also supports IP traffic (for administrator access or backup traffic), the interface port must be enabled before you configure IPMI.

- Above the **Network Ports** table, click **Configure**.
The Configure Port dialog box appears.
- 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.
- Enable a disabled IPMI network port by selecting the network port in the **Network Ports** table, and clicking **Enable**.
- Disable a disabled IPMI network port by selecting the network port in the **Network Ports** table, and clicking **Disable**.
- Click **Apply**.

Preparing for remote power management and console monitoring with the CLI

Remote console monitoring uses the Serial Over Lan (SOL) feature to enable viewing of text-based console output without a serial server. You must use the CLI to set up a system for remote power management and console monitoring.

Remote console monitoring is typically used in combination with the `ipmi remote power cycle` command to view the remote system’s boot sequence. This procedure

should be used on every system for which you might want to remotely view the console during the boot sequence.

Procedure

1. Connect the console to the system directly or remotely.
 - Use the following connectors for a direct connection.
 - DIN-type connectors for a PS/2 keyboard
 - USB-A receptacle port for a USB keyboard
 - DB15 female connector for a VGA monitor

Note

Systems DD4200, DD4500, and DD7200 do not support direct connection, including KVM.

-
- For a serial connection, use a standard DB9 male or micro-DB9 female connector. Systems DD4200, DD4500, and DD7200 provide a female micro-DB9 connector. A null modem cable with male micro-DB9 and standard female DB9 connectors is included for a typical laptop connection.
 - For a remote IPMI/SOL connection, use the appropriate RJ45 receptacle as follows.
 - For DD990 systems, use default port eth0d.
 - For other systems, use the maintenance or service port. For port locations, refer to the system documentation, such as a hardware overview or installation and setup guide.
2. To support remote console monitoring, use the default BIOS settings.
 3. To display the IPMI port name, enter `ipmi show config`.
 4. To enable IPMI, enter `ipmi enable {port | all}`.
 5. To configure the IPMI port, enter `ipmi config port { dhcp | ipaddress ipaddr netmask mask gateway ipaddr }`.

Note

If the IPMI port also supports IP traffic (for administrator access or backup traffic), the interface port must be enabled with the `net enable` command before you configure IPMI.

-
6. If this is the first time using IPMI, run `ipmi user reset` to clear IPMI users that may be out of synch between two ports, and to disable default users.
 7. To add a new IPMI user, enter `ipmi user add user`.
 8. To set up SOL, do the following:
 - a. Enter `system option set console lan`.
 - b. When prompted, enter `y` to reboot the system.

Managing power with DD System Manager

After IPMI is properly set up on a remote system, you can use DD System Manager as an IPMI initiator to log into the remote system, view the power status, and change the power status.

Procedure

1. Select **Maintenance > IPMI**.
2. Click **Login to Remote System**.

The IPMI Power Management dialog box appears.

3. Enter the remote system IPMI IP address or hostname and the IPMI username and password, then click **Connect**.
4. View the IPMI status.

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).

5. To change the IPMI power status, click the appropriate button.
 - **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.

Managing power with the CLI

You can manage power on a remote system and start remote console monitoring using the CLI.

Note

The remote system must be properly set up before you can manage power or monitor the system.

Procedure

1. Establish a CLI session on the system from which you want to monitor a remote system.
 2. To manage power on the remote system, enter `ipmi remote power {on | off | cycle | status} ipmi-target <ipaddr | hostname> user user`.
 3. To begin remote console monitoring, enter `ipmi remote console ipmi-target <ipaddr | hostname> user user`.
-

Note

The user name is an IPMI user name defined for IPMI on the remote system. DD OS user names are not automatically supported by IPMI.

4. To disconnect from a remote console monitoring session and return to the command line, enter the at symbol (@).
5. To terminate remote console monitoring, enter the tilde symbol (~).

CHAPTER 4

Monitoring Data Domain Systems

This chapter includes:

- [Viewing individual system status and identity information](#)..... 154
- [Health Alerts panel](#).....157
- [Viewing and clearing current alerts](#)..... 157
- [Viewing the alerts history](#)..... 158
- [Viewing hardware component status](#)..... 159
- [Viewing system statistics](#).....163
- [Viewing active users](#)..... 164
- [History report management](#)..... 164
- [Viewing the Task Log](#).....169
- [Viewing the system High Availability status](#)..... 170

Viewing individual system status and identity information

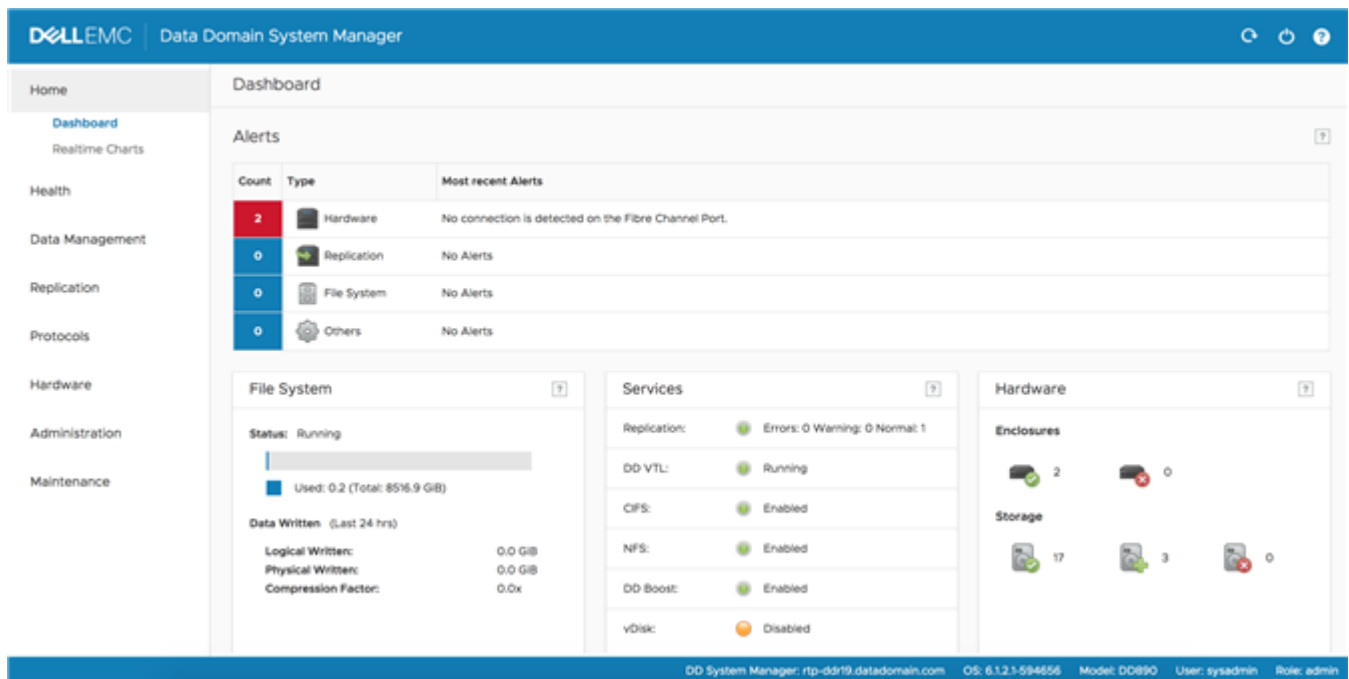
The **Dashboard** area displays summary information and status for alerts, the file system, licensed services, and hardware enclosures. The **Maintenance** area displays additional system information, including the system uptime and system and chassis serial numbers.

The system name, software version, and user information appear in the footer at all times.

Procedure

1. To view system dashboard, select **Home > Dashboard**.

Figure 4 System dashboard



2. To view the system uptime and identity information, select **Maintenance > System**.

The system uptime and identification information appears in the System area.

Dashboard Alerts area

The Dashboard Alerts area 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 anywhere in the alerts area to display more information on the current alerts.

Table 58 Dashboard Alerts column descriptions

Column	Description
Count	A count of the current alerts for the subsystem type specified in the adjacent

Table 58 Dashboard Alerts column descriptions (continued)

Column	Description
	column. The background color indicates the severity of the alert.
Type	The subsystem that generated the alert.
Most recent alerts	The text of the most recent alert for the subsystem type specified in the adjacent column

Dashboard File System area

The Dashboard File System area displays statistics for the entire file system. Click anywhere in the File System area to display more information.

Table 59 File System area label descriptions

Column	Description
Status	The current status of the file system.
X.Xx	The average compression reduction factor for the file system.
Used	The total file system space being used.
Data Written: Pre-compression	The data quantity received by the system prior to compression.
Data Written: Post-compression	The data quantity stored on the system after compression.

Dashboard Services area

The Dashboard Services area displays the status of the replication, DD VTL, CIFS, NFS, DD Boost, and vDisk services. Click on a service to display detailed information about that service.

Table 60 Services area column descriptions

Column	Description
Left column	The left column lists the services that may be used on the system. These service can include replication, DD VTL, CIFS, NFS, DD Boost, vDisk.
Right column	The right column shows the operational status of the service. For most services, the status is enabled, disabled, or not licensed. The replication service row displays the number of replication contexts that are in normal, warning, and error states. A color coded box displays green for normal operation, yellow for

Table 60 Services area column descriptions (continued)

Column	Description
	warning situations, or red when errors are present).

Dashboard HA Readiness area

In high-availability (HA) systems, the HA panel indicates whether the system can fail over from the active node to the standby node if necessary.

You can click on the **HA panel** to navigate to the **High Availability** section under **HEALTH**.

Dashboard Hardware area

The Dashboard Hardware area displays the status of the system enclosures and drives. Click anywhere in the Hardware area to display more information on these components.

Table 61 Hardware area label descriptions

Label	Description
Enclosures	The enclosure icons display the number of enclosures operating in the normal (green checkmark) and degraded (red X) states.
Storage	The storage icons display the number of disk drives operating in the normal (green checkmark), spare (green +), or failed (red X) state.

Maintenance System area

The Maintenance System area displays the system model number, DD OS version, system uptime, and system and chassis serial numbers.

Table 62 System area label descriptions

Label	Description
Model Number	The model number is the number assigned to the Data Domain system.
Version	The version is the DD OS version and build number of the software running on the system.
System Uptime	The system uptime displays how long the system has been running since the last system start. The time in parenthesis indicates when the system uptime was last updated.
System Serial No.	The system serial number is the serial number assigned to the system. On newer systems,

Table 62 System area label descriptions (continued)

Label	Description
	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.
Chassis Serial No.	The chassis serial number is the serial number on the current system chassis.

Health Alerts panel

Alerts are messages from system services and subsystems that report system events. The Health > Alerts panel displays tabs that allow you to view current and non-current alerts, the configured alert notification groups, and the configuration for those who want to receive daily alert summary reports.

Alerts are also sent as SNMP traps. See the *MIB Quick Reference Guide* or the SNMP MIB for the full list of traps.

Viewing and clearing current alerts

The Current Alerts tab displays a list of all the current alerts and can display detailed information for a selected alert. An alert is automatically removed from the Current Alerts list when the underlying situation is corrected or when manually cleared.

Procedure

1. To view all of the current alerts, select **Health > Alerts > Current Alerts**.
2. To limit the number of entries in the current alert list, do the following.
 - a. In the Filter By area, select a **Severity** and **Class** to expose only alerts that pertain to those choices.
 - b. Click **Update**.

All alerts not matching the Severity and Class are removed from the list.
3. To display additional information for a specific alert in the **Details** area, click the alert in the list.
4. To clear an alert, select the alert checkbox in the list and click **Clear**.

A cleared alert no longer appears in the current alerts list, but it can be found in the alerts history list.
5. To remove filtering and return to the full listing of current alerts, click **Reset**.

Current Alerts tab

The Current Alerts tab displays a list of alerts and detailed information about a selected alert.

Table 63 Alerts list, column label descriptions

Item	Description
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.

Table 64 Details area, row label descriptions

Item	Description
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.
Tenant Units	Lists affected tenant units.
Description	More descriptive information about the alert.
Action	A suggestion to remedy the alert.
Object Info	Additional information about the affected object.
SNMP OID	SNMP object ID.

Viewing the alerts history

The Alerts History tab displays a list of all the cleared alerts and can display detailed information for a selected alert.

Procedure

1. To view all of the alerts history, select **Health > Alerts > Alerts History**.
2. To limit the number of entries in the current alert list, do the following.
 - a. In the Filter By area, select a **Severity** and **Class** to expose only alerts that pertain to those choices.
 - b. Click **Update**.

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

3. To display additional information for a specific alert in the **Details** area, click the alert in the list.
4. To remove filtering and return to the full listing of cleared alerts, click **Reset**.

Alerts History tab

The Alerts History tab displays a list of cleared alerts and details about a selected alert.

Table 65 Alerts list, column label descriptions

Item	Description
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	Whether the status is posted or cleared. A posted alert is not cleared.

Table 66 Details area, row label descriptions

Item	Description
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.
Tenant Units	Lists affected tenant units.
Additional Information	More descriptive information about the alert.
Status	Whether the status is posted or cleared. A posted alert is not cleared.
Description	More descriptive information about the alert.
Action	A suggestion to remedy the alert.

Viewing hardware component status

The Hardware Chassis panel displays a block drawing of each enclosure in a system, including the chassis serial number and the enclosure status. 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.

On systems running DD OS 5.5.1 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.

Procedure

1. Select **Hardware > Chassis**.

The Chassis view shows the system enclosures. Enclosure 1 is the system controller, and the rest of the enclosures appear below Enclosure 1.

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

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

Fan status

Fans are numbered and correspond to their location in the chassis. Hover over a system fan to display a tooltip for that device.

Table 67 Fan tooltip, column label descriptions

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 status

Data Domain systems and some components are configured to operate within a specific temperature range, which is defined by a temperature profile that is not configurable. Hover over the Temperature box to display the temperature tooltip.

Table 68 Temperature tooltip, column label descriptions

Item	Description
Description	The location within the chassis being measured. The components listed depend on the model and are often shown as abbreviations. Some examples are: <ul style="list-style-type: none"> • CPU 0 Temp (Central Processing Unit) • MLB Temp 1 (main logic board) • BP middle temp (backplane) • LP temp (low profile of I/O riser FRU) • FHFL temp (full height full length of I/O riser FRU)

Table 68 Temperature tooltip, column label descriptions (continued)

Item	Description
	<ul style="list-style-type: none"> FP temp (front panel)
C/F	The C/F column displays temperature in degrees Celsius and Fahrenheit. When the description for a CPU specifies <i>relative</i> (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"> OK—The temperature is acceptable Critical—The temperature is higher than the shutdown temperature. Warning—The temperature is higher than the warning temperature (but lower than the shutdown temperature). Dash (-) —No temperature thresholds are configured for this component, so there is no status to report.

Management panel status

DD6300, DD6800, and DD9300 systems have a fixed management panel with an Ethernet port for the management network on the rear of the chassis. Hover over the Ethernet port to display a tooltip.

Table 69 Management panel tooltip, column label descriptions

Item	Description
Description	The type of NIC installed in the management panel.
Vendor	The manufacturer of the management NIC.
Ports	The name of the management network (Ma).

SSD status (DD6300 only)

The DD6300 supports up to two SSDs in slots on the rear of the chassis. The SSD slots are numbered and correspond to their location in the chassis. Hover over an SSD to display a tooltip for that device.

Table 70 SSD tooltip, column label descriptions

Item	Description
Description	The name of the SSD.
Status	The state of the SSD.
Life Used	The percentage of the rated operating life the SSD has used.

Power supply status

The tooltip 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 slot status

The PCI slots shown in the chassis view indicate the number of PCI slots and the numbers of each slot. Tooltips provide component status for each card in a PCI slot. For example, the tooltip for one NVRAM card model displays the memory size, temperature data, and battery levels.

NVRAM status

Hover over NVRAM to display information about the Non-Volatile RAM, batteries, and other components.

Table 71 NVRAM tooltip, dolumn label descriptions

Item	Description
Component	<p>The items in the component list depend on the NVRAM installed in the system and can include the following items.</p> <ul style="list-style-type: none"> • Firmware version • Memory size • Error counts • Flash controller error counts • Board temperature • CPU temperature • Battery number (The number of batteries depends on the system type.) • Current slot number for NVRAM
C/F	<p>Displays the temperature for select components in the Celsius/ Fahrenheit format.</p>
Value	<p>Values are provided for select components and describe the following.</p> <ul style="list-style-type: none"> • Firmware version number • Memory size value in the displayed units • Error counts for memory, PCI, and controller • Flash controller error counts sorted in the following groups: configuration errors (Cfg Err), panic conditions (Panic), Bus Hang, bad block warnings (Bad Blk Warn), backup errors (Bkup Err), and restore errors (Rstr Err) • Battery information, such percent charged and status (enabled or disabled)

Viewing system statistics

The Realtime Charts panel displays up to seven graphs that show real-time subsystem performance statistics, such as CPU usage and disk traffic.

Procedure

1. Select **Home > Realtime Charts**.

The Performance Graphs area displays the currently selected graphs.

2. To change the selection of graphs to display, select and clear the checkboxes for graphs in the list box.
3. To view specific data-point information, hover over a graph point.
4. 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.

Performance statistics graphs

The performance statistics graphs display statistics for key system components and features.

DD Boost Active Connections

The DD Boost Active Connections graph displays the number of active DD Boost connections for each of the past 200 seconds. Separate lines within the graph display counts for Read (recovery) connections and Write (backup) connections.

DD Boost Data Throughput

The DD Boost Data Throughput graph displays the bytes/second transferred for each of the past 200 seconds. Separate lines within the graph display the rates for data read from the system by DD Boost clients and data written to the system by DD Boost clients.

Disk

The Disk graph displays 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 system.

File System Operations

The File System Operations graph displays the number of operations per second that occurred for each of the past 200 seconds. Separate lines within the graph display the NFS and CIFS operations per second.

Network

The Network graph displays 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.

Recent CPU Usage

The Recent CPU Usage graph displays the percentage of CPU usage for each of the past 200 seconds.

Replication (DD Replicator must be licensed)

The Replication graph displays the amount of replication data traveling over the network for each of the last 200 seconds. Separate lines display the In and Out data as follows:

- **In:** The total number of units of measurement, such as kilobytes per second, received by this side from the other side of the DD 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 DD 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.

Viewing active users

The Active Users tab displays the names of users who are logged into the system and statistics about the current user sessions.

Procedure

1. Select **Administration > Access > Active Users**.

The Active Users list appears and displays information for each user.

Table 72 Active Users list, column label descriptions

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 login. GUI appears for DD System Manager users.

Note

To manage local users, click **Go to Local Users**.

History report management

DD System Manager enables you to 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, and view daily and cumulative reports on the file system.

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 deduplication 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 New Report area lists the types of reports you can generate on your system.

Note

Replication reports can only be created if the system contains a replication license and a valid replication context is configured.

File System Cumulative Space Usage report

The File System Cumulative Space Usage Report displays 3 charts that detail space usage 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.

Table 73 File System—Usage chart label descriptions

Item	Description
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 was written. The time displayed on this report changes based upon the Duration selection when the chart was created.
Total Compression Factor	The total compression factor reports the compression ratio.

Table 74 File System—Consumption chart label descriptions

Item	Description
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.
Used (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.
Capacity	The total capacity on a Data Domain system.
Cleaning	Cleaning is the Cleaning cycle (start and end time for each cleaning cycle). Administrators can use this information to choose the best time for space cleaning the best throttle setting.

Table 75 File System Weekly Cumulative Capacity chart label descriptions

Item	Description
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.
Compression Factor	The total compression factor. This is indicated by a black line on the report.

File System Daily Space Usage report

The File System Daily Space Usage report displays five charts that detail space usage during the specified duration. This report is used to analyze daily activities.

Table 76 File System Daily Space Usage chart label descriptions

Item	Description
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.

Table 77 File System Daily Capacity Utilization chart label descriptions

Item	Description
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.

Table 78 File System Weekly Capacity Utilization chart label descriptions

Item	Description
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.

Table 79 File System Compression Summary chart label descriptions

Item	Description
Time	The period of data collection for this report.
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.

Table 80 File System Cleaning Activity chart label descriptions

Item	Description
Start Time	The time the cleaning activity started.
End Time	The time the cleaning activity finished.
Duration (Hours)	The total time required for cleaning in hours.
Space Reclaimed	The space reclaimed in Gibibytes (GiB).

Replication Status report

The Replication Status report displays three charts that provide 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.

Table 81 Replication Context Summary chart label descriptions

Item	Description
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.

Table 82 Replication Context Error Status chart label descriptions

Item	Description
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.

Table 83 Replication Destination Space Availability chart label descriptions

Item	Description
Destination	Destination system name.
Space Availability (GiB)	Total amount of storage available.

Replication Summary report

The Replication Summary report provides 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.

Table 84 Replication Summary report label descriptions

Item	Description
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.

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.

Procedure

1. Select **Health > Jobs**.

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, do either of the following.
 - Click **Update** to update the task log.
 - Click **Reset** to display all tasks and remove any filters that were set.
4. To display detailed information about a task, select the task in the task list.

Table 85 Detailed Information, label descriptions

Item	Description
System	The system name.
Task Description	A description of the task.
Status	The status of the task (completed, failed, or in progress).
Start Time	The date and time the task started.

Table 85 Detailed Information, label descriptions (continued)

Item	Description
End Time	The date and time the task ended.
Error Message	An applicable error message, if any.

Viewing the system High Availability status

You can use the **High Availability** panel to see detailed information about the HA status of the system and whether the system can perform failover if necessary.

Procedure

1. Select **Health > High Availability** on the DD System Manager.

The **Health High Availability** screen appears.

A green check mark indicates the system is operating normally and ready for failover.

The screen shows the active node, which is typically Node 0.

2. Hover the cursor over a node to see its status.

The node is highlighted in blue if it is active.

3. Click the drop-down menu in the banner if you want to change the view from the active node to the standby node, which is typically Node 1.

High Availability status

The **Health High Availability (HA)** view informs you about the system status using a diagram of the nodes and their connected storage. In addition, you can also see any current alerts as well as detailed information about the system.

You can determine if the active node and the storage are operational by hovering the cursor over them. Each is highlighted in blue when operating normally. The standby node should appear gray.

You can also filter the alerts table by clicking on a component. Only alerts related to the selected components will be displayed.

Figure 5 Health/High Availability indicators




Table 86 High Availability indicators

Item	Description
HA System bar	Displays a green check mark when the system is operating normally and ready for failover.
Failover to Node 0	Allows you to manually fail over to the standby node.
Take Node 1 Offline	Allows you to take the active node offline if necessary.
System Information	Lists the Data Domain system model, the system type, the version of the Data Domain operating system version in use, and the applied HA license.
HA Manager	Displays the nodes, their attached storage, the HA interconnect, and the cabling.
Severity	Indicates the severity of any alerts that could impact the system's HA status.
Component	Indicates which component is affected.
Class	Indicates the class of the alert received such as hardware, environment, and others.

Table 86 High Availability indicators (continued)

Item	Description
Post Time	Indicates the time and date the alert was posted.

CHAPTER 5

File System

This chapter includes:

- [File system overview](#) 174
- [Monitoring file system usage](#) 181
- [Managing file system operations](#) 189
- [Fast copy operations](#) 196

File system overview

Learn how to use the file system.

How the file system stores data

Data Domain storage capacity is best managed by keeping multiple backups and 20% empty space to accommodate backups until the next cleaning. Space use is primarily affected by the size and compressibility of data, and the retention period.

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, 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 reports 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 the section regarding types of compression.) 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 `fileysys 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

Data Domain 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. Two other types of local compression are available, `gzfast` and `gz`. Both provide increased compression over `lz`, but at the cost of additional CPU load. Local compression options provide a trade-off between slower performance and space usage. It is also possible to turn off local compression. To change compression, see the section regarding changing local compression.

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

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 the section regarding modifying a cleaning schedule).

Data Domain recommends running the cleaning operation once a week.

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 collection replication, data in a replication context on the source system that has not been replicated cannot be processed for file system cleaning. If file system cleaning is not able to complete because the source and destination systems are out of synch, the system reports the status of the cleaning operation as `partial`, and only limited system statistics are available for the cleaning operation. If collection replication is disabled, the amount of data that cannot be processed for file system cleaning increases because the replication source and destination systems remain out of synch. The KB article *Data Domain: An overview of Data Domain File System (DDFS) clean/garbage collection (GC) phases*, available from the Online Support site at <https://support.emc.com> provides additional information.

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

Interfaces supported by the file system.

- NFS
- CIFS
- DD Boost
- DD VTL

Supported backup software

Guidance for setting up backup software and backup servers to use with a Data Domain systems is available at support.emc.com.

Data streams sent to a Data Domain system

For optimal performance, Data Domain recommends limits on simultaneous streams between Data Domain systems and your backup servers.

A data stream, in the context of the following table, 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 Replication source or destination stream refers to a directory replication operation or a DD Boost file replication stream associated with a file replication operation.

Table 87 Data streams sent to a Data Domain system

Model	RAM/ NVRAM	Backup write streams	Backup read streams	Repl ^a source streams	Repl ^a dest streams	Mixed
DD140, DD160, 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
DD620, DD630, DD640	8 GB / 0.5 GB or 1 GB	20	16	20	20	w<=20; r<=16; ReplSrc<=30; ReplDest<=20; ReplDest+w<=20; Total<=30
DD640, DD670	16 GB or 20 GB / 1 GB	90	30	60	90	w<=90; r<=30; ReplSrc<=60; 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<=90
DD860	72 GB ^b / 1 GB	90	50	90	90	w<=90; r<=50; ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; Total<=90
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 ^b / 4 GB	540	150	270	540	w<=540; r<=150; ReplSrc<=270; ReplDest<=540; ReplDest +w<=540; Total<=540
DD2200	8 GB	35	6	18	20	w<=35; r<=6; ReplSrc<=18; ReplDest<=20; ReplDest+w<=35; Total<=35
DD2200	16 GB	60	16	30	60	w<=60; r<=16; ReplSrc<=30; ReplDest<=60; ReplDest+w<=60; Total<=60
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 ^b / 4 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest +w<=270; Total<=270
DD4500	192 GB ^b / 4 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest +w<=270; Total<=270
DD7200	128 or 256 GB ^b / 4 GB	540	150	270	540	w<=540; r<=150; ReplSrc<=270; ReplDest<=540; ReplDest +w<=540; Total<=540
DD9500	256/512 GB	1885	300	540	1080	w<=1885; r<=300; ReplSrc<=540; ReplDest<=1080; ReplDest+w<=1080; Total<=1885

Table 87 Data streams sent to a Data Domain system (continued)

Model	RAM/ NVRAM	Backup write streams	Backup read streams	Repl ^a source streams	Repl ^a dest streams	Mixed
DD9800	256/768 GB	1885	300	540	1080	w<=1885; r<=300; ReplSrc<=540; ReplDest<=1080; ReplDest+w<=1080; Total<=1885
DD6300	48/96 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest +w<=270; Total<=270
DD6800	192 GB	400	110	220	400	w<=400; r<=110; ReplSrc<=220; ReplDest<=400; ReplDest +w<=400; Total<=400
DD9300	192/384 GB	800	220	440	800	w<=800; r<=220; ReplSrc<=440; ReplDest<=800; ReplDest +w<=800; Total<=800
DD VE 8 TB	8 GB / 512 MB	20	16	20	20	w<= 20 ; r<= 16 ReplSrc<=20; ReplDest<=20; ReplDest+w<=20; w+r+ReplSrc <=20;Total<=20
DD VE 16 TB	16 GB / 512 MB or 24 GB / 1 GB	45	30	45	45	w<= 45 ; r<= 30 ReplSrc<=45; ReplDest<=45; ReplDest+w<=45; w+r+ReplSrc <=45;Total<=45
DD VE 32 TB	24 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 48 TB	36 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 64 TB	48 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 96 TB	64 GB / 2 GB	180	50	90	180	w<= 180 ; r<= 50 ReplSrc<=90; ReplDest<=180; ReplDest +w<=180; w+r+ReplSrc <=180;Total<=180
DD3300 4 TB	12 GB (virtual memory) / 512 MB	20	16	30	20	w<= 20 ; r<= 16 ReplSrc<=30; ReplDest<=20; ReplDest+w<=20; w+r+ReplSrc <=30;Total<=30
DD3300 8 TB	32 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD3300 16 TB	32 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD3300 32 TB	46 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=140

Table 87 Data streams sent to a Data Domain system (continued)

- a. DirRepl, OptDup, MTreeRepl streams
- b. The Data Domain Extended Retention software option is available only for these devices with extended (maximum) memory

File system limitations

File system limitations, including: limits on the number of files, the battery, and so on.

Limits on number of files in a Data Domain system

Consequences and considerations of storing more than 1 billion files.

Data Domain recommends storing no more than 1 billion files on a system. Storing a larger number of files can adversely affect the performance and the length of cleaning, and 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 may take from a few minutes to several hours depending upon the number of files in the system.

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.

When the file system passes the billion file limit, several processes or operations might be adversely affected, for example:

- Cleaning may take a very long time to complete, perhaps several days.
- AutoSupport operations may take more time.
- Any process or command that needs to enumerate all the files.

If there are many small files, other considerations arise:

- 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 100 to 200 files per second depending upon the number of MTrees and CIFS connections. 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 250,000. Larger directory sizes might experience slower responses to metadata operations such as listing the files in the directory and opening or creating a file.

Limits on the battery

For systems that use NVRAM, the operating system creates a low battery alert if the battery charge falls below 80% capacity, and the file system is disabled.

NOTICE

The Data Domain DD2200 system does not use NVRAM so firmware calculations decide whether the battery charge is sufficient to save the data and disable the file system if there is a loss of AC power.

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.

Limited access during HA failover

Access to files may be interrupted for up to 10 minutes during failover on High Availability systems. (DD Boost and NFS require additional time.)

Monitoring file system usage

View real-time data storage statistics.

The File System view has tabs and controls that provide access to real-time data storage statistics, cloud unit information, encryption information, and graphs of space usage amounts, consumption factors, and data written trends. There are also controls for managing file system cleaning, expansion, copying, and destruction.

Accessing the file system view

This section describes the file system functionality.

Procedure

- Select **Data Management > File System**.

About the File System Status panel

Display the status of file system services.

To access the File System Status panel, click **Data Management > File System > Show Status of File System Services**.

File System

The **File System** field contains an **Enable/Disable** link 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.
- Destroyed—if the file system is deleted.
- Error—if there is an error condition, such as a problem initializing the file system.

Cloud File Recall

The **Cloud File Recall** field contains a **Recall** link to initiate a file recall from the Cloud Tier. A **Details** link is available if any active recalls are underway. For more information, see the "Recalling a File from the Cloud Tier" topic.

Physical Capacity Measurement

The **Physical Capacity Measurement** field contains an **Enable** button when physical capacity measurement status is disabled. When enabled, the system displays **Disable** and **View** buttons. Click **View** to see currently running physical capacity measurements: MTree, priority, submit time, start time, and duration.

Data Movement

The **Data Movement** field contains **Start/Stop** buttons and shows the date the last data movement operation finished, the number of files copied, and the amount of data copied. The system displays a **Start** button when the data movement operation is available, and a **Stop** when a data movement operation is running.

Active Tier Cleaning

The **Active Tier Cleaning** field contains a **Start/Stop** 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
```

Cloud Tier Cleaning

The **Cloud Tier Cleaning** field contains a **Start/Stop** 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 tab

Click the Summary tab to show space usage statistics for the active and cloud tiers and to access controls for viewing file system status, configuring file system settings, performing a Fast Copy operation, expand capacity, and destroy the file system.

For each tier, space usage statistics include:

- **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 modifying a cleaning schedule procedure 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.
- **Pre-Compression** (GiB)—Data written before compression.
- **Total Compression Factor (Reduction %)**—Pre-Comp / Post-Comp.
- **Cleanable** (GiB)—The amount of space that could be reclaimed if a cleaning were run.

For Cloud Tier, the **Cloud File Recall** field contains a **Recall** link to initiate a file recall from the Cloud Tier. A **Details** link is available if any active recalls are underway. For more information, see the "Recalling a File from the Cloud Tier" topic.

Separate panels provide the following statistics for the last 24 hours for each tier:

- **Pre-Compression** (GiB)—Data written before compression.
- **Post-Compression** (GiB)—Storage used after compression.
- **Global Compression Factor**—(Pre-Compression / (Size after global compression)).
- **Local Compression Factor**—(Size after global compression) / Post-Compression).
- **Total Compression Factor (Reduction %)**—[(Pre-Comp - Post-Comp) / Pre-Comp] * 100.

About file system settings

Display and change system options as well as the current cleaning schedule.

To access the File System Settings dialog, click **Data Management > File System > Settings**.

Table 88 General settings

General settings	Description
Local Compression Type	The type of local compression in use. <ul style="list-style-type: none"> • See the section regarding types of compression for an overview. • See the section regarding changing local compression
Cloud Tier Local Comp	The type of compression in use for the cloud tier. <ul style="list-style-type: none"> • See the section regarding types of compression for an overview. • See the section regarding changing local compression
Report Replica as Writable	How applications see a replica. <ul style="list-style-type: none"> • See the section regarding changing read-only settings
Staging Reserve	Manage disk staging. <ul style="list-style-type: none"> • See the section regarding working with disk staging • See the section regarding configuring disk staging

Table 88 General settings (continued)

General settings	Description
Marker Type	Backup software markers (tape markers, tag headers, or other names are used) in data streams. See the section regarding tape marker settings
Throttle	See the section regarding setting the physical capacity measurement throttle.
Cache	Physical Capacity Cache initialization cleans up the caches and enhances the measuring speed.

You can adjust the workload balance of the file system to increase performance based on your usage.

Table 89 Workload Balance settings

Workload Balance settings	Description
Random workloads (%)	Instant access and restores perform better using random workloads.
Sequential workloads (%)	Traditional backups and restores perform better with sequential workloads.

Table 90 Data Movement settings

Data movement policy settings	Description
File Age Threshold	When data movement starts, all files that have not been modified for the specified threshold number of days will be moved from the active to the retention tier.
Schedule	Days and times data is moved.
Throttle	The percentage of available resources the system uses for data movement. A throttle value of 100% is the default throttle and means that data movement will not be throttled.

Table 91 Cleaning settings

Cleaning schedule settings	Description
Time	The date time cleaning operations run. <ul style="list-style-type: none"> See the section regarding modifying a cleaning schedule
Throttle	The system resources allocation. <ul style="list-style-type: none"> See the section regarding throttling the cleaning operation

About the Cloud Units tab

Display summary information for cloud units, add and modify cloud units, and manage certificates.

The Cloud Units tab on the File System page is shown only when the optional DD Cloud Tier license is enabled. This view lists summary information (status, network bandwidth, read access, local compression, data movement and data status) the name of the cloud provider, the used capacity, and the licensed capacity. Controls are provided for editing the cloud unit, managing certificates, and adding a new cloud unit.

About the Retention Units tab

Display the retention unit and its state, status, and size.

The Retention Units tab on the File System page is shown only when the optional DD Extended Retention license is enabled. This view lists the retention unit and shows its 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.

Select the diamond symbol to the right of a column heading to sort the order of the values in reverse.

About the DD Encryption tab

Display encryption status, progress, algorithms, and so on.

Table 92 DD Encryption settings

Setting	Description
DD System	Status can be one of the following: <ul style="list-style-type: none"> • Not licensed—No other information provided. • Not configured—Encryption is licensed but not configured. • Enabled—Encryption is enabled and running. • Disabled—Encryption is disabled.
Active Tier	View encryption status for the active tier: <ul style="list-style-type: none"> • Enabled—Encryption is enabled and running. • Disabled—Encryption is disabled.
Cloud Unit	View encryption status per cloud unit: <ul style="list-style-type: none"> • Enabled—Encryption is enabled and running. • Disabled—Encryption is disabled.
Encryption Progress	View encryption status details for the active tier regarding the application of changes and re-encryption of data. Status can be one of the following: <ul style="list-style-type: none"> • None • Pending • Running • Done

Table 92 DD Encryption settings (continued)

Setting	Description
	<p>Click View Details to display the Encryption Status Details dialog that includes the following information for the Active Tier:</p> <ul style="list-style-type: none"> Type (Example: Apply Changes when encryption has already been initiated, or Re-encryption when encryption is a result of compromised data—perhaps a previously destroyed key.) Status (Example: Pending) Details: (Example: Requested on December xx/xx/xx and will take after the next system clean).
Encryption Algorithm	<p>The algorithm used to encrypt the data:</p> <ul style="list-style-type: none"> AES 256-bit (CBC) (default) AES 256-bit (GCM) (more secure but slower) AES 128-bit (CBC) (not as secure as 256-bit) AES 128-bit (GCM) (not as secure as 256-bit) <p>See Changing the Encryption Algorithm for details.</p>
Encryption Passphrase	When configured, shows as “*****.” To change the passphrase, see Managing the System Passphrase.
File System Lock	
Status	<p>The File System Lock status is either:</p> <ul style="list-style-type: none"> Unlocked—The feature is not enabled. Locked—The feature is enabled.
Key Management	
Key Manager	Either the internal Data Domain Embedded Key Manager or the optional RSA Data Protection Manager (DPM) Key Manager. Click Configure to switch between key managers (if both are configured), or to modify Key Manager options.
Server	The name of the RSA Key Manager Server.
Server Status	Online or offline, or the error messages returned by the RSA Key Manager Server.
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 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> <hr/> <p>Note</p> <p>The Data Domain system supports only key classes configured to return the current key.</p> <hr/>
Port	The port number of the RSA server.

Table 92 DD Encryption settings (continued)

Setting	Description
FIPS mode	Whether or not the imported host certificate is FIPS compliant. The default mode is enabled.
Encryption Keys	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), its state (see Working with the RSA DPM Key Manager, DPM Encryption Key States Supported by Data Domain), and the amount of the data encrypted with the key. The system displays the last updated time for key information above the right column. Selected keys in the list can be: <ul style="list-style-type: none"> • Synchronized so the list shows new keys added to the RSA server (but are not usable until the file system is restarted). • Deleted. • Destroyed.

About the space usage view (file system)

Display a visual (but static) representation of data use for the file system at certain points in time.

Click **Data Management > File System > Charts**. Select **Space Usage** from the Chart drop-down list.

Click a point on a graph line to display data at that point. The lines of the graph denote measurements 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 a Date Range (that is, 1w, 1m, 3m, 1y, or All) above the graph lets you change the number of days of data shown on the graph, from one week to all.

About the consumption view

Display space used over time, in relation to total system capacity.

Click **Data Management > File System > Charts**. Select **Consumption** from the Chart drop-down list.

Click a point on a graph line to display data at that point. The lines of the graph denote measurements 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 a Date Range (that is, 1w, 1m, 3m,1y, or All) above the graph lets you change the number of days of data shown on the graph, from one week to all.

About the daily written view (file system)

Display the flow of data over time. The data amounts are shown over time for pre- and post-compression amounts.

Click **Data Management > File System > Charts**. Select **Daily Written** from the Chart drop-down list.

Click a point on a graph line to display a box with data at that point. The lines on the graph denote measurements for:

- **Pre-Comp Written**—The total amount of data written to the file system by backup servers. Pre-compressed data on the file system is what a backup server sees as the total uncompressed data held by the file system.
- **Post-Comp Written**—The total amount of data written to the file system after compression has been performed, as shown in GiBs.
- **Total Comp Factor**—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 a Date Range (that is, 1w, 1m, 3m,1y, or All) above the graph lets you change the number of days of data shown on the graph, from one week to all.

When the file system is full or nearly full

Data Domain systems have three progressive levels of being full. As each level is reached, more operations are progressively disallowed. At each level, deleting data and then performing a file system cleaning operation makes disk space available.

Note

The process of deleting files and removing snapshots does not immediately reclaim disk space, the next cleaning operation reclaims 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 send these alerts, add the user to the alert emailing list.

Note

To join the alert email list, see [Viewing and Clearing Alerts](#).

Managing file system operations

This section describes file system cleaning, sanitization, and performing basic operations.

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

Create a file system from the Data Management > File System page using the Summary tab.

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 (see the section on viewing system storage information for more information). 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 **Data Management > File System > Summary > Create**.

The File System Create Wizard is launched. Follow the instructions provided.

Enabling or disabling 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 **Data Management > File System > Summary**.
2. For **File System**, click **Enable** or **Disable**.
3. On the confirmation dialog, click **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" do not clear enough space for normal operations.

A file system may not be expandable, however, for these reasons:

- The file system is not enabled.
- There are no unused disks or enclosures in the Active, Retention, or Cloud tiers.
- An expanded storage license is not installed.
- There are not enough capacity licenses installed.

DD6300 systems support the option to use ES30 enclosures with 4 TB drives (43.6 TiB) at 50% utilization (21.8 TiB) in the active tier if the available licensed capacity is exactly 21.8 TiB. The following guidelines apply to using partial capacity shelves.

- No other enclosure types or drive sizes are supported for use at partial capacity.
 - A partial shelf can only exist in the Active tier.
 - Only one partial ES30 can exist in the Active tier.
 - Once a partial shelf exists in a tier, no additional ES30s can be configured in that tier until the partial shelf is added at full capacity.
-

Note

This requires licensing enough additional capacity to use the remaining 21.8 TiB of the partial shelf.

- If the available capacity exceeds 21.8 TB, a partial shelf cannot be added.
- Deleting a 21 TiB license will not automatically convert a fully-used shelf to a partial shelf. The shelf must be removed, and added back as a partial shelf.

To expand the file system:

Procedure

1. Select **Data Management > File System > Summary > Expand Capacity**.

The Expand File System Capacity wizard is launched. The **Storage Tier** drop-down list always contains Active Tier, and it may contain either Extended Retention Tier or Cloud Tier as a secondary choice. The wizard displays the current capacity of the file system for each tier as well as how much additional storage space is available for expansion.

Note

File system capacity can be expanded only if the physical disks are installed on the system and file system is enabled.

2. From the **Storage Tier** drop-down list, select a tier.
3. In the **Addable Storage** area, select the storage devices to use and click **Add to Tier**.
4. Follow the instructions in the wizard. When the confirmation page is displayed, click **Close**.

Destroying the file system

Destroying the file system should be done only under the direction of Customer Support. This action deletes all data in the 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 clean 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.


CAUTION

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. Select **Data Management > File System > Summary > 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

This section describes how to start, stop, and modify cleaning schedules.

Starting cleaning

To immediately start a cleaning operation.

Procedure

1. Select **Data Management > File System > Summary > Settings > Cleaning**.

The Cleaning tab of the File System Setting dialog displays the configurable settings for each tier.

2. For the active tier:
 - a. In the Throttle % text box, enter a system throttle amount. This is the percentage of CPU usage dedicated to cleaning. The default is 50 percent.
 - b. In the Frequency drop-down list, select one of these frequencies: Never, Daily, Weekly, Biweekly, and Monthly. The default is Weekly.
 - c. For At, configure a specific time.
 - d. For On, select a day of the week.
3. For the cloud tier:
 - a. In the Throttle % text box, enter a system throttle amount. This is the percentage of CPU usage dedicated to cleaning. The default is 50 percent.
 - b. In the Frequency drop-down list, select one of these frequencies: Never, After every 'N' Active Tier cleans.

Note

If a cloud unit is inaccessible when cloud tier cleaning runs, the cloud unit is skipped in that run. Cleaning on that cloud unit occurs in the next run if the cloud unit becomes available. The cleaning schedule determines the duration between two runs. If the cloud unit becomes available and you cannot wait for the next scheduled run, you can start cleaning manually.

4. Click **Save**.

Stopping cleaning

To immediately stop a cleaning operation.

Procedure

1. Select **Data Management > File System > Summary > Settings > Cleaning**.
The Cleaning tab of the File System Setting dialog displays the configurable settings for each tier.
2. For the active tier:
 - a. In the Frequency drop-down list, select Never.
3. For the cloud tier:
 - a. In the Frequency drop-down list, select Never.
4. Click **Save**.

Performing sanitization

To comply with government guidelines, system sanitization, also called data shredding, must be performed when classified or sensitive data is written to any system that is not approved to store such data.

When an incident occurs, the system administrator must take immediate action to thoroughly eradicate the data that was accidentally written. The goal is to effectively restore the storage device to a state as if the event never occurred. If the data leakage is with sensitive data, the entire storage will need to be sanitized using Data Domain Professional Services' Secure Data erasure practice.

The Data Domain sanitization command exists to enable the administrator to delete files at the logical level, whether a backup set or individual files. Deleting a file in most

file systems consists of just flagging the file or deleting references to the data on disk, freeing up the physical space to be consumed at a later time. However, this simple action introduces the problem of leaving behind a residual representation of underlying data physically on disks. Deduplicated storage environments are not immune to this problem.

Shredding data in a system implies eliminating the residual representation of that data and thus the possibility that the file may be accessible after it has been shredded. Data Domain's sanitization approach ensures is compliant with the 2007 versions of Department of Defense (DoD) 5220.22 of the following specifications:

- *US Department of Defense 5220.22-M Clearing and Sanitization Matrix*
- *National Institute of Systems and Technology (NIST) Special Publication 800-88 Guidelines for Media Sanitization*

Sanitizing deduplicated data

Data Domain systems sanitize data in place, in its native deduplicated state.

Deduplication storage systems extract common data patterns from files sent to the system and store only unique copies of these patterns, referencing all the redundant instances. Because these data patterns or segments may potentially be shared among many files in the system, the sanitization process must first determine whether each of the segments of the contaminated file are shared with a clean file and then erase only those segments that are not shared, along with any contaminated metadata.

All storage tiers, caches, unused capacity, and free space are cleared so that every copy of every segment that belongs exclusively to the deleted files is eradicated. The system reclaims and overwrites all of the storage occupied by these segments to effectively restore the storage device to a state as if the contaminated files never existed in that system.

Sanitization level 1: data clearing or shredding

If the data you need to remove is unclassified, as defined in the "US Department of Defense 5220.22-M Clearing and Sanitization Matrix," Level 1 sanitization can be used to overwrite the affected storage once. This provides the basis for handling most data shredding and system sanitization cases.

The Data Domain system sanitization feature ensures that every copy of every segment that belongs only to erased files is overwritten using a single-pass zeroization mechanism. Clean data in the system being sanitized is online and available to users.

Procedure

1. Delete the contaminated files or backups through the backup software or corresponding client. In the case of backups, be sure to manage the backup software appropriately to ensure that related files on that image are reconciled, catalog records are managed as required, and so forth.
2. Run the `system sanitize start` command on the contaminated Data Domain system to cause all previously used space in it to be overwritten once (see the figure below).
3. Wait for the affected system to be sanitized. Sanitization can be monitored by using the `system sanitize watch` command.

If the affected Data Domain system has replication enabled, all the systems containing replicas need to be processed in a similar manner. Depending on how much data exists in the system and how it is distributed, the `system`

`sanitize` command could take some time. However, during this time, all clean data in the system is available to users.

Sanitization level 2: full system sanitization

If the data you need to remove is classified, as defined in the "US Department of Defense 5220.22-M Clearing and Sanitization Matrix," Level 2 sanitization, or full system sanitization, is now required.

Data Domain recommends Blancco for multi-pass overwrites with any overwrite pattern and a certificate. This provides the basis for handling universal Department of Defense requirements where complete system sanitization is required. For more information, go to:

https://www.emc.com/auth/rcoll/servicekitdocument/cp_datadomainsdataerase_psbasdddde.pdf

Modifying basic settings

Change the type of compression used, marker types, Replica write status, and Staging Reserve percentage, as described in this section.

Changing local compression

Use the General tab of the File System Settings dialog to configure the local compression type.

Note

Do not change the type of local compression unless it is necessary.

Procedure

1. Select **Data Management > File System > Summary > Settings > General**.
2. From the Local Compression Type drop-down list, select a compression type.

Table 93 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.

3. Click **Save**.

Changing read-only settings

Change the replica to writable. Some backup applications must see the replica as writable to do a restore or vault operation from the replica.

Procedure

1. Select **Data Management > File System > Summary > Settings > General**.
2. In the Report Replica as Writable area, toggle between **Disabled** and **Enabled** as appropriate.
3. Click **Save**.

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 NetWorker and Symantec's NetBackup (NBU), it does not require a license, and is disabled by default.

Note

The DD 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.

Configuring disk staging

Enable disk staging and specify the staging reserve percentage.

Procedure

1. Select **Data Management > File System > Summary > Settings > General**.
2. In the Staging Reserve area, toggle between **Disabled** and **Enabled** as appropriate.

3. If Staging Reserve is enabled, enter a value in the % of Total Space box.

This value represents the percentage of the total disk space to be reserved for disk staging, typically 20 to 30%.

4. Click **Save**.

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 DD 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 *How EMC Data Domain Systems Integrate into the Storage Environment*. You can use these matrices and integration guides to troubleshoot vendor-related issues.

SSD Random workload share

The value for the threshold at which to cap random I/O on the Data Domain system can be adjusted from the default value to accommodate changing requirements and I/O patterns.

By default, the Data Domain system sets the SSD random workload share at 40%. This value can be adjusted up or down as needed. Select **Data Management > File System > Summary > Settings > Workload Balance**, and adjust the slider.

Click **Save**.

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.

Performing a fast copy operation

Copy a file or directory tree from a Data Domain system source directory to another destination on the Data Domain system.

Procedure

1. Select **Data Management > File System > Summary > Fast Copy**.

The Fast Copy dialog 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 a lower case L 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

MTrees

This chapter includes:

- [MTrees overview](#) 200
- [Monitoring MTree usage](#) 207
- [Managing MTree operations](#) 210

MTrees overview

An MTree is a logical partition of the file system.

You can use MTrees in the following ways: for DD Boost storage units, DD VTL pools, or an NFS/CIFS share. MTrees allow granular management of snapshots, quotas, and DD 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

There can be up to the maximum configurable MTrees designated for MTree replication contexts.

Do not place user files in the top-level directory of an MTree.

MTree limits

MTree limits for Data Domain systems

Table 94 Supported MTrees

Data Domain System	DD OS Version	Supported configurable MTrees	Supported concurrently active MTrees
DD9800	6.0 and later	256	256
DD9500	5.7 and later	256	256
DD6800, DD9300	6.0 and later	128	128
DD6300	6.0 and later	100	32
DD990, DD4200, DD4500, DD7200	5.7 and later	128	128
All other DD systems	5.7 and later	100	Up to 32 based on the model
DD9500	5.6	100	64
DD990, DD890	5.3 and later	100	Up to 32 based on the model
DD7200, DD4500, DD4200	5.4 and later	100	Up to 32 based on the model
All other DD systems	5.2 and later	100	Up to 14 based on the model

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 DD 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 the section regarding MTree quota configuration for more information.

Quota enforcement

Enable or disable quota enforcement.

About the MTree panel

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.

Select **Data Management > MTree**.

- Select 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.

Table 95 MTree overview information

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

View important file system statistics.

View detail information

Select an MTree to view information.

Table 96 MTree detail information for a selected MTree

Item	Description
Full Path	The pathname of the MTree.
Pre-Comp Used	The current amount of raw data from the backup application that has been written to the MTree.
Status	<p>The status of the MTree (combinations are supported). Status can be:</p> <ul style="list-style-type: none"> • D: Deleted • RO: Read-only • RW: Read/write • RD: Replication destination • RLCE: DD Retention Lock Compliance enabled • RLCD: DD Retention Lock Compliance disabled • RLGE: DD Retention Lock Governance enabled • RLGD: DD Retention Lock Governance disabled
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"> • Yes—The MTree or its parent directory is shared. • Partial—The subdirectory under this MTree is shared. • No—This MTree and its parent or subdirectory are not shared. <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"> • Yes—The MTree or its parent directory is exported. • Partial—The subdirectory under this MTree is exported. • No—This MTree and its parent or subdirectory are not exported. <p>Click the NFS link to go to the NFS view.</p>

Table 96 MTree detail information for a selected MTree (continued)

Item	Description
DD Boost Storage Unit	<p>The DD Boost export status. Status can be:</p> <ul style="list-style-type: none"> • Yes—The MTree is exported. • No—This MTree is not exported. • Unknown—There is no information. <p>Click the DD Boost link to go to the DD Boost view.</p>
VTL Pool	<p>If applicable, the name of the DD VTL pool that was converted to an MTree.</p>
vDisk Pool	<p>vDisk report status. Status can be:</p> <ul style="list-style-type: none"> • Unknown— vDisk service is not enabled. • No— vDisk service is enabled but the MTree is not a vDisk pool. • Yes— vDisk service is enabled and the MTree is a vDisk pool.
Physical Capacity Measurements	
Used (Post-Comp)	<p>MTree space that is used after compressed data has been ingested.</p>
Compression	<p>Global Comp-factor.</p>
Last Measurement Time	<p>Last time the system measured the MTree.</p>
Schedules	<p>Number of schedules assigned.</p> <p>Click Assign to view and assign schedules to the MTree.</p> <ul style="list-style-type: none"> • Name: The schedule name. • Status: Enabled or Disabled • Priority: <ul style="list-style-type: none"> ▪ Normal— Submits a measurement task to the processing queue. ▪ Urgent— Submits a measurement task to the front of the processing queue. • Schedule: Time the task runs. • MTree Assignments: Number of MTrees the schedule is assigned to.
Submitted Measurements	<p>Displays the post compression status for the MTree.</p> <p>Click Measure Now to submit a manual post compression job for the MTree and select a priority for the job.</p> <ul style="list-style-type: none"> • 0— No measurement job submitted. • 1— 1 measurement job running. • 2— 2 measurement jobs running.
Snapshots	<p>Displays these statistics:</p>

Table 96 MTree detail information for a selected MTree (continued)

Item	Description
	<ul style="list-style-type: none"> Total Snapshots Expired Unexpired Oldest Snapshot Newest Snapshot Next Scheduled Assigned Snapshot Schedules <p>Click Total Snapshots to go to the Data Management > Snapshots view.</p> <p>Click Assign Schedules to configure snapshot schedules.</p>

View MTree replication information

Display MTree replication configuration.

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.

Table 97 MTree replication information

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.
Sync As Of	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.

- 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 the section about creating a snapshot schedule for instructions).

Table 98 MTree snapshot information

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 DD Retention Lock software options, summary information about the DD Retention Lock configuration displays.

Note

For information on how to manage DD Retention Lock for an MTree, see the section about working with DD Retention Lock.

Table 99 DD Retention Lock information

Item	Description
Status	Indicates whether DD Retention Lock is enabled or disabled.
Retention Period	Indicates the minimum and maximum DD Retention Lock time periods.
UUID	Shows either: <ul style="list-style-type: none"> the unique identification number generated for an MTree when the MTree is enabled for DD Retention Lock that the DD Retention Lock on a file in the MTree has been reverted

Enabling and managing DD Retention Lock settings

Use the DD Retention Lock area of the GUI to modify retention lock periods.

Procedure

1. Select **Data Management > MTree > Summary**.
2. In the Retention Lock area, click **Edit**.
3. In the Modify Retention Lock dialog box, select **Enable** to enable DD 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 panel.
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 DD Retention Lock summary area.

About the space usage view (MTrees)

Display a visual representation of data usage for an MTree at certain points in time.

Select **Data Management > MTree > Space Usage**.

- 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.

Note

For the MTrees Space Usage view, the system displays only pre-compressed information. Data can be shared between MTrees so compressed usage for a single MTree cannot be provided.

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 daily written view (MTrees)

Display the flow of data over the last 24 hours. Data amounts are shown over time for pre- and post-compression.

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 Written**—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 Written**—The total amount of data written to the MTree after compression has been performed, as shown in GiBs.
- **Total Comp Factor**—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 (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.

Below the Daily Written graph, the following totals display for the current duration value:

- Pre-Comp Written
- Post-Comp Written
- Global-Comp Factor
- Local-Comp Factor
- Total-Comp Factor

Monitoring MTree usage

Display space usage and data written trends for an MTree.

Procedure

- Select **Data Management > MTree**.

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 DD Retention Lock for an MTree.

The MTree view has an MTree overview panel and three tabs which are described in detail in these sections.

- [About the MTree panel](#) on page 201
- [About the summary view](#) on page 202
- [About the space usage view \(MTrees\)](#) on page 206
- [About the daily written view \(MTrees\)](#) on page 206

Note

Physical capacity measurement (PCM) provides space usage information for MTrees. For more information about PCM, see the section regarding understanding physical capacity measurement.

Understanding physical capacity measurement

Physical capacity measurement (PCM) provides space usage information for a sub-set of storage space. From the DD System Manager, PCM provides space usage

information for MTrees, but from the command line interface you can view space usage information for MTrees, tenants, tenant units, and pathsets.

Once a path is selected for PCM, all paths underneath it are automatically included. Do not select a child path after its parent path is already selected. For example, if `/data/col1/mtree3` is selected, do not select any subdirectories under `mtree3`.

The *Data Domain Operating System Command Reference Guide* provides more information about using PCM from the command line.

Enabling, disabling, and viewing physical capacity measurement

Physical capacity measurement provides space usage information for an MTree.

Procedure

1. Select **Data Management > File System > File System**.
The system displays the Summary tab in the File System panel.
2. Click **Enable** to the right of **Physical Capacity Measurement Status** to enable PCM.
3. Click **Details** to the right of **Physical Capacity Measurement Status** to view currently running PCM jobs.
 - **MTree**: The MTree that PCM is measuring.
 - **Priority**: The priority (normal or urgent) for the task.
 - **Submit Time**: The time the task was requested.
 - **Duration**: The length of time PCM ran to accomplish of the task.
4. Click **Disable** to the right of **Physical Capacity Measurement Status** to disable PCM and cancel all currently running PCM jobs.

Initializing physical capacity measurement

Physical capacity measurement (PCM) initialization is a one-time action that can take place only if PCM is enabled and the cache has not been initialized. It cleans the caches and enhances measuring speed. During the initialization process, you can still manage and run PCM jobs.

Procedure

1. Select **Data Management > File System > Configuration**.
2. Click **Initialize** under Physical Capacity Measurement to the right of Cache.
3. Click **Yes**.

Managing physical capacity measurement schedules

Create, edit, delete, and view physical capacity measurement schedules. This dialog only displays schedules created for MTrees and schedules that currently have no assignments.

Procedure

1. Select **Data Management > MTree > Manage Schedules**.
 - Click **Add (+)** to create a schedule.
 - Select a schedule and click **Modify** (pencil) to edit the schedule.
 - Select a schedule and click **Delete (X)** to delete a schedule.

- Optionally, click the heading names to sort by schedule: **Name**, **Status** (Enabled or Disabled) **Priority** (Urgent or Normal), **Schedule** (schedule timing), and **MTree Assignments** (the number of MTrees the schedule is assigned to).

Creating physical capacity measurement schedules

Create physical capacity measurement schedules and assign them to MTrees.

Procedure

- Select **Data Management > MTree > Manage Schedules**.
- Click **Add (+)** to create a schedule.
- Enter the name of the schedule.
- Select the status:
 - Normal**: Submits a measurement task to the processing queue.
 - Urgent**: Submits a measurement task to the front of the processing queue.
- Select how often the schedule triggers a measurement occurrence: every **Day**, **Week**, or **Month**.
 - For **Day**, select the time.
 - For **Week**, select the time and day of the week.
 - For **Month**, select the time, and days during the month.
- Select MTree assignments for the schedule (the MTrees that the schedule will apply to):
- Click **Create**.
- Optionally, click on the heading names to sort by schedule: **Name**, **Status** (Enabled or Disabled) **Priority** (Urgent or Normal), **Schedule** (schedule timing), and **MTree Assignments** (the number of MTrees the schedule is assigned to).

Editing physical capacity measurement schedules

Edit a physical capacity measurement schedule.

Procedure

- Select **Data Management > MTree > Manage Schedules**.
- Select a schedule and click **Modify** (pencil).
- Modify the schedule and click **Save**.
Schedule options are described in the Creating physical capacity measurement schedules topic.
- Optionally, click the heading names to sort by schedule: **Name**, **Status** (Enabled or Disabled) **Priority** (Urgent or Normal), **Schedule** (schedule timing), and **MTree Assignments** (the number of MTrees the schedule is assigned to).

Assigning physical capacity measurement schedules to an MTree

Attach schedules to an MTree.

Before you begin

Physical capacity measurement (PCM) schedules must be created.

Note

Administrators can assign up to three PCM schedules to an MTree.

Procedure

1. Select **Data Management > MTree > Summary**.
2. Select MTrees to assign schedules to.
3. Scroll down to the Physical Capacity Measurements area and click **Assign** to the right of Schedules.
4. Select schedules to assign to the MTree and click **Assign**.

Starting physical capacity measurement immediately

Start the measurement process as soon as possible.

Procedure

1. Select **Data Management > MTree > Summary**.
2. Scroll down to the Physical Capacity Measurements area and click **Measure Now** to the right of Submitted Measurements.
3. Select **Normal** (Submits a measurement task to the processing queue), or **Urgent** (Submits a measurement task to the front of the processing queue).
4. Click **Submit**.

Setting the physical capacity measurement throttle

Set the percentage of system resources that are dedicated to physical capacity measurement.

Procedure

1. Select **Data Management > File System > Configuration**.
2. In the Physical Capacity Measurement area, click **Edit** to the left of Throttle.
- 3.

Option	Description
Click Default	Enters the 20% system default.
Type throttle percent	The percentage of system resources that are dedicated to physical capacity measurement.

4. Click **Save**.

Managing MTree operations

This section describes MTree creation, configuration, how to enable and disable MTree quotas, and so on.

Creating an MTree

An MTree is a logical partition of the file system. Use MTrees in for DD Boost storage units, DD VTL pools, or an NFS/CIFS share.

MTrees are created in the area `/data/col1/mtree_name`.

Procedure

1. Select **Data Management > MTree**.
2. In the MTree overview area, click **Create**.
3. 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 (')
4. Set storage space restrictions for the MTree to prevent it from consuming excessive space. Enter a soft or hard limit quota setting, or both. 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.

5. 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

Set the storage space restriction for an MTree, Storage Unit, or DD VTL pool.

The **Data Management > Quota** page shows the administrator how many MTrees have no soft or hard quotas set. For MTrees with quotas set, the page shows the percentage of pre-compressed soft and hard limits used.

Consider the following information when managing quotas.

- 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 DD VTL, DD Boost, CIFS, and NFS.
- Snapshots are not counted.
- Quotas cannot be set on the `/data/coll/backup` directory.
- The maximum quota value allowed is 4096 PiB.

Configure MTree quotas

Use the MTree tab or the Quota tab to configure MTree quotas.

Procedure

1. Select one of the following menu paths:
 - Select **Data Management > MTree**.
 - Select **Data Management > Quota**.
2. Select only one MTree in the MTree tab, or one or more MTrees in the Quota tab.

Note

Quotas cannot be set on the `/data/coll/backup` directory.

3. In the MTree tab, click the **Summary** tab, and then click the **Configure** button in the Quota area.
4. In the Quota tab, click the **Configure Quota** button.

Configuring MTree quotas

Enter values for hard and soft quotas and select the unit of measurement.

Procedure

1. 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.
2. Click **OK**.

Deleting an MTree

Removes the MTree from the MTree table. The MTree data is deleted at the next 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 by the cleaning operation.

Procedure

1. Select **Data Management > MTree**.
2. Select an MTree.
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.

Undeleting an MTree

Undelete retrieves a deleted MTree and its data and places it back in the MTree table.

An undelete of an MTree retrieves a deleted MTree and its data and places it back in the MTree table.

An undelete is possible only if file cleaning has not been run after the MTree was marked for deletion.

Note

You can also use this procedure to undelete a storage unit.

Procedure

1. Select **Data Management > MTree > More Tasks > Undelete**.
2. Select the checkboxes of the MTrees you wish to bring back and click **OK**.
3. Click **Close** in the Undelete MTree Status dialog box after viewing the progress.

The recovered MTree displays in the MTree table.

Renaming an MTree

Use the Data Management MTree GUI to rename MTrees.

Procedure

1. Select **Data Management > MTree**.
2. Select an MTree in the MTree table.
3. Select the Summary tab.
4. In the Detailed Information overview area, click **Rename**.
5. Enter the name of the MTree in the New MTree Name text box.

See the section about creating an MTree for a list of allowed characters.

6. Click **OK**.

The renamed MTree displays in the MTree table.

CHAPTER 7

Snapshots

This chapter includes:

- [Snapshots overview](#) 216
- [Monitoring snapshots and their schedules](#) 217
- [Managing snapshots](#) 218
- [Managing snapshot schedules](#) 220
- [Recover data from a snapshot](#) 222

Snapshots overview

This chapter describes how to use the snapshot feature with MTrees.

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, and you can manage MTree snapshots and schedules and display information about the status of existing snapshots.

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 regarding viewing MTree snapshot information.

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

This section provides detailed and summary information about the status of snapshots and snapshot schedules.

About the snapshots view

The topics in this section describe the Snapshot view.

Snapshots overview panel

View the total number of snapshots, the number of expired snapshots, unexpired snapshots, and the time of the next cleaning.

Select **Data Management > Snapshots**.

Table 100 Snapshot overview panel 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 cleaning 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 cleaning operation will be performed.

Snapshots view

View snapshot information by name, by MTree, creation time, whether it is active, and when it expires.

The Snapshots tab displays a list of snapshots and lists the following information.

Table 101 Snapshot 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"> • Name—Name of the snapshot (wildcards are accepted). • Year—Drop-down list to select the year.
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

View the days snapshots will be taken, the times, the time they will be retained, and the naming convention.

Table 102 Snapshot schedule 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”).

1. Select a schedule in the Schedules tab. The Detailed Information area appears listing the MTrees that share the same schedule with the selected MTree.
2. Click the **Add/Remove** button to add or remove MTrees from schedule list.

Managing snapshots

This section describes how to manage snapshots.

Creating a snapshot

Create a snapshot when an unscheduled snapshot is required.

Procedure

1. Select **Data Management > Snapshots** to open the Snapshots view.
2. In the Snapshots view, click **Create**.
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 panel and click **Add**.
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**.

Modifying a snapshot expiration date

Modify snapshot expiration dates to remove them or extend their life for auditing or compliance.

Procedure

1. Select **Data Management Snapshots** to open the Snapshots view.
 2. Click the checkbox of the snapshot entry in the list and click **Modify Expiration Date**.
-

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**.

Renaming a snapshot

Use the Snapshot tab to rename a snapshot.

Procedure

1. Select **Data Management > Snapshots** to open the Snapshots view.
2. Select the checkbox of the snapshot entry in the list and click **Rename**.
3. In the Name text field, enter a new name.
4. Click **OK**.

Expiring a snapshot

Snapshots cannot be deleted. To release disk space, expire snapshots and they will be deleted in the next cleaning cycle after the expiry date.

Procedure

1. Select **Data Management > Snapshots** to open the Snapshots view.
 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

Set up and manage a series of snapshots that will be automatically taken at regular intervals (a snapshot schedule).

Multiple snapshot schedules can be active at the same time.

Note

If multiple snapshots with the same name are scheduled to occur at the same time, only one is retained. Which one is retained is indeterminate, thus only one of the snapshots with that name should be scheduled for a given time.

Creating a snapshot schedule

Create a weekly or monthly snapshot schedule using the Data Management GUI.

Procedure

1. Select **Data Management > Snapshots > Schedules** to open the Schedules view.
2. Click **Create**.
3. In the **Name** text field, enter the name of the schedule.
4. In the **Snapshot Name Pattern** text box, enter a name pattern.
 Enter a string of characters and variables that translates to a snapshot name (for example, `scheduled-%Y-%m-%d-%H-%m`, translates to "scheduled-2012-04-12-17-33"). Use alphabetic characters, numbers, `_`, `-`, and variables that translate into current values.
5. Click **Validate Pattern & Update Sample**.
6. Click **Next**.
7. Select the date when the schedule will 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 will 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 panel, then click **Add** and **OK**.

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.

Modifying a snapshot schedule

Change the snapshot schedule name, date, and retention period.

Procedure

1. In the schedule list, select the schedule and click **Modify**.
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.

Deleting a snapshot schedule

Delete a snapshot schedule from the schedule list.

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

Use the fastcopy operation to retrieve data stored in a snapshot. See the section regarding fast copy operations.

CHAPTER 8

CIFS

This chapter includes:

- [CIFS overview](#) 224
- [Configuring SMB signing](#) 224
- [Performing CIFS setup](#) 225
- [Working with shares](#) 227
- [Managing access control](#) 232
- [Monitoring CIFS operation](#) 237
- [Performing CIFS troubleshooting](#) 240

CIFS overview

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/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, at the least, need access to the `/data/coll/backup` directory. 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, 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 DD Manager and the `cifs` command.

Note

- The DD System Manager **Protocols > 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 to display CIFS statistics and status. For complete information about the `cifs` command, see the *Data Domain Operating System Command Reference Guide*.
 - For information about the initial system configuration, see the *Data Domain Operating System 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 support.emc.com web site. Search for the complete name of the document using the Search field.
-

Configuring SMB signing

On a DD OS version that supports it, you can configure the SMB signing feature using the CIFS option called server signing.

This feature is disabled by default because it degrades performance. When enabled, SMB signing can cause a 29 percent (reads) to 50 percent (writes) throughput performance drop, although individual system performance will vary. There are three possible values for SMB signing: disabled, auto and mandatory:

- When SMB signing is set to disabled, SMB signing is disabled, this is the default.

- When SMB signing is set to required, SMB signing is required, and both computers in the SMB connection must have SMB signing enabled.

SMB Signing CLI Commands

```
cifs option set "server-signing" required
```

Sets server signing to required.

```
cifs option reset "server-signing"
```

Resets server signing to the default (disabled).

As a best practice, whenever you change the SMB signing options, disable and then enable (restart) CIFS service using the following CLI commands:

```
cifs disable
cifs enable
```

The DD System Manager interface displays whether the SMB signing option is disabled or set to auto or mandatory. To view this setting in the interface, navigate to: **Protocols > CIFS > Configuration tab**. In the Options area, the value for the SMB signing option will be disabled, auto or mandatory reflecting the value set using the CLI commands.

Performing CIFS setup

This section contains instructions about enabling CIFS services, naming the CIFS server, and so on.

HA systems and CIFS

HA systems are compatible with CIFS; however, if a CIFS job is in progress during a failover, the job will need to be restarted.

"/ddvar is an ext3 file system, and cannot be shared like a normal MTree-based share. The information in /ddvar will become stale when the active node fails over to the standby node because the filehandles are different on the two nodes. If /ddvar is mounted to access log files or upgrade the system, unmount and remount /ddvar if a failover has occurred since the last time /ddvar was mounted."

Preparing clients for access to Data Domain systems

Find documentation online.

Procedure

1. Log into the Online Support (support.emc.com) web site.
2. In the Search field, enter the name of the document that you are looking for.
3. Select the appropriate document, such as the *CIFS and Data Domain Systems Tech Note*.
4. Follow the instructions in the document.

Enabling CIFS services

Enable the client to access the system using the CIFS protocol.

After configuring a client for access to Data Domain systems, 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 **Protocols > 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 the section regarding setting authentication parameters.

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

Change the NetBIOS hostname with the CLI.

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

Set the Data Domain authentication parameters for working with CIFS.

Click the Configure link in to the left of the Authentication label in the Configuration tab. The system will navigate to the **Administration > Access > Authentication** tab where you can configure authentication for Active Directory, Kerberos, Workgroups, and NIS.

Setting CIFS options

View CIFS configuration, restrict anonymous connections.

Procedure

1. Select **Protocols > CIFS > Configuration**.
2. In the Options area, click **Configure Options**.
3. To restrict anonymous connections, click the checkbox of the **Enable** option in the Restrict Anonymous Connections area.

- In the Log Level area, click the drop-down list to select the level number.

The level is an integer from 1 (one) to 5 (five). One is the default system level that sends the least-detailed level of CIFS-related log messages, five results in the most detail. Log messages are stored in the file `/ddvar/log/debug/cifs/cifs.log`.

Note

A log level of 5 degrades system performance. Click the **Default** in the Log Level area after debugging an issue. This sets the level back to 1.

- In the Server Signing area, select:
 - Enabled** to enable server signing
 - Disabled** to disable server signing
 - Required** when server signing is required

Disabling CIFS services

Prevent clients from accessing the Data Domain system.

Procedure

- Select **Protocols > CIFS**.
- In the Status area, click **Disable**.
- 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/coll/backup`

A Data Domain system supports a maximum number of 3000 CIFS shares.¹ And 600 simultaneous connections are allowed. However, the maximum number of connections supported is based on system memory. See the section regarding setting the maximum open files on a connection for more information.

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.

1. May be affected by hardware limitations.

Procedure

1. Select **Protocols > CIFS** tabs to navigate to the CIFS view.
2. Ensure authentication has been configured, as described in the section regarding setting authentication parameters.
3. On the CIFS client, set shared directory permissions or security options.
4. On the CIFS view, click the Shares tab.
5. Click **Create**.
6. In the Create Shares dialog box, enter the following information:

Table 103 Shares dialog box information

Item	Description
Share Name	A descriptive name for the share.
Directory Path	The path to the target directory (for example, /data/col1/backup/dir1).
	Note col1 uses the lower case letter L followed by the number 1.
Comment	A descriptive comment about the share.

Note

The share name can be a maximum of 80 characters and cannot contain the following characters: \ / : * ? " < > | + [] ; , = or extended ASCII characters.

7. 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**.
Consider the following when entering the client name.
 - No blanks or tabs (white space) characters are allowed.
 - 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.
 Repeat this step for each client that you need to configure.
8. 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.
9. Click **OK**.

The newly created share appears at the end of the list of shares, located in the center of the Shares panel.

Modifying a share on a Data Domain system

Change share information and connections.

Procedure

1. Select **Protocols > CIFS > Shares** to navigate to the CIFS view, Shares tab.
2. Click the checkbox next the share that you wish to modify in the Share Name list.
3. Click **Modify**.
4. 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.
 - 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 the 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**.
5. 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.
6. Click **OK**.

Creating a share from an existing share

Create a share from an existing share 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 tab, click the checkbox for the share you wish to use as the source.
2. Click **Create From**.
3. Modify the share information, as described in the section about modifying a share on a Data Domain system.

Disabling a share on a Data Domain system

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**.
3. Click **Close**.

Enabling a share on a Data Domain system

Enable one or more existing shares.

Procedure

1. In the Shares tab, click the checkbox of the shares you wish to enable in the Share Name list.
2. Click **Enable**.
3. Click **Close**.

Deleting a share on a Data Domain system

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 appears.
3. Click **OK**.
The shares are removed.

Performing MMC administration

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

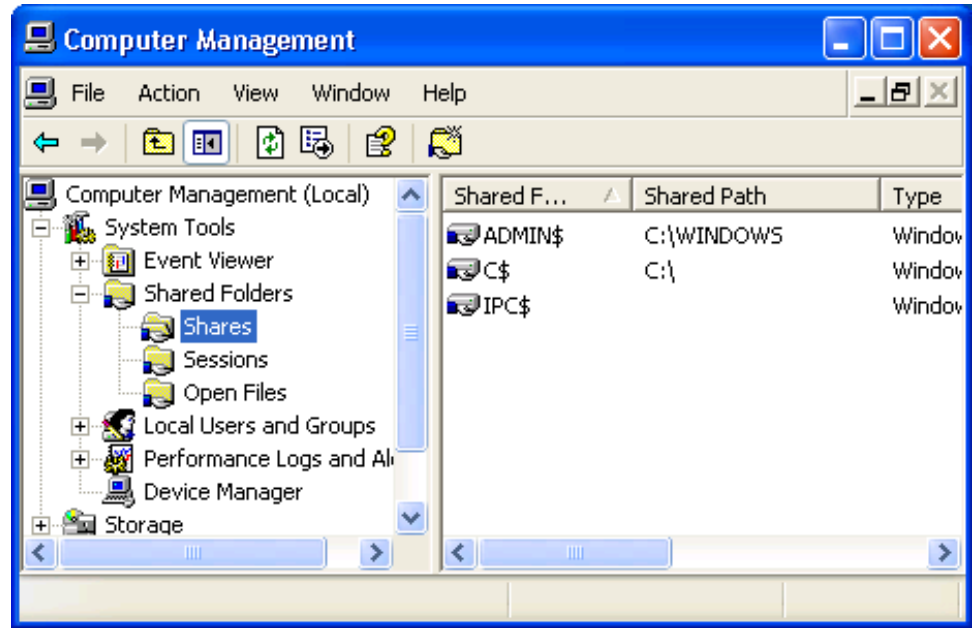
Use CIFS to connect to a Data Domain system and create a read-only backup subfolder.

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. For more information, see the section on creating a `/data/col1/backup` subfolder as read-only.

Figure 6 Computer Management dialog box



Creating a `\data\col1\backup` subfolder as read-only

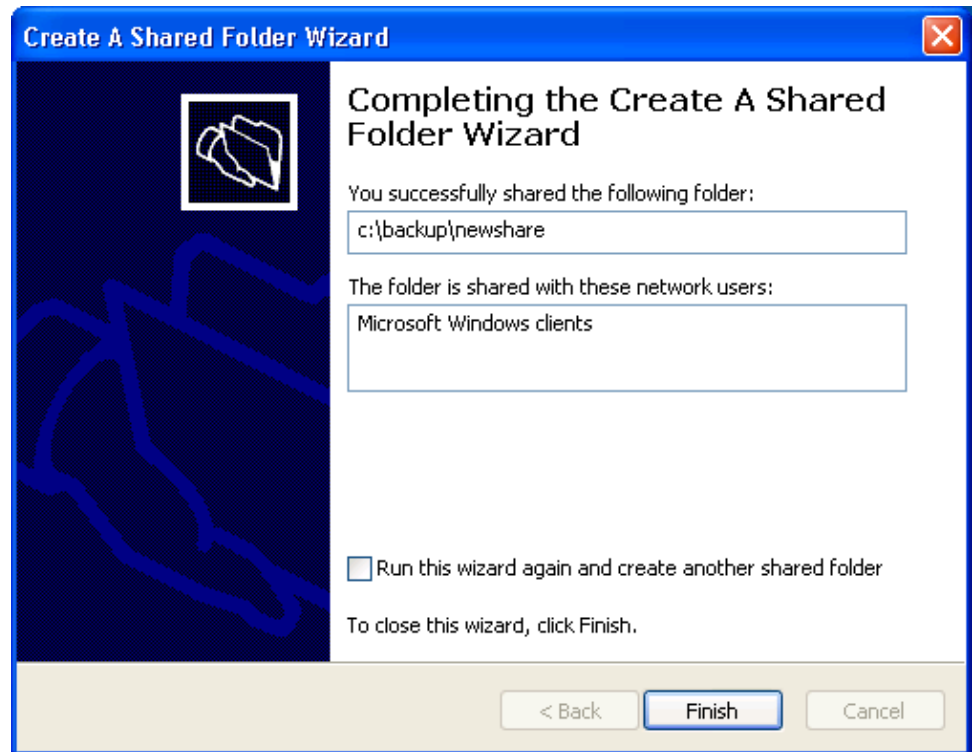
Enter a path, share name, and select permissions.

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\col1\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 7 Completing the Create a Shared Folder Wizard

7. The Completing dialog 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.

Displaying CIFS information

Display information about shared folders, sessions, and open files.

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 panel.

Managing access control

Access shared from a Windows client, provide administrative access, and allow access from trusted domain users.

Accessing shares from a Windows client

Use the command line to map a share.

Procedure

- From the Windows client use this DOS command:

```
net use drive: backup-location
```

For example, enter:

```
# \\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.

Providing domain users administrative access

Use the command line to add CIFS and include the domain name in the ssh instruction.

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" @dd22
```

Allowing administrative access to a Data Domain system for domain users

Use the command line to map a DD system default group number, and then enable CIFS administrative access.

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, and you can have up to 50 groups (dd admin group1 to dd admin group50).

Note

For a description of DD OS user roles and Windows groups, see the section about managing Data Domain systems.

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.

Restricting administrative access from Windows

Use the command line to prohibit access to users without a DD account.

Procedure

- 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

This sections contains information about ACLs, setting DACL and SACL permissions using Windows Explorer, and so on.

NT access control lists

Access control lists (ACLs) are enabled by default on the Data Domain system.



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:

Table 104 Permissions

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 as follows:

Table 105 Permissions Detail

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 file system, 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 file system, either through commands or through the CIFS protocol using the Windows Explorer GUI.

Setting DACL permissions using the Windows Explorer

Use Explorer properties settings to select DACL permissions.

Procedure

1. Right-click the file or folder and select **Properties**.
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**.

Setting SACL permissions using the Windows Explorer

Use Explorer properties settings to select SACL permissions.

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**.

Viewing or changing the current owner security ID (owner SID)

Use the Advanced Security Settings for ACL dialog box.

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.

This option 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.

Consider the following information while managing this option.

- 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 convert the metadata, consult your contracted support provider.

Monitoring CIFS operation

Monitoring CIFS Operation topics.

Displaying CIFS status

View and enable/disable CIFS status.

Procedure

1. In the DD System Manager, select **Protocols > CIFS**.
 - Status is either enabled and running, or disabled but CIFS authentication is running.
To enable CIFS, see the section regarding enabling CIFS services. To disable CIFS, see the section regarding disabling CIFS services.
 - **Connections** lists the tally of open connections and open files.

Table 106 Connections Details information

Item	Description
Open Connections	Open CIFS connections
Connection Limit	Maximum allowed connections
Open Files	Current open files
Max Open Files	Maximum number of open files on a Data Domain system

2. Click **Connection Details** to see more connection information.

Table 107 Connections Details information

Item	Description
Sessions	Active CIFS sessions
Computer	IP address or computer name connected with DDR for the session
User	User operating the computer connected with the DDR
Open Files	Number of open files for each session
Connection Time	Connection length in minutes
User	Domain name of computer
Mode	File permissions
Locks	Number of locks on the file
Files	File location

Display CIFS configuration

This section displays CIFS Configuration.

Authentication configuration

The information in the Authentication panel changes, depending on the type of authentication that is configured.

Click the Configure link in to the left of the Authentication label in the Configuration tab. The system will navigate to the **Administration > Access > Authentication** page where you can configure authentication for Active Directory, Kerberos, Workgroups, and NIS.

Active directory configuration

Table 108 Active directory configuration information

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 Controllers	The name of the configured domain controllers display 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.
Short Domain Name	The short domain name displays.

Workgroup configuration

Table 109 Workgroup configuration authentication information

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 shares information

This section displays shares information.

Viewing configured shares

View the list of configured shares.

Table 110 Configured shares information

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>Note</p> <p>col1 uses the lower case letter L 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

Display detailed information about a share by clicking a share name in the share list.

Table 111 Share information

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).
	<p>Note</p> <p>col1 uses the lower case letter L followed by the number 1.</p>
Directory Path Status	Indicates whether the configured directory path exists on the DDR. Possible values are Path Exists or Path Does Not Exist, the later indicating an incorrect or incomplete CIFS configuration.
Max Connections	The maximum number of connections allowed to the share at one time. The default value is Unlimited.
Comment	The comment that was configured when the share was created.
Share Status	The status of the share: either enabled or disabled.

- 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.

Displaying CIFS statistics

Use the command line to 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.

Displaying clients current activity

Use the command line to display CIFS sessions and open files information.

Procedure

- Enter: `cifs show active`

Results

Table 112 Sessions

Computer	User	Open files	Connect time (sec)	Idle time (sec)
::ffff: 10.25.132.84	ddve-25179109\sysadmin	1	92	0

Table 113 Open files

User	Mode	Locks	File
ddve-25179109\sysadmin	1	0	C:\data\col1\backup

Setting the maximum open files on a connection

Use the command line to set the maximum number of files that can be open concurrently.

Procedure

- Enter: `cifs option set max-global-open-files value`.

The *value* for the maximum global open files can be between 1 and the open files maximum limit. The maximum limit is based on the DDR system memory. For systems with greater than 12 GB, the maximum open files limit is 30,000. For systems with less than or equal to 12 GB, the maximum open files limit is 10,000.

Table 114 Connection and maximum open file limits

DDR Models	Memory	Connection Limit	Open File Maximum Limit
DD620, DD630, DD640	8 GB	300	10,000
DD640	16 GB	600	30,000
DD640	20 GB	600	30,000
DD860	36 GB	600	30,000
DD860, DD860ArT	72 GB	600	30,000
	96 GB	600	30,000
	128 GB	600	30,000
	256 GB	600	30,000

Note

The system has a maximum limit of 600 CIFS connections and 250,000 open files. However, if the system runs out of open files, the number of files can be increased.

Note

File access latencies are affected by the number of files in a directory. To the extent possible, we recommend directory sizes of less than 250,000. Larger directory sizes might experience slower responses to metadata operations such as listing the files in the directory and opening or creating a file.

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 DD System Manager, **Administration > Settings > Time and Date Settings** tab 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

Use the command line on a Windows domain controller to synchronize with an NTP server.

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 the section about working with time and date settings.

Synchronize from an NTP server

Configure the time server synchronization, as described in the section regarding working with time and date settings.

CHAPTER 9

NFS

This chapter includes:

- [NFS overview](#)244
- [Managing NFS client access to the Data Domain system](#)..... 245
- [Displaying NFS information](#)..... 249
- [Integrating a DDR into a Kerberos domain](#)..... 250
- [Add and delete KDC servers after initial configuration](#)..... 251

NFS overview

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/coll/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

On Data Domain systems, the `/ddvar/core` is on a separate partition. If you mount `/ddvar` only, you will not be able to navigate to `/ddvar/core` from the `/ddvar` mountpoint.

Clients, such as backup servers that perform backup and restore operations with a Data Domain System, need access to the `/backup` or `/data/coll/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 *Data Domain Operating System 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 *Data Domain Operating System 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.
-

HA systems and NFS

HA systems are compatible with NFS. If a NFS job is in progress during a failover, the job will **not** need to be restarted.

Note

/ddvar is an ext3 file system, and cannot be shared like a normal MTree-based share. The information in /ddvar will become stale when the active node fails over to the standby node because the filehandles are different on the two nodes. If /ddvar is mounted to access log files or upgrade the system, unmount and remount /ddvar if a failover has occurred since the last time /ddvar was mounted.

To create valid NFS exports that will failover with HA, the export needs to be created from the Active HA node, and generally shared over the failover network interfaces.

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.

Enabling NFS services

Enable NFS services to allow the client to access the system using the NFS protocol.

Procedure

1. Select **Protocols > NFS**.
The NFS view opens displaying the Exports tab.
2. Click **Enable**.

Disabling NFS services

Disable NFS services to prevent the client access to the system using the NFS protocol.

Procedure

1. Select the **Protocols > NFS** tabs.
The NFS view opens displaying the Exports tab.
2. Click **Disable**.

Creating an export

You can use Data Domain System Manager's Create button on the NFS view or use the Configuration Wizard to specify the NFS clients that can access the `/backup`, `/data/coll/backup`, `/ddvar`, `/ddvar/core` areas, or the `/ddvar/ext` area if it exists.

A Data Domain system supports a maximum of 2048 exports², with the number of connections scaling in accordance with system memory.

2. May be affected by hardware limitations.

Note

You have to assign client access to each export separately and remove access from each export separately. For example, a client can be removed from `/ddvar` and still have access to `/data/coll/backup`.

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 **ProtocolsNFS**.

The NFS view opens displaying the Exports tab.

2. Click **Create**.
 3. 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.

4. In the Clients area, select an existing client or click the + icon to create a client. The Client 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, an IPv4 or IPv6 IP address, an IPv4 address with either a netmask or prefix length, an IPv6 address with prefix 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.

General:

- Read-only permission (ro).
- Allow connections from ports below 1024 (secure) (default).

Anonymous UID/GID:

- Map requests from UID (user identifier) or GID (group identifier) 0 to the anonymous UID/GID (root _squash).
- Map all user requests to the anonymous UID/GID (all _squash).
- Use Default Anonymous UID/GID.

Allowed Kerberos Authentication Modes:

- Unauthenticated connections (sec=sys). Select to not use authentication.
- Authenticated Connections (sec=krb5).

Note

Integrity and Privacy are supported, although they might slow performance considerably.

c. Click **OK**.

5. Click **OK** to create the export.

Modifying an export

Change the directory path, domain name, and other options using the GUI.

Procedure

1. Select **Protocols > NFS**.

The NFS view opens displaying the Exports tab.

2. Click the checkbox of an export in the NFS Exports table.

3. Click **Modify**.

4. Modify the pathname in the Directory Path text box.

5. In the Clients area, select another client and click the pencil icon (modify), or click the + icon to create a client.

a. Enter a server name in the Client 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, an IPv4 or IPv6 IP address, an IPv4 address with either a netmask or prefix length, an IPv6 address with prefix 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.

General:

- Read-only permission (ro).
- Allow connections from ports below 1024 (secure) (default).

Anonymous UID/GID:

- Map requests from UID (user identifier) or GID (group identifier) 0 to the anonymous UID/GID (root_squash).
- Map all user requests to the anonymous UID/GID (all_squash).
- Use Default Anonymous UID/GID.

Allowed Kerberos Authentication Modes:

- Unauthenticated connections (sec=sys). Select to not use authentication.
- Authenticated Connections (sec=krb5).

Note

Integrity and Privacy are not supported.

c. Click **OK**.

6. Click **OK** to modify the export.

Creating an export from an existing export

Create an export from an existing export and then modify it as needed.

Procedure

1. In the NFS Exports tab, click the checkbox of the export you wish to use as the source.
2. Click **Create From**.
3. Modify the export information, as described in section about modifying an export.

Deleting an export

Delete an export from the NFS Exports tab.

Procedure

1. In the NFS Exports tab, click the checkbox of the export you wish to delete.
2. Click **Delete**.
3. Click **OK** and **Close** to delete the export.

Displaying NFS information

The topics in this section describe how to use the DD System Manager to monitor NFS client status and NFS configuration.

Viewing NFS status

Display whether NFS is active and Kerberos is enabled.

Procedure

- Click **Protocols > NFS**.

The top panel shows the operational status of NFS; for example, whether NFS is currently active and running, and whether Kerberos mode is enabled.

Note

Click **Configure** to view the **Administration > Access > Authentication** tab where you can configure Kerberos authentication.

Viewing NFS exports

See the list of clients allowed to access the Data Domain System.

Procedure

1. Click **Protocols > NFS**.

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.

2. 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, the system displays a list of clients.

Use the Filter By text box to sort by mount path.

Click **Update** for the system to refresh the table and use the filters supplied.

Click **Reset** for the system to clear the Path and Client filters.

Viewing active NFS clients

Display all clients that have been connected in the past 15 minutes and their mount path.

Procedure

- Select the **Protocols > NFS > Active Clients** tab.

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.

Click **Update** for the system to refresh the table and use the filters supplied.

Click **Reset** for the system to clear the Path and Client filters.

Integrating a DDR into a Kerberos domain

Set the domain name, the host name, and the DNS server for the DDR.

Enable the DDR to use the authentication server as a Key Distribution Center (for UNIX) and as a Distribution Center (for Windows Active Directory).

CAUTION

The examples provided in this description are specific to the operating system (OS) used to develop this exercise. You must use commands specific to your OS.

Note

For UNIX Kerberos mode, a keytab file must be transferred from the Key Distribution Center (KDC) server, where it is generated, to the DDR. If you are using more than one DDR, each DDR requires a separate keytab file. The keytab file contains a shared secret between the KDC server and the DDR.

Note

When using a UNIX KDC, the DNS server does not have to be the KDC server, it can be a separate server.

Procedure

1. Set the host name and the domain name for the DDR, using DDR commands.

```
net set hostname <host>
net set {domainname <local-domain-name>}
```

Note

The host name is the name of the DDR.

2. Configure NFS principal (node) for the DDR on the Key Distribution Center (KDC).

Example:

```
addprinc nfs/hostname@realm
```

Note

Hostname is the name for the DDR.

3. Verify that there are nfs entries added as principals on the KDC.

Example:

```
listprincs
nfs/hostname@realm
```

4. Add the DDR principal into a keytab file.

Example:

```
ktadd <keytab_file> nfs/hostname@realm
```

- Verify that there is an nfs keytab file configured on the KDC.

Example:

```
klist -k <keytab_file>
```

Note

The <keytab_file> is the keytab file used to configure keys in a previous step.

- Copy the keytab file from the location where the keys for NFS DDR are generated to the DDR in the /ddvar/ directory.

Table 115 Keytab destination

Copy file from:	Copy file to:
<keytab_file> (The keytab file configured in a previous step.)	/ddvar/

- Set the realm on the DDR, using the following DDR command:

```
authentication kerberos set realm <home realm> kdc-type <unix, windows.> kdc <IP address of server>
```

- When the kdc-type is UNIX, import the keytab file from /ddvar/ to /ddr/etc/, where the Kerberos configuration file expects it. Use the following DDR command to copy the file:

```
authentication kerberos keytab import
```

NOTICE

This step is required only when the kdc-type is UNIX.

Kerberos setup is now complete.

- To add a NFS mount point to use Kerberos, use the nfs add command.

See the *Data Domain Operating System Command Reference Guide* for more information.

- Add host, NFS and relevant user principals for each NFS client on the Key Distribution Center (KDC).

Example: `listprincs`

```
host/hostname@realm
nfs/hostname@realm
root/hostname@realm
```

- For each NFS client, import all its principals into a keytab file on the client.

Example:

```
ktadd -k <keytab_file> host/hostname@realm
ktadd -k <keytab_file> nfs/hostname@realm
```

Add and delete KDC servers after initial configuration

After you have integrated a DDR into a Kerberos domain, and thereby enabled the DDR to use the authentication server as a Key Distribution Center (for UNIX) and as a

Distribution Center (for Windows Active Directory), you can use the following procedure to add or delete KDC servers.

Procedure

1. Join the DDR to a Windows Active Directory (AD) server or a UNIX Key Distribution Center (KDC).

```
authentication kerberos set realm <home-realm> kdc-type {windows
[kdcs <kdc-list>] | unix kdcs <kdc-list>}
```

Example: `authentication kerberos set realm krb5.test kdc-type unix kdcs nfskrb-kdc.krb5.test`

This command joins the system to the `krb5.test` realm and enables Kerberos authentication for NFS clients.

Note

A keytab generated on this KDC must exist on the DDR to authenticate using Kerberos.

2. Verify the Kerberos authentication configuration.

```
authentication kerberos show config
```

```
Home Realm:      krb5.test
KDC List:        nfskrb-kdc.krb5.test
KDC Type:        unix
```

3. Add a second KDC server.

```
authentication kerberos set realm <home-realm> kdc-type {windows
[kdcs <kdc-list>] | unix kdcs <kdc-list>}
```

Example: `authentication kerberos set realm krb5.test kdc-type unix kdcs ostqa-sparc2.krb5.test nfskrb-kdc.krb5.test`

Note

A keytab generated on this KDC must exist on the DDR to authenticate using Kerberos.

4. Verify that two KDC servers are added.

```
authentication kerberos show config
```

```
Home Realm:      krb5.test
KDC List:        ostqa-sparc2.krb5.test, nfskrb-
kdc.krb5.test
KDC Type:        unix
```

5. Display the value for the Kerberos configuration key.

```
reg show config.keberos
```

```
config.kerberos.home_realm = krb5.test
config.kerberos.home_realm.kdc1 = ostqa-sparc2.krb5.test
config.kerberos.home_realm.kdc2 = nfskrb-kdc.krb5.test
config.kerberos.kdc_count = 2
config.kerberos.kdc_type = unix
```

6. Delete a KDC server.

Delete a KDC server by using the `authentication kerberos set realm <home-realm> kdc-type {windows [kdcs <kdc-list>] | unix kdcs <kdc-list>}` command without listing the KDC server that you want to delete.

For example, if the existing KDC servers are kdc1, kdc2, and kdc3, and you want to remove kdc2 from the realm, you could use the following example:

```
authentication kerberos set realm <realm-name> kdc-type  
<kdc_type> kdc kdc1,kdc3
```


CHAPTER 10

NFSv4

This chapter includes:

• Introduction to NFSv4	256
• ID Mapping Overview	257
• External formats	257
• Internal Identifier Formats	258
• When ID mapping occurs	259
• NFSv4 and CIFS/SMB Interoperability	260
• NFS Referrals	261
• NFSv4 and High Availability	263
• NFSv4 Global Namespaces	263
• NFSv4 Configuration	264
• Kerberos and NFSv4	265
• Enabling Active Directory	268
• LDAP and NFSv4	270

Introduction to NFSv4

Because NFS clients are increasingly using NFSv4.x as the default NFS protocol level, Data Domain systems can now employ NFSv4 instead of requiring the client to work in a backwards-compatibility mode.

In Data Domain systems, clients can work in mixed environments in which NFSv4 and NFSv3 must be able to access the same NFS exports.

The Data Domain NFS server can be configured to support NFSv4 and NFSv3, depending on site requirements. You can make each NFS export available to only NFSv4 clients, only NFSv3 clients, or both.

Several factors might affect whether you choose NFSv4 or NFSv3:

- **NFS client support**
Some NFS clients may support only NFSv3 or NFSv4, or may operate better with one version.
- **Operational requirements**
An enterprise might be strictly standardized to use either NFSv4 or NFSv3.
- **Security**
If you require greater security, NFSv4 provides a greater security level than NFSv3, including ACL and extended owner and group configuration.
- **Feature requirements**
If you need byte-range locking or UTF-8 files, you should choose NFSv4.
- **NFSv3 submounts**
If your existing configuration uses NFSv3 submounts, NFSv3 might be the appropriate choice.

NFSv4 compared to NFSv3 on Data Domain systems

NFSv4 provides enhanced functionality and features compared to NFSv3.

The following table compares NFSv3 features to those for NFSv4.

Feature	NFSv3	NFSv4
Standards-based Network Filesystem	Yes	Yes
Kerberos support	Yes	Yes
Kerberos with LDAP	Yes	Yes
Quota reporting	Yes	Yes
Multiple exports with client-based access lists	Yes	Yes
ID mapping	Yes	Yes
UTF-8 character support	No	Yes
File/directory-based Access Control Lists (ACL)	No	Yes
Extended owner/group (OWNER@)	No	Yes
File share locking	No	Yes
Byte range locking	No	Yes
DD-CIFS integration (locking, ACL, AD)	No	Yes

Feature	NFSv3	NFSv4
Stateful file opens and recovery	No	Yes
Global namespace and pseudoFS	No	Yes
Multi-system namespace using referrals	No	Yes

NFSv4 ports

You can enable or disable NFSv4 and NFSv3 independently. In addition, you can move NFS versions to different ports; both versions do not need to occupy the same port.

With NFSv4, you do not need to restart the Data Domain file system if you change ports. Only an NFS restart is required in such instances.

Like NFSv3, NFSv4 runs on Port 2049 as the default if it is enabled.

NFSv4 does not use portmapper (Port 111) or mountd (Port 2052).

ID Mapping Overview

NFSv4 identifies owners and groups by a common external format, such as `joe@example.com`. These common formats are known as identifiers, or IDs.

Identifiers are stored within an NFS server and use internal representations such as ID 12345 or ID S-123-33-667-2. The conversion between internal and external identifiers is known as ID mapping.

Identifiers are associated with the following:

- Owners of files and directories
- Owner groups of files and directories
- Entries in Access Control Lists (ACLs)

Data Domain systems use a common internal format for NFS and CIFS/SMB protocols, which allows files and directories to be shared between NFS and CIFS/SMB. Each protocol converts the internal format to its own external format with its own ID mapping.

External formats

The external format for NFSv4 identifiers follows NFSv4 standards (for example, RFC-7530 for NFSv4.0). In addition, supplemental formats are supported for interoperability.

Standard identifier formats

Standard external identifiers for NFSv4 have the format `identifier@domain`. This identifier is used for NFSv4 owners, owner-groups, and access control entries (ACEs). The domain must match the configured NFSv4 domain that was set using the `nfs option` command.

The following CLI example sets the NFSv4 domain to `mycorp.com` for the Data Domain NFS server:

```
nfs option set nfs4-domain mycorp.com
```

See client-specific documentation you have for setting the client NFS domain. Depending on the operating system, you might need to update a configuration file (for example, `/etc/idmapd.conf`) or use a client administrative tool.

Note

If you do not set the default value, it will follow the DNS name for the Data Domain system.

Note

The filesystem must be restarted after changing the DNS domain for the `nfs4-domain` to automatically update.

ACE extended identifiers

For ACL ACE entries, Data Domain NFS servers also support the following standard NFSv4 ACE extended identifiers defined by the NFSv4 RFC:

- OWNER@, The current owner of the file or directory
- GROUP@, the current owner group of the file or directory.
- The special identifiers INTERACTIVE@, NETWORK@, DIALUP@, BATCH@, ANONYMOUS@, AUTHENTICATED@, SERVICE@.

Alternative formats

To allow interoperability, NFSv4 servers on Data Domain systems support some alternative identifier formats for input and output.

- Numeric identifiers; for example, “12345”.
- Windows compatible Security identifiers (SIDs) expressed as “S-NNN-NNN-...”

See the sections on input mapping and output mapping for more information about restrictions to these formats.

Internal Identifier Formats

The Data Domain filesystem stores identifiers with each object (file or directory) in the filesystem. All objects have a numeric user ID (UID) and group ID (GID). These, along with a set of mode bits, allow for traditional UNIX/Linux identification and access controls.

Objects created by the CIFS/SMB protocol, or by the NFSv4 protocol when NFSv4 ACLs are enabled, also have an extended security descriptor (SD). Each SD contains the following:

- An owner security identifier (SID)
- An owner group SID
- A discretionary ACL (DACL)
- (Optional) A system ACL (SACL)

Each SID contains a relative ID (RID) and a distinct domain in a similar manner to Windows SIDs. See the section on NFSv4 and CIFS interoperability for more information on SIDs and the mapping of SIDs.

When ID mapping occurs

The Data Domain NFSv4 server performs mapping in the following circumstances:

- **Input mapping**
The Data Domain NFS server receives an identifier from an NFSv4 client. See [Input mapping](#) on page 259.
- **Output mapping:**
An identifier is sent from the Data Domain NFS server to the NFSv4 client. See [Output mapping](#) on page 259.
- **Credential mapping**
The RPC client credentials are mapped to an internal identity for access control and other operations. See [Credential mapping](#) on page 260.

Input mapping

Input mapping occurs when an NFSv4 client sends an identifier to the Data Domain NFSv4 server - setting up the owner or owner-group of a file, for example. Input mapping is distinct from credential mapping. For more information on credential mapping, see xxxx

Standard format identifiers such as `joe@mycorp.com` are converted into an internal UID/GID based on the configured conversion rules. If NFSv4 ACLs are enabled, a SID will also be generated, based on the configured conversion rules.

Numeric identifiers (for example, “12345”) are directly converted into corresponding UID/GIDs if the client is not using Kerberos authentication. If Kerberos is being used, an error will be generated as recommended by the NFSv4 standard. If NFSv4 ACLs are enabled, a SID will be generated based on the conversion rules.

Windows SIDs (for example, “S-NNN-NNN-...”) are validated and directly converted into the corresponding SIDs. A UID/GID will be generated based on the conversion rules.

Output mapping

Output mapping occurs when the NFSv4 server sends an identifier to the NFSv4 client; for example, if the server returns the owner or owner-group of a file.

1. If configured, the output might be the numeric ID.
This can be useful for NFSv4 clients that are not configured for ID mapping (for example, some Linux clients).
2. Mapping is attempted using the configured mapping services, (for example, NIS or Active Directory).
3. The output is a numeric ID or SID string if mapping fails and the configuration is allowed.
4. Otherwise, nobody is returned.

The `nfs` option `nfs4-idmap-out-numeric` configures the mapping on output:

- If `nfs` option `nfs4-idmap-out-numeric` is set to `map-first`, mapping will be attempted. On error, a numeric string is output if allowed. This is the default.
- If `nfs` option `nfs4-idmap-out-numeric` is set to `always`, output will always be a numeric string if allowed.

- If `nfs` option `nfs4-idmap-out-numeric` is set to `never`, mapping will be attempted. On error, `nobody@nfs4-domain` is the output. If the RPC connection uses GSS/Kerberos, a numeric string is never allowed and `nobody@nfs4-domain` is the output.

The following example configures the Data Domain NFS server to always attempt to output a numeric string on output. For Kerberos the name `nobody` is returned:

```
nfs option set nfs4-idmap-out-numeric always
```

Credential mapping

The NFSv4 server provides credentials for the NFSv4 client.

These credentials perform the following functions:

- Determine the access policy for the operation; for example, the ability to read a file.
- Determine the default owner and owner-group for new files and directories.

Credentials sent from the client may be `john_doe@mycorp.com`, or system credentials such as `UID=1000, GID=2000`. System credentials specify a UID/GID along with auxiliary group IDs.

If NFSv4 ACLs are disabled, then the UID/GID and auxiliary group IDs are used for the credentials.

If NFSv4 ACLs are enabled, then the configured mapping services are used to build an extended security descriptor for the credentials:

- SIDs for the owner, owner-group, and auxiliary group mapped and added to the Security Descriptor (SD).
- Credential privileges, if any, are added to the SD.

NFSv4 and CIFS/SMB Interoperability

The security descriptors used by NFSv4 and CIFS are similar from an ID mapping perspective, although there are differences.

You should be aware of the following to ensure for optimal interoperability:

- Active Directory should be configured for both CIFS and NFSv4, and the NFS ID mapper should be configured to use Active Directory for ID mapping.
- If you are using CIFS ACLs extensively, you can usually improve compatibility by also enabling NFSv4 ACLs.
 - Enabling NFSv4 ACLs allows NFSv4 credentials to be mapped to the appropriate SID when evaluating DACL access.
- The CIFS server receives credentials from the CIFS client, including default ACL and user privileges.
 - In contrast, the NFSv4 server receives a more limited set of credentials, and constructs credentials at runtime using its ID mapper. Because of this, the filesystem might see different credentials.

CIFS/SMB Active Directory Integration

The Data Domain NFSv4 server can be configured to use the Windows Active Directory configuration that is set with the Data Domain CIFS server.

The Data Domain system is mapped to use Active Directory if possible. This functionality is disabled by default, but you can enable it using the following command:

```
nfs option set nfs4-idmap-active-directory enabled
```

Default DACL for NFSv4

NFSv4 sets a different default DACL (discretionary access control list) than the default DACL supplied by CIFS.

Only OWNER@, GROUP@ and EVERYONE@ are defined in the default NFSv4 DACL. You can use ACL inheritance to automatically add CIFS-significant ACEs by default if appropriate.

System Default SIDs

Files and directories created by NFSv3, and NFSv4 without ACLs, use the default system domain, sometimes referred to as the default UNIX domain:

- User SIDs in the system domain have format S-1-22-1-N, where N is the UID.
- Group SIDs in the system domain have format S-1-22-2-N, when N is the GID.
For example, a user with UID 1234 would have an owner SID of S-1-22-1-1234.

Common identifiers in NFSv4 ACLs and SIDs

The EVERYONE@ identifier and other special identifiers (such as BATCH@, for example) in NFSv4 ACLs use the equivalent CIFS SIDS and are compatible.

The OWNER@ and GROUP@ identifiers have no direct correspondence in CIFS; they appear as the current owner and current owner-group of the file or directory.

NFS Referrals

The referral feature allows an NFSv4 client to access an export (or filesystem) in one or multiple locations. Locations can be on the same NFS server or on different NFS servers, and use either the same or different path to reach the export.

Because referrals are an NFSv4 feature, they apply only to NFSv4 mounts.

Referrals can be made to any server that uses NFSv4 or later, including the following:

- A Data Domain system running NFS with NFSv4 enabled
- Other servers that support NFSv4 including Linux servers, NAS appliances, and VNX systems.

A referral can use an NFS export point with or without a current underlying path in the Data Domain filesystem.

NFS exports with referrals can be mounted through NFSv3, but NFSv3 clients will not be redirected since referrals are a NFSv4 feature. This characteristic is useful in scaleout systems to allow exports to be redirected at a file-management level.

Referral Locations

NFSv4 referrals always have one or more locations.

These locations consist of the following:

- A path on a remote NFS server to the referred filesystem.
- One or more server network addresses that allow the client to reach the remote NFS server.

Typically when multiple server addresses are associated with the same location, those addresses are found on the same NFS server.

Referral location names

You can name each referral location within an NFS export. You can use the name to access the referral as well as to modify or delete it.

A referral name can contain a maximum of 80 characters from the following character sets:

- a-z
- A-Z
- 0-9
- "."
- ","
- "_"
- "-"

Note

You can include spaces as long as those spaces are embedded within the name. If you use embedded spaces, you must enclose the entire name in double quotes.

Names that begin with "." are reserved for automatic creation by the Data Domain system. You can delete these names but you cannot create or modify them using the command line interface (CLI) or system management services (SMS).

Referrals and Scaleout Systems

NFSv4 referrals and locations can better enable access if you are scaling out your Data Domain systems

Because your Data Domain system might or might not already contain a global namespace, the following two scenarios describe how you might use NFSv4 referrals:

- Your Data Domain system does not contain a global namespace.
 - You can use NFSv4 referrals to build that global namespace. System administrators can build these global namespaces, or you can use smart system manager (SM) element building referrals as necessary.
- Your Data Domain system already has a global namespace.
 - If your system has a global namespace with MTrees placed in specific nodes, NFS referrals can be created to redirect access to those MTrees to the nodes

added to the scaled-out system. You can create these referrals or have them performed automatically within NFS if the necessary SM or file manager (FM) information is available.

See the *Data Domain Operating System Administration Guide* for more information about MTrees.

NFSv4 and High Availability

With NFSv4, protocol exports (for example, `/data/coll/<mtree>`) are mirrored in a High Availability (HA) setup. However, configuration exports such as `/ddvar` are not mirrored.

The `/ddvar` filesystem is unique to each node of an HA pair. As a result, `/ddvar` exports and their associated client access lists are not mirrored to the standby node in an HA environment.

The information in `/ddvar` becomes stale when the active node fails over to the standby node. Any client permissions granted to `/ddvar` on the original active node must be recreated on the newly active node after a failover occurs.

You must also add any additional `/ddvar` exports and their clients (for example, `/ddvar/core`) that were created on the original active node to the newly active node after a failover occurs.

Finally, any desired `/ddvar` exports must be unmounted from the client and then remounted after a failover occurs.

NFSv4 Global Namespaces

The NFSv4 server provides a virtual directory tree known as a PseudoFS to connect NFS exports into a searchable set of paths.

The use of a PseudoFS distinguishes NFSv4 from NFSv3, which uses the MOUNTD auxiliary protocol.

In most configurations, the change from NFSv3 MOUNTD to NFSv4 global namespace is transparent and handled automatically by the NFSv4 client and server.

NFSv4 global namespaces and NFSv3 submounts

If you use NFSv3 export submounts, the global namespaces characteristic of NFSv4 might prevent submounts from being seen on the NFSv4 mount.

Example 1 NFSv3 main exports and submount exports

If NFSv3 has a main export and a submount export, these exports might use the same NFSv3 clients yet have different levels of access:

Export	Path	Client	Options
Mt1	<code>/data/coll/mt1</code>	<code>client1.example.com</code>	<code>ro</code>
Mt1-sub	<code>/data/coll/mt1/subdir</code>	<code>client1.example.com</code>	<code>rw</code>

In the previous table, the following applies to NFSv3:

Example 1 NFSv3 main exports and submount exports (continued)

- If client1.example.com mounts /data/coll/mt1, the client gets read-only access.
- If client1.example.com mounts /data/coll/mt1/subdir, the client gets read-write access.

NFSv4 operates in the same manner in regard to highest-level export paths. For NFSv4, client1.example.com navigates the NFSv4 PseudoFS until it reaches the highest-level export path, /data/coll/mt1, where it gets read-only access.

However, because the export has been selected, the submount export (Mt1-sub) is not part of the PseudoFS for the client and read-write access is not given.

Best practice

If your system uses NFSv3 exports submounts to give the client read-write access based on the mount path, you must consider this before using NFSv4 with these submount exports.

With NFSv4, each client has an individual PseudoFS.

Export	Path	Client	Options
Mt1	/data/coll/mt1	client1.example.com	ro
Mt1-sub	/data/coll/mt1/subdir	client2.example.com	rw

NFSv4 Configuration

The default Data Domain system configuration only enables NFSv3. To use NFSv4, you must first enable the NFSv4 server.

Enabling the NFSv4 Server

Procedure

1. Enter `nfs enable version 4` to enable NFSv4:

```
# nfs enable version 4
NFS server version(s) 3:4 enabled.
```

2. (Optional) If you want to disable NFSv3, enter `nfs disable version 3`.

```
# nfs disable version 3
NFS server version(s) 3 disabled.
NFS server version(s) 4 enabled.
```

After you finish

After the NFSv4 server is enabled, you might need to perform additional NFS configuration tasks specifically for your site. These tasks can include performing the following actions on the Data Domain system:

- Setting the NFSv4 domain
- Configuring NFSv4 ID mapping
- Configuring ACL (Access Control Lists)

Setting the default server to include NFSv4

The Data Domain NFS command option `default-server-version` controls which NFS version is enabled when you enter the `nfs enable` command without specifying a version.

Procedure

1. Enter the `nfs option set default-server-version 3:4` command:

```
# nfs option set default-server-version 3:4
NFS option 'default-server-version' set to '3:4'.
```

Updating existing exports

You can update existing exports to change the NFS version used by your Data Domain system.

Procedure

1. Enter the `nfs export modify all` command:

```
# nfs export modify all clients all options version=version
number
```

To ensure all existing clients have either version 3, 4, or both, you can modify the NFS version to the appropriate string. The following example shows NFS modified to include versions 3 and 4:

```
#nfs export modify all clients all options version=3:4
```

For more information about the `nfs export` command, see the *Data Domain Operating System Command Reference Guide* for more information.

Kerberos and NFSv4

Both NFSv4 and NFSv3 use the Kerberos authentication mechanism to secure user credentials.

Kerberos prevents user credentials from being spoofed in NFS packets and protects them from tampering en route to the Data Domain system.

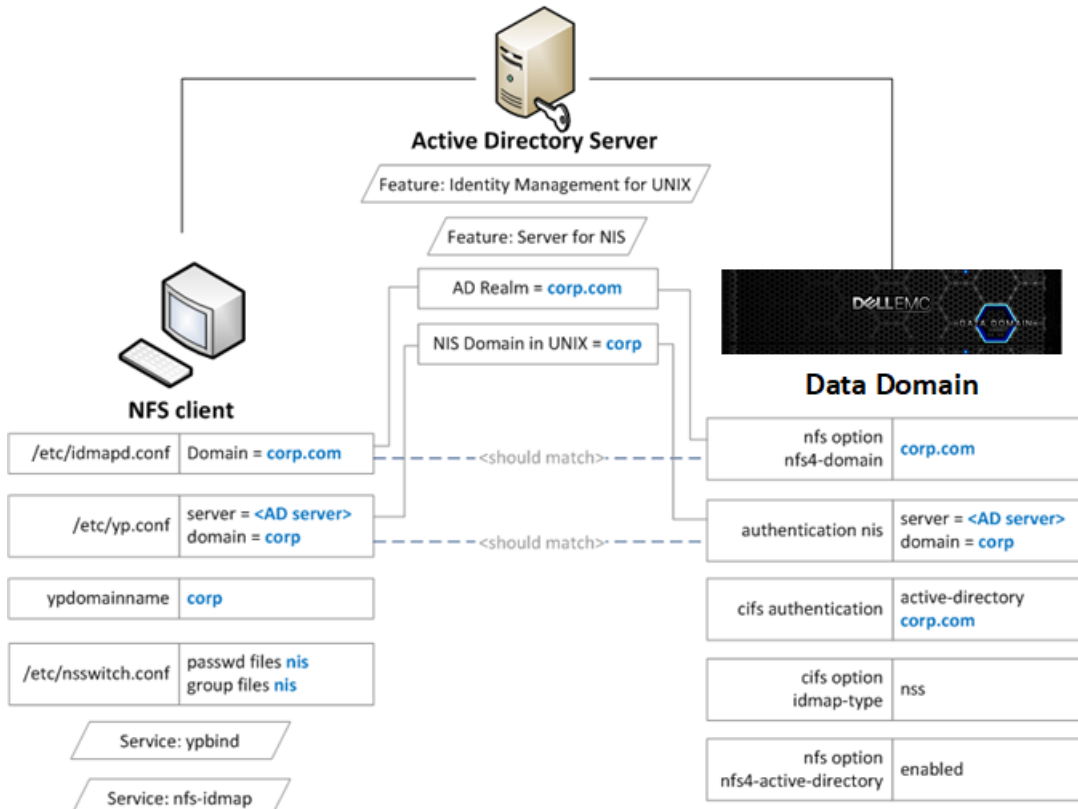
There are distinct types of Kerberos over NFS:

- Kerberos 5 (`sec=krb5`)
Use Kerberos for user credentials.
- Kerberos 5 with integrity (`sec=krb5i`)
Use Kerberos and check the integrity of the NFS payload using an encrypted checksum.
- Kerberos 5 with security (`sec=krb5p`)
Use Kerberos 5 with integrity and encrypt the entire NFS payload.

Note

`krb5i` and `krb5p` can both cause performance degradation due to additional computational overhead on both the NFS client and the Data Domain system.

Figure 8 Active Directory Configuration



You employ existing commands used for NFSv3 when configuring your system for Kerberos. See the `nfsv3` chapter of the *Data Domain Command Reference Guide* for more information.

Configuring Kerberos with a Linux-Based KDC

Before you begin

You should ensure that all your systems can access the Key Distribution Center (KDC).

If the systems cannot reach the KDC, check the domain name system (DNS) settings.

The following steps allow you to create keytab files for the client and the Data Domain system:

- In Steps 1-3 you create the keytab file for the Data Domain system.
- In Steps 4-5 you create the keytab file for the client.

Procedure

1. Create the `nfs/<ddr_dns_name>@<realm>` service principal.

```
kadmin.local: addprinc -randkey nfs/ddr12345.<domain-name>@<domain-name>
```

2. Export `nfs/<ddr_dns_name>@<realm>` to a keytab file.

```
kadmin.local: ktadd -k /tmp/ddr.keytab nfs/ddr12345.corp.com@CORP.COM
```

3. Copy the keytab file to the Data Domain system at the following location:

```
/ddr/var/krb5.keytab
```

4. Create one of the following principal for the client and export that principal to the keytab file:

```
nfs/<client_dns_name>@<REALM>
root/<client_dns_name>@<REALM>
```

5. Copy the keytab file to the client at the following location:

```
/etc/krb5.keytab
```

Note

It is recommended that you use an NTP server to keep the time synchronized on all entities.

Configuring the Data Domain System to Use Kerberos Authentication

Procedure

1. Configure the KDC and Kerberos realm on the Data Domain system by using the authentication command:

```
# authentication kerberos set realm <realm> kdc-type unix kdc
<kdc-server>
```

2. Import the keytab file:

```
# authentication kerberos keytab import
```

3. (Optional) Configure the NIS server by entering the following commands:

```
# authentication nis servers add <server>
# authentication nis domain set <domain-name>
# authentication nis enable
# filesys restart
```

4. (Optional) Make the `nfs4-domain` the same as the Kerberos realm using the `nfs option` command:

```
nfs option set nfs4-domain <kerberos-realm>
```

5. Add a client to an existing export by adding `sec=krb5` to the `nfs export add` command:

```
nfs export add <export-name> clients * options
version=4,sec=krb5
```

Configuring Clients

Procedure

1. Configure the DNS server and verify that forward and reverse lookups are working.

2. Configure the KDC and Kerberos realm by editing the `/etc/krb5.conf` configuration file.

You might need to perform this step based on the client operating system you are using.

3. Configure NIS or another external name mapping service.
4. (Optional) Edit the `/etc/idmapd.conf` file to ensure it is the same as the Kerberos realm.

You might need to perform this step based on the client operating system you are using.

5. Verify the keytab file `/etc/krb5.keytab` contains an entry for the `nfs/` service principal or the `root/` principal.

```
[root@fc22 ~]# klist -k
Keytab name: FILE:/etc/krb5.keytab
KVNO Principal
-----
-----
3 nfs/fc22.domain-name@domain-name
```

6. Mount the export using the `sec=krb5` option.

```
[root@fc22 ~]# mount ddr12345.<domain-name>:/data/coll/
mtree1 /mnt/nfs4 -o sec=krb5,vers=4
```

Enabling Active Directory

Configuring Active Directory authentication makes the Data Domain system part of a Windows Active Directory realm. CIFS clients and NFS clients use Kerberos authentication.

Procedure

1. Join an active directory realm using the `cifs set` command:

```
# cifs set authentication active-directory <realm>
```

Kerberos is automatically set up on the Data Domain system. the required `nfs/` service principal is automatically created on the KDC.

2. Configure NIS using the `authentication nis` command:

```
# authentication nis servers add <windows-ad-server>
# authentication nis domain set <ad-realm>
# authentication nis enable
```

3. Configure CIFS to use NSS for ID mapping by using `cifs` commands:

```
# cifs disable
# cifs option set idmap-type nss
# cifs enable
# filesys restart
```

4. Set the `nfs4-domain` to be the same as the Active Directory realm:

```
# nfs option set nfs4-domain <ad-realm>
```

5. Enable Active Directory for NFSv4 id mapping by using the `nfs` command:

```
# nfs option set nfs4-idmap-active-directory enabled
```

Configuring Active Directory

Procedure

1. Install the Active Directory Domain Services (AD DS) role on the Windows server.
2. Install the Identity Management for UNIX components.

```
C:\Windows\system32>Dism.exe /online /enable-feature /
featurename:adminui /all
C:\Windows\system32>Dism.exe /online /enable-feature /
featurename:nis /all
```

3. Verify the NIS domain is configured on the server.

```
C:\Windows\system32>nisadmin
The following are the settings on localhost

Push Interval : 1 days
Logging Mode  : Normal

NIS Domains
NIS Domain in AD  Master server  NIS Domain in UNIX
-----
corp              win-ad-server  corp
```

4. Assign AD users and groups UNIX UID/GIDs for the NFSv4 server.
 - a. Go to **Server Manager > Tools > Active Directory**.
 - b. Open the **Properties** for an AD user or group.
 - c. Under the UNIX Attributes tab, fill in the NIS domain, UID, and Primary GID fields.

Configuring clients on Active Directory

Procedure

1. Create a new AD user on the AD server to represent the NFS client's service principal.
2. Create the nfs/ service principal for the NFS client.

```
> ktpass -princ nfs/<client_dns_name>@<REALM> -mapuser nfsuser -
pass **** -out nfsclient.keytab
/crytp rc4-hmac-nt /ptype KRB5_NT_PRINCIPAL
```

3. (Optional) Copy the keytab file to /etc/krb5.keytab on the client.

The need to perform this step depends on which client OS you are using.

LDAP and NFSv4

Enabling LDAP allows you to use an existing OpenLDAP server or deployment with DDR for NFSv4 ID mapping, NFSv3 Kerberos with LDAP, or NFSv4 Kerberos with LDAP.

Configure LDAP servers

You can configure one or more LDAP servers at the same time.

Note

LDAP must be disabled when making any changes to the configuration.

Specify the LDAP server in one of the following formats:

- IPv4 address—`10.26.16.250`
- IPv4 address with port number—`10.26.16.251:400`
- IPv6 address—`[::ffff:9.53.96.21]`
- IPv6 address with port number—`[::ffff:9.53.96.21]:400`
- Hostname—`myldapserver`
- Hostname with port number—`myldapserver:400`

When configuring multiple servers:

- Separate each server with a space.
- The first server listed when using the `authentication ldap servers add` command becomes the primary server.
- If any of the servers cannot be configured, the command fails for all servers listed.

Procedure

1. Add one or more LDAP servers by using the `authentication ldap servers add` command:

```
# authentication ldap servers add 10.26.16.250 10.26.16.251:400
LDAP server(s) added
LDAP Server(s):      2
#   IP Address/Hostname
---   -----
1.   10.26.16.250 (primary)
2.   10.26.16.251:400
---   -----
```

2. Remove one or more LDAP servers by using the `authentication ldap servers del` command:

```
# authentication ldap servers del 10.26.16.251:400
LDAP server(s) deleted.
LDAP Servers: 1
#   Server
-   -----
1   10.26.16.250 (primary)
-   -----
```

3. Remove all LDAP servers by using the `authentication ldap servers reset` command:

```
# authentication ldap servers reset
LDAP server list reset to empty.
```

Configure the LDAP base suffix

The base suffix is the base DN for search and is where the LDAP directory begins searching.

Procedure

1. Set the LDAP base suffix by using the `authentication ldap base set` command:

```
# authentication ldap base set "dc=anvil,dc=team"
LDAP base-suffix set to "dc=anvil,dc=team".
```

2. Reset the LDAP base suffix by using the `authentication ldap base reset` command:

```
# authentication ldap base reset
LDAP base-suffix reset to empty.
```

Configure LDAP client authentication

Configure the account (Bind DN) and password (Bind PW) that is used to authenticate with the LDAP server and make queries.

You should always configure the Bind DN and password. Normally, LDAP servers require authenticated bind by default. If `client-auth` is not set, anonymous access is requested, providing no name or password. The output of `authentication ldap show` is as follows:

```
# authentication ldap show
LDAP configuration
  Enabled:          yes (*)
  Base-suffix:     dc=u2,dc=team
  Binddn:          (anonymous)
  Server(s):       1

#  Server
-  -----  -----
1  10.207.86.160  (primary)
-  -----  -----

Secure LDAP configuration
  SSL Enabled:     no
  SSL Method:      off
  tls_reqcert:    demand

(*) Requires a filesystem restart for the configuration to take effect.
```

If `binddn` is set using `client-auth` CLI, but `bindpw` is not provided, unauthenticated access is requested.

```
# authentication ldap client-auth set binddn
"cn=Manager,dc=u2,dc=team"
Enter bindpw:
** Bindpw is not provided. Unauthenticated access would be requested.
LDAP client authentication binddn set to "cn=Manager,dc=u2,dc=team".
```

Procedure

1. Set the Bind DN and password by using the `authentication ldap client-auth set binddn` command:

```
# authentication ldap client-auth set binddn
"cn=Administrator,cn=Users,dc=anvil,dc=team"
```

```
Enter bindpw:
LDAP client authentication binddn set to
"cn=Administrator,cn=Users,dc=anvil,dc=team".
```

2. **Reset the Bind DN and password by using the authentication ldap client-auth reset command:**

```
# authentication ldap client-auth reset
LDAP client authentication configuration reset to empty.
```

Enable LDAP

Before you begin

An LDAP configuration must exist before enabling LDAP. Additionally, you must disable NIS, ensure that the LDAP server is reachable, and be able to query the root DSE of the LDAP server.

Procedure

1. **Enable LDAP by using the authentication ldap enable command:**

```
# authentication ldap enable
```

The details of the LDAP configuration are displayed for you to confirm before continuing. To continue, type **yes** and restart the file system for LDAP configuration to take effect.

2. **View the current LDAP configuration by using the authentication ldap show command:**

```
# authentication ldap show
LDAP configuration
    Enabled:          no
    Base-suffix:     dc=anvil,dc=team
    Binddn:
cn=Administrator,cn=Users,dc=anvil,dc=team
    Server(s):       2
#   Server
-   -----
1   10.26.16.250    (primary)
2   10.26.16.251:400
-   -----

Secure LDAP configuration
    SSL Enabled:     no
    SSL Method:      off
    tls_reqcert:     demand
```

Basic LDAP and secure LDAP configuration details are displayed.

3. **View the current LDAP status by using the authentication ldap status command:**

```
# authentication ldap status
```

The LDAP status is displayed. If the LDAP status is not **good**, the problem is identified in the output. For example:

```
# authentication ldap status
Status: invalid credentials
```

or

```
# authentication ldap status
Status: invalid DN syntax
```


4. Disable LDAP by using the `authentication ldap disable` command:

```
# authentication ldap disable
LDAP is disabled.
```

Enable secure LDAP

You can configure DDR to use secure LDAP by enabling SSL.

Before you begin

If there is no LDAP CA certificate and `tls_reqcert` is set to `demand`, the operation fails. Import an LDAP CA certificate and try again.

If `tls_reqcert` is set to `never`, an LDAP CA certificate is not required. For more information, see [Configure LDAP server certificate verification with imported CA certificates](#) on page 273.

Procedure

1. Enable SSL by using the `authentication ldap ssl enable` command:

```
# authentication ldap ssl enable
Secure LDAP is enabled with 'ldaps' method.
```

The default method is secure LDAP, or *ldaps*. You can specify other methods, such as TLS:

```
# authentication ldap ssl enable method start_tls
Secure LDAP is enabled with 'start_tls' method.
```

2. Disable SSL by using the `authentication ldap ssl disable` command:

```
# authentication ldap ssl disable
Secure LDAP is disabled.
```

Configure LDAP server certificate verification with imported CA certificates

You can change the TLS request certificate behavior.

Procedure

1. Change the TLS request certificate behavior by using the `authentication ldap ssl set tls_reqcert` command.

Do not verify the certificate:

```
# authentication ldap ssl set tls_reqcert never
"tls_reqcert" set to "never". LDAP server certificate will not
be verified.
```

Verify the certificate:

```
# authentication ldap ssl set tls_reqcert demand
"tls_reqcert" set to "demand". LDAP server certificate will be
verified.
```

2. Reset the TLS request certificate behavior by using the `authentication ldap ssl reset tls_reqcert` command. The default behavior is `demand`:

```
# authentication ldap ssl reset tls_reqcert
tls_reqcert has been set to "demand". LDAP Server certificate
will be verified with imported CA certificate. Use "adminaccess"
CLI to import the CA certificate.
```

Manage CA certificates for LDAP

You can import or delete certificates and show current certificate information.

Procedure

1. Import a CA certificate for LDAP server certificate verification by using the `adminaccess certificate import` command.

Specify `ldap` for `ca` application:

```
# adminaccess certificate import{host application {all | aws-
federal | ddbost | https| keysecure | rkm | <application-
list>}| ca application { ldap }} [file <file-name>] Import host
or ca certificate
```

2. Delete a CA certificate for LDAP server certificate verification by using the `adminaccess certificate delete` command.

Specify `ldap` for application:

```
# adminaccess certificate delete
{ subject <subject-name> | fingerprint <fingerprint>}
[application { ldap }]
```

Specify `ldap` for `imported-ca` application:

```
# adminaccess certificate delete
{ imported-host application { all | aws-federal | ddbost |
https
| keysecure | rkm | <application-list>}
| imported-ca application { ldap }]
```

3. Show current CA certificate information for LDAP server certificate verification by using the `adminaccess certificate show` command:

```
# adminaccess certificate show imported-ca ldap
```

CHAPTER 11

Storage Migration

This chapter includes:

- [Storage migration overview](#)276
- [Migration planning considerations](#)..... 276
- [Viewing migration status](#)..... 278
- [Evaluating migration readiness](#)..... 279
- [Migrating storage using DD System Manager](#)..... 279
- [Storage migration dialog descriptions](#)..... 280
- [Migrating storage using the CLI](#)..... 283
- [CLI storage migration example](#)..... 284

Storage migration overview

Storage migration supports the replacement of existing storage enclosures with new enclosures that may offer higher performance, higher capacity, and a smaller footprint.

After new enclosures are installed, you can migrate the data from the older enclosures to the new enclosures while the system continues to support other processes such as data access, expansion, cleaning, and replication. The storage migration does require system resources, but you can control this with throttle settings that give the migration a relatively higher or lower priority. You can also suspend a migration to make more resources available to other processes, then resume the migration when resource demand is lower.

During the migration, the system uses data on the source and destination enclosures. New data is written to the new enclosures. Non-migrated data is updated on the source enclosures, and migrated data is updated on the destination enclosures. If the migration is interrupted, the migration can resume migrating blocks that have not been marked as migrated.

During the migration, each block of data is copied and verified, the source block is freed and marked as migrated, and the system index is updated to use the new location. New data that was destined to land in the source block will now be redirected to destination block. All new data block allocations that would have been allocated from source are allocated from the destination.

The Migration copy process is done at the shelf level, not the logical data level, so all disk sectors on the source shelf are accessed and copied over regardless of whether there is data on them. Therefore, the Storage Migration Utility cannot be used to shrink a logical data footprint.

Note

Because the data set is divided between the source and destination enclosures during migration, you cannot halt a migration and resume use of only the source enclosures. Once started, the migration must complete. If a failure, such as a faulty disk drive, interrupts the migration, address the issue and resume the migration.

Depending on the amount of data to migrate and the throttle settings selected, a storage migration can take days or weeks. When all data is migrated, the finalize process, which must be manually initiated using the `storage migration finalize` command, restarts the filesystem. During the restart, the source enclosures are removed from the system configuration and the destination enclosures become part of the filesystem. When the finalize process is complete, the source enclosures can be removed from the system.

After a storage migration, the disk shelf numbers reported by DD OS might not be sequential. This is because shelf numbering is tied to the serial number of each individual disk shelf. KB article 499019, *Data Domain: Storage enclosure numbering is not sequential*, available on <https://support.emc.com>, provides additional details. In DD OS version 5.7.3.0 and later, the `enclosure show persistent-id` command described in the KB article requires administrator access, not SE access.

Migration planning considerations

Consider the following guidelines before starting a storage migration.

- Storage migration requires a single-use license and operates on system models supported by DD OS version 5.7 or later.

Note

Multiple storage migration operations require multiple licenses. However, multiple source enclosures can be migrated to multiple destination enclosures during a single operation.

-
- Storage migration is based on capacity, not enclosure count. Therefore:
 - One source enclosure can be migrated to one destination enclosure.
 - One source enclosure can be migrated to multiple destination enclosures.
 - Multiple source enclosures can be migrated to one destination enclosure.
 - Multiple source enclosures can be migrated to multiple destination enclosures.
 - The destination enclosures must:
 - Be new, unassigned and unlicensed shelves.
 - Be supported on the DD system model.
 - Contain at least as much usable capacity as the enclosures they are replacing.

Note

It is not possible to determine the utilization of the source shelf. The Data Domain system performs all calculations based on the capacity of the shelf.

-
- The DD system model must have sufficient memory to support the active tier storage capacity of the new enclosures.
 - Data migration is not supported for disks in the system controller.

-  **CAUTION**

Do not upgrade DD OS until the in-progress storage migration is complete.

-
- Storage migration cannot start when the file system is disabled or while a DD OS upgrade is in progress, another migration is in progress, or a RAID reconstruction is in progress.

Note

If a storage migration is in progress, a new storage migration license is required to start a new storage migration operation after the in-progress migration completes. The presence or absence of a storage migration license is reported as part of the upgrade precheck.

-
- All specified source enclosures must be in the same tier (active or archive).
 - There can be only one disk group in each source enclosure, and all disks in the disk group must be installed in within the same enclosure.
 - All disks in each destination enclosure must be of the same type (for example, all SATA or all SAS).
 - After migration begins, the destination enclosures cannot be removed.
 - Source enclosures cannot be removed until migration is complete and finalized.
 - The storage migration duration depends on the system resources (which differ for different system models), the availability of system resources, and the data quantity to migrate. Storage migration can take days or weeks to complete.

DS60 shelf considerations

The DS60 dense shelf can hold 60 disks, allowing the customer to use the full amount of space in the rack. The drives are accessed from the top of the shelf, by extending the shelf from the cabinet. Due to the weight of the shelves, approximately 225 lbs when fully loaded, read this section before proceeding with a storage migration to DS60 shelves.

Be aware of the following considerations when working with the DS60 shelf:

CAUTION

- **Loading shelves at the top of the rack may cause the shelf to tip over.**
 - **Validate that the floor can support the total weight of the DS60 shelves.**
 - **Validate that the racks can provide enough power to the DS60 shelves.**
 - **When adding more than five DS60s in the first rack, or more than six DS60s in the second rack, stabilizer bars and a ladder are required to maintain the DS60 shelves.**
-

Viewing migration status

DD System Manager provides two ways to view storage migration status.

Procedure

1. Select **Hardware > Storage**.

In the Storage area, review the Storage Migration Status line. If the status is Not Licensed, you must add a license before using any storage migration features. If the storage migration license is installed, the status can be one of the following: None, Starting, Migrating, Paused by User, Paused by System, Copy Completed - Pending Finalization, Finalizing, Failed during Copy, or Failed during Finalize.

2. If a storage migration is in progress, click **View Storage Migration** to view the progress dialogs.
-

Note

The migration status shows the percentage of blocks transferred. In a system with many free blocks, the free blocks are not migrated, but they are included in the progress indication. In this situation, the progress indication will climb quickly and then slow when the data migration starts.

3. When a storage migration is in progress, you can also view the status by selecting **Health > Jobs**.

Evaluating migration readiness

You can use the system to evaluate storage migration readiness without committing to start the migration.

Procedure

1. Install the destination enclosures using the instructions in the product installation guides.
2. Select **Administration** > **Licenses** and verify that the storage migration license is installed.
3. If the storage migration license is not installed, click **Add Licenses** and add the license.
4. Select **Hardware** > **Storage**, then click **Migrate Data**.
5. In the Select a Task dialog, select **Estimate**, then click **Next**.
6. In the Select Existing Enclosures dialog, use the checkboxes to select each of the source enclosures for the storage migration, then click **Next**.
7. In the Select New Enclosures dialog, use the checkboxes to select each of the destination enclosures for the storage migration, then click **Next**.

The Add Licenses button allows you to add storage licenses for the new enclosures as needed, without interrupting the current task.

8. In the Review Migration Plan dialog, review the estimated migration schedule, then click **Next**.
9. Review the precheck results in the Verify Migration Preconditions dialog, then click Close.

Results

If any of the precheck tests fail, resolve the issue before you start the migration.

Migrating storage using DD System Manager

The storage migration process evaluates system readiness, prompts you to confirm that you want to start the migration, migrates the data, and then prompts you to finalize the process.

Procedure

1. Install the destination enclosures using the instructions in the product installation guides.
2. Select **Administration** > **Licenses** and verify that the storage migration license is installed.
3. If the storage migration license is not installed, click **Add Licenses** and add the license.
4. Select **Hardware** > **Storage**, then click **Migrate Data**.
5. In the Select a Task dialog, select **Migrate**, then click **Next**.
6. In the Select Existing Enclosures dialog, use the checkboxes to select each of the source enclosures for the storage migration, then click **Next**.

7. In the Select New Enclosures dialog, use the checkboxes to select each of the destination enclosures for the storage migration, then click **Next**.

The Add Licenses button allows you to add storage licenses for the new enclosures as needed, without interrupting the current task.

8. In the Review Migration Plan dialog, review the estimated migration schedule, then click **Start**.

9. In the Start Migration dialog, click **Start**.

The Migrate dialog appears and updates during the three phases of the migration: Starting Migration, Migration in Progress, and Copy Complete.

10. When the Migrate dialog title displays Copy Complete and a filesystem restart is acceptable, click **Finalize**.

Note

This task restarts the filesystem and typically takes 10 to 15 minutes. The system is unavailable during this time.

Results

When the migration finalize task is complete, the system is using the destination enclosures and the source enclosures can be removed.

Storage migration dialog descriptions

The DD System Manager dialog descriptions provide additional information on storage migration. This information is also available by clicking the help icon in the dialogs.

Select a Task dialog

The configuration in this dialog determines whether the system will evaluate storage migration readiness and stop, or evaluate readiness and begin storage migration.

Select **Estimate** to evaluate system readiness and stop.

Select **Migrate** to start migration after the system evaluation. Between the system evaluation and the start of the migration, a dialog prompts you to confirm or cancel the storage migration.

Select Existing Enclosures dialog

The configuration in this dialog selects either the active or the retention tier and the source enclosures for the migration.

If the DD Extended Retention feature is installed, use the list box to select either the **Active Tier** or **Retention Tier**. The list box does not appear when DD Extended Retention is not installed.

The Existing Enclosures list displays the enclosures that are eligible for storage migration. Select the checkbox for each of the enclosures to migrate. Click **Next** when you are ready to continue.

Select New Enclosures dialog

The configuration in this dialog selects the destination enclosures for the migration. This dialog also displays the storage license status and an **Add Licenses** button.

The Available Enclosures list displays the enclosures that are eligible destinations for storage migration. Select the checkbox for each of the desired destination enclosures.

The license status bar represents all of the storage licenses installed on the system. The green portion represents licenses that are in use, and the and clear portion represents the licensed storage capacity available for destination enclosures. If you need to install additional licenses to support the selected destination controllers, click **Add Licenses**.

Click **Next** when you are ready to continue.

Review Migration Plan dialog

This dialog presents an estimate of the storage migration duration, organized according to the three stages of storage migration.

Stage 1 of the storage migration runs a series of tests to verify that the system is ready for the migration. The test results appear in the Verify Migration Preconditions dialog.

During Stage 2, the data is copied from the source enclosures to the destination enclosures. When a large amount of data is present, the copy can take days or weeks to complete because the copy takes place in the background, while the system continues to serve backup clients. A setting in the Migration in Progress dialog allows you to change the migration priority, which can speed up or slow down the migration.

Stage 3, which is manually initiated from the Copy Complete dialog, updates the system configuration to use the destination enclosures and removes the configuration for the source controllers. During this stage, the file system is restarted and the system is unavailable to backup clients.

Verify Migration Preconditions dialog

This dialog displays the results of the tests that execute before the migration starts.

The following list shows the test sequence and provides additional information on each of the tests.

P1. This system's platform is supported.

Older DD system models do not support storage migration.

P2. A storage migration license is available.

A storage migration license is required.

P3. No other migration is currently running.

A previous storage migration must complete before you can start another.

P4. The current migration request is the same as the interrupted migration request.

Resume and complete the interrupted migration.

P5. Check the disk group layout on the existing enclosures.

Storage migration requires that each source enclosure contain only one disk group, and all the disks in the group must be in that enclosure.

P6. Verify the final system capacity.

The total system capacity after migration and the removal of the source enclosures must not exceed the capacity supported by the DD system model.

P7. Verify the replacement enclosures' capacity.

The usable capacity of the destination enclosures must be greater than that of the source enclosures.

P8. Source enclosures are in the same active tier or retention unit.

The system supports storage migration from either the active tier or the retention tier. It does not support migration of data from both tiers at the same time.

P9. Source enclosures are not part of the head unit.

Although the system controller is listed as an enclosure in the CLI, storage migration does not support migration from disks installed in the system controller.

P10. Replacement enclosures are addable to storage.

All disks in each destination enclosure must be of the same type (for example, all SATA or all SAS).

P11. No RAID reconstruction is occurring in the source controllers.

Storage migration cannot start while a RAID reconstruction is in progress.

P12. Source shelf belongs to a supported tier.

The source disk enclosure must be part of a tier supported on the migration destination.

Migration progress dialogs

This series of dialogs presents the storage migration status and the controls that apply at each stage.

Migrate - Starting Migration

During the first stage, the progress is shown on the progress bar and no controls are available.

Migrate - Migration in Progress

During the second stage, data is copied from the source enclosures to the destination enclosures and the progress is shown on the progress bar. Because the data copy can take days or weeks to complete, controls are provided so that you can manage the resources used during migration and suspend migration when resources are needed for other processes.

You can click **Pause** to suspend the migration and later click **Resume** to continue the migration.

The **Low**, **Medium**, and **High** buttons define throttle settings for storage migration resource demands. A low throttle setting gives storage migration a lower resource priority, which results in a slower migration and requires fewer system resources. Conversely, A high throttle setting gives storage migration a higher resource priority, which results in a faster migration and requires more system resources. The medium setting selects an intermediate priority.

You do not have to leave this dialog open for the duration of the migration. To check the status of the migration after closing this dialog, select **Hardware > Storage** and view the migration status. To return to this dialog from the Hardware/Storage page, click **Manage Migration**. The migration progress can also be viewed by selecting **Health > Jobs**.

Migrate - Copy Complete

When the copy is complete, the migration process waits for you to click **Finalize**. During this final stage, , which takes 10 to 15 minutes, the filesystem is restarted and the system is not available. It is a good practice to start this stage during a maintenance window or a period of low system activity.

Migrating storage using the CLI

A migration simply requires moving all of the allocated blocks from the blocksets formatted over source DGs (e.g., source blocksets) to the blocksets formatted over destination DGs (e.g., destination blocksets). Once all of the allocated blocks have been moved from the source blocksets, those blocksets can be removed from the file system, their disks can be removed from their storage tier, and the physical disks and enclosures can be removed from the DDR.

Note

The preparation of new enclosures for storage migration is managed by the storage migration process. Do not prepare destination enclosures as you would for an enclosure addition. For example, use of the `filesys expand` command is appropriate for an enclosure addition, but this command prevents enclosures from being used as storage migration destinations.

A DS60 disk shelf contains four disk packs, of 15 disks each. When a DS60 shelf is the migration source or destination, the disk packs are referenced as `enclosure:pack`. In this example, the source is enclosure 7, pack 2 (7:2), and the destination is enclosure 7, pack 4 (7:4).

Procedure

1. Install the destination enclosures using the instructions in the product installation guides.
2. Check to see if the storage migration feature license is installed.

```
# elicence show
```

3. If the license is not installed, update the elicence to add the storage migration feature license

```
# elicence update
```

4. View the disk states for the source and destination disks.

```
# disk show state
```

The source disks should be in the active state, and the destination disks should be in the unknown state.

5. Run the storage migration precheck command to determine if the system is ready for the migration.

```
# storage migration precheck source-enclosures 7:2 destination-enclosures 7:4
```

6. View the migration throttle setting.

```
storage migration option show throttle
```

7. When the system is ready, begin the storage migration.

```
# storage migration start source-enclosures 7:2 destination-enclosures 7:4
```

8. Optionally, view the disk states for the source and destination disks during the migration.

```
# disk show state
```

During the migration, the source disks should be in the migrating state, and the destination disks should be in the destination state.

9. Review the migration status as needed.

```
# storage migration status
```

10. View the disk states for the source and destination disks.

```
# disk show state
```

During the migration, the source disks should be in the migrating state, and the destination disks should be in the destination state.

11. When the migration is complete, update the configuration to use the destination enclosures.

Note

This task restarts the file system and typically takes 10 to 15 minutes. The system is unavailable during this time.

```
storage migration finalize
```

12. If you want to remove all data from each of the source enclosures, remove the data now.

```
storage sanitize start enclosure <enclosure-id>[:<pack-id>]
```

Note

The storage sanitize command does not produce a certified data erasure. Data Domain offers certified data erasure as a service. For more information, contact your Data Domain representative.

13. View the disk states for the source and destination disks.

```
# disk show state
```

After the migration, the source disks should be in the unknown state, and the destination disks should be in the active state.

Results

When the migration finalize task is complete, the system is using the destination storage and the source storage can be removed.

CLI storage migration example

elicense show

```
# elicense show
Feature licenses:
## Feature      Count Mode                Expiration Date
-----
1  REPLICATION  1    permanent (int)  n/a
2  VTL          1    permanent (int)  n/a
-----
```

elicense update

```
# elicense update mylicense.lic
New licenses: Storage Migration
Feature licenses:
## Feature          Count    Mode          Expiration Date
-----
1  REPLICATION      1      permanent (int) n/a
2  VTL               1      permanent (int) n/a
3  Storage Migration 1      permanent (int)
-----
** This will replace all existing Data Domain licenses on the system with the above EMC ELMS
licenses.
Do you want to proceed? (yes|no) [yes]: yes
eLicense(s) updated.
```

disk show state

```
# disk show state
Enclosure      Disk
Row(disk-id)  1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
-----
1              .  .  .  .
2              U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
3              U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
4              U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
5              v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
6              U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
7
              Pack 1  Pack 2  Pack 3  Pack 4
E(49-60)      U  U  U  .  .  .  U  U  U  U  U  U
D(37-48)      U  U  U  .  .  .  U  U  U  U  U  U
C(25-36)      U  U  U  .  .  .  U  U  U  U  U  U
B(13-24)      U  U  U  .  .  .  U  U  U  U  U  U
A( 1-12)      U  U  U  .  .  .  U  U  U  U  U  U
-----

Legend  State          Count
-----
.       In Use Disks   18
s       Spare Disks    1
v       Available Disks 15
U       Unknown Disks 105
-----
```

storage migration precheck

```
#storage migration precheck  source-enclosures 2  destination-enclosures 11

Source enclosures:
Disks      Count    Disk      Disk      Enclosure  Enclosure
          Group    Size      Model      Serial No.
-----
2.1-2.15  15      dg1      1.81 TiB  ES30      APM00111103820
-----
Total source disk size: 27.29 TiB

Destination enclosures:
Disks      Count    Disk      Disk      Enclosure  Enclosure
          Group    Size      Model      Serial No.
-----
11.1-11.15  15      unknown  931.51 GiB ES30      APM00111103840
-----
Total destination disk size: 13.64 TiB

1 "Verifying platform support.....PASS"
2 "Verifying valid storage migration license exists.....PASS"
3 "Verifying no other migration is running.....PASS"
4 "Verifying request matches interrupted migration.....PASS"
5 "Verifying data layout on the source shelves.....PASS"
6 "Verifying final system capacity.....PASS"
7 "Verifying destination capacity.....PASS"
8 "Verifying source shelves belong to same tier.....PASS"
```

Storage Migration

```
9 "Verifying enclosure 1 is not used as source.....PASS"
10 "Verifying destination shelves are addable to storage.....PASS"
11 "Verifying no RAID reconstruction is going on in source shelves.....PASS"
Migration pre-check PASSED
```

Expected time to migrate data: 8 hrs 33 min

storage migration show history

```
# storage migration show history
```

Id	Source Enclosure*	Source Enclosure Serial No.	Dest Enclosure*	Dest Enclosure Serial No.	Status	Start Time	End Time
2	9:0	SHU952400106A23	7:0	SHU95240840055B	Finalized	Sat Aug 8 11:59:37 2015	Mon Aug 10 11:10:11 2015
1	9:0	SHU952400106A23	7:0 8:0	SHU95240840055B SHU9524084004LR	Finalized	Thu Aug 6 16:39:55 2015	Fri Aug 7 10:28:07 2015

(*) Enclosure ids at migration start time.

storage migration start

```
#storage migration start source-enclosures 2 destination-enclosures 11
```

Source enclosures:

Disks	Count	Disk Group	Disk Size	Enclosure Model	Enclosure Serial No.
2.1-2.15	15	dg1	1.81 TiB	ES30	APM00111103820

Total source disk size: 27.29 TiB

Destination enclosures:

Disks	Count	Disk Group	Disk Size	Enclosure Model	Enclosure Serial No.
11.1-11.15	15	unknown	931.51 GiB	ES30	APM00111103840

Total destination disk size: 13.64 TiB

Expected time to migrate data: 84 hrs 40 min

```
** Storage migration once started cannot be aborted.
Existing data on the destination shelves will be overwritten.
Do you want to continue with the migration? (yes|no) [no]: yes
```

Performing migration pre-check:

```
1 Verifying platform support.....PASS
2 Verifying valid storage migration license exists.....PASS
3 Verifying no other migration is running.....PASS
4 Verifying request matches interrupted migration.....PASS
5 Verifying data layout on the source shelves.....PASS
6 Verifying final system capacity.....PASS
7 Verifying destination capacity.....PASS
8 Verifying source shelves belong to same tier.....PASS
9 Verifying enclosure 1 is not used as source.....PASS
10 Verifying destination shelves are addable to storage.....PASS
11 Verifying no RAID reconstruction is going on in source shelves.....PASS
```

Migration pre-check PASSED

Storage migration will reserve space in the filesystem to migrate data.
Space reservation may add up to an hour or more based on system resources.

Storage migration process initiated.
Check storage migration status to monitor progress.

storage migration status

```
# storage migration status
```

Id	Source Enclosure(s)	Destination Enclosure(s)	State	Percent Complete	Estimated Time to Complete	Current Throttle Setting
5	7:2	7:4	migrating	45%	30 hrs 18 mins	high

disk show state, migration in progress

```
# disk show state
```

Enclosure	Disk														
Row(disk-id)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1
2	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U
3	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U
4	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U
5	v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
6	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U
7	-----														
	Pack 1			Pack 2			Pack 3			Pack 4					
E(49-60)	U	U	U	m	m	s	U	U	U	s	d	d	d	d	d
D(37-48)	U	U	U	m	m	m	U	U	U	d	d	d	d	d	d
C(25-36)	U	U	U	m	m	m	U	U	U	d	d	d	d	d	d
B(13-24)	U	U	U	m	m	m	U	U	U	d	d	d	d	d	d
A(1-12)	U	U	U	m	m	m	U	U	U	d	d	d	d	d	d

Legend	State	Count													
.	In Use Disks	4													
s	Spare Disks	2													
v	Available Disks	15													
U	Unknown Disks	90													
m	Migrating Disks	14													
d	Destination Disks	14													

storage migration finalize

```
# storage migration finalize

Storage migration finalize restarts the filesystem.
This can take several minutes and the filesystem is unavailable until the operation completes.
Do you want to continue? (yes|no) [no]: yes

Performing migration finalization pre-check:
(P1) Verifying storage migration is ready for finalization...PASS
(P2) Verifying there are no foreign disks.....PASS
(P3) Verifying data layout on the source shelves.....PASS

Migration finalization pre-check PASSED
Finalizing the storage migration with id 5:

Notifying filesystem to finalize migration...

Done.

Disabling the filesystem
Please wait.....
The filesystem is now disabled.
Removing source enclosures from filesystem...

Done.

Removing source enclosures from storage tier...

Done.

Enabling the filesystem
Please wait.....
The filesystem is now enabled.
Storage migration with id 5 from enclosure(s) 7.2 to enclosure(s) 7.4 has been finalized.
```

disk show state, migration complete

```
# disk show state
Enclosure      Disk
  Row(disk-id)  1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
-----
1              .  .  .  .
2              U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
3              U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
4              U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
5              v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
6              U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
7
              |-----|
              | Pack 1 | Pack 2 | Pack 3 | Pack 4 |
              |-----|
E(49-60)      | U  U  U | U  U  U | U  U  U | s  .  . |
D(37-48)      | U  U  U | U  U  U | U  U  U | .  .  . |
C(25-36)      | U  U  U | U  U  U | U  U  U | .  .  . |
B(13-24)      | U  U  U | U  U  U | U  U  U | .  .  . |
A( 1-12)      | U  U  U | U  U  U | U  U  U | .  .  . |
              |-----|

Legend  State      Count
-----
.       In Use Disks  18
s       Spare Disks   1
v       Available Disks 15
U       Unknown Disks 105
-----
```

Note

Currently storage migration is only supported on the active node. Storage migration is not supported on the standby node of an HA cluster.

CHAPTER 12

Metadata on Flash

This chapter includes:

- [Overview of Metadata on Flash \(MDoF\)](#) 290
- [MDoF licensing and capacity](#) 290
- [SSD cache tier](#) 291
- [MDoF SSD cache tier - system management](#) 291
- [SSD alerts](#) 295

Overview of Metadata on Flash (MDoF)

MDoF creates caches for file system metadata using flash technologies. The SSD Cache is a low latency, high input/output operations per second (IOPS) cache to accelerate metadata and data access.

Note

The minimum software version required is DD OS 6.0.

Caching the file system metadata on SSDs improves I/O performance for both traditional and random workloads.

For traditional workloads, offloading random access to metadata from HDDs to SSDs allows the hard drives to accommodate streaming write and read requests.

For random workloads, SSD cache provides low latency metadata operations, which allows the HDDs to serve data requests instead of cache requests.

Read cache on SSD improves random read performance by caching frequently accessed data. Writing data to NVRAM combined with low latency metadata operations to drain the NVRAM faster improve random write latency. The absence of cache does not prevent file system operation, it only impacts file system performance.

When the cache tier is first created, a file system restart is only required if the cache tier is being added after the file system is running. For new systems that come with cache tier disks, no file system restart is required if the cache tier is created before enabling the file system for the first time. Additional cache can be added to a live system, without the need to disable and enable the file system.

Note

DD9500 systems that were upgraded from DD OS 5.7 to DD OS 6.0 require a one-time file system restart after creating the cache tier for the first time.

One specific condition with regard to SSDs is when the number of spare blocks remaining gets close to zero, the SSD enters a read only condition. When a read only condition occurs, DD OS treats the drive as read-only cache and sends an alert.

MDoF is supported on the following Data Domain systems:

- DD6300
- DD6800
- DD9300
- DD9500
- DD9800
- DD VE instances, including DD3300 systems, in capacity configurations of 16 TB and higher (SSD Cache Tier for DD VE)

MDoF licensing and capacity

A license enabled through ELMS is necessary for using the MDoF feature; the SSD Cache license will not be enabled by default.

The following table describes the various SSD capacity licenses and the SSD capacities for the given system:

Table 116 SSD capacity licenses per system

Model	Memory	Number of SSDs	SSD capacity
DD6300	48 GB (Base)	1	800 GB
	96 GB (Expanded)	2	1600 GB
DD6800	192 GB (Base)	2	1600 GB
	192 GB (Expanded)	4	3200 GB
DD9300	192 GB (Base)	5	4000 GB
	384 GB (Expanded)	8	6400 GB
DD9500	256 GB (Base)	8	6400 GB
	512 GB (Expanded)	15	12000 GB
DD9800	256 GB (Base)	8	6400 GB
	768 GB (Expanded)	15	12000 GB

SSD Cache Tier for DD VE

DD VE instances and DD3300 systems do not require a license for the SSD Cache Tier. The maximum supported SSD capacity is 1% of the Active Tier capacity.

The following table describes the various SSD capacity licenses and the SSD capacities for the given system:

Table 117 DD VE and DD3300 SSD capacity

Capacity configuration	Maximum SSD capacity
DD VE 16 TB	160 GB
DD VE 32 TB	320 GB
DD VE 48 TB	480 GB
DD VE 64 TB	640 GB
DD VE 96 TB	960 GB
DD3300 16 TB	160 GB
DD3300 32 TB	320 GB

SSD cache tier

The SSD cache tier provides the SSD cache storage for the file system. The file system draws the required storage from the SSD cache tier without active intervention from the user.

MDoF SSD cache tier - system management

Be aware of the following considerations for SSD cache:

- When SSDs are deployed within a controller, those SSDs are treated as internal root drives. They display as enclosure 1 in the output of the `storage show all` command.
- Manage individual SSDs with the `disk` command the same way HDDs are managed.
- Run the `storage add` command to add an individual SSD or SSD enclosure to the SSD cache tier.
- The SSD cache tier space does not need to be managed. The file system draws the required storage from the SSD cache tier and shares it among its clients.
- The `filesys create` command creates an SSD volume if SSDs are available in the system.

Note

If SSDs are added to the system later, the system should automatically create the SSD volume and notify the file system. SSD Cache Manager notifies its registered clients so they can create their cache objects.

- If the SSD volume contains only one active drive, the last drive to go offline will come back online if the active drive is removed from the system.

The next section describes how to manage the SSD cache tier from Data Domain System Manager, and with the DD OS CLI.

Managing the SSD cache tier

Storage configuration features allow you to add and remove storage from the SSD cache tier.

Procedure

1. Select **Hardware > Storage > Overview**.
2. Expand the **Cache Tier** dialog.
3. Click **Configure**.

The maximum amount of storage that can be added to the active tier depends on the DD controller used.

Note

The licensed capacity bar shows the portion of licensed capacity (used and remaining) for the installed enclosures.

4. Select the checkbox for the Shelf to be added.
5. Click the **Add to Tier** button.
6. 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**.

CLI Equivalent

When the cache tier SSDs are installed in the head unit:

a. Add the SSDs to the cache tier.

```
# storage add disks 1.13,1.14 tier cache
Checking storage requirements...done
Adding disk 1.13 to the cache tier...done

Updating system information...done

Disk 1.13 successfully added to the cache tier.

Checking storage requirements...
done
Adding disk 1.14 to the cache tier...done

Updating system information...done

Disk 1.14 successfully added to the cache tier.
```

b. Verify the state of the newly added SSDs.

```
# disk show state
Enclosure  Disk
-----  -----
1          1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
-----  -----
1          .  .  .  .  s  .  .  s  s  s  s  s  v  v
2          U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
3          U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
-----  -----

Legend    State                Count
-----  -----
.         In Use Disks          6
s         Spare Disks          6
v         Available Disks     2
U         Unknown Disks       30
-----  -----

Total 44 disks
```

When the cache tier SSDs are installed in an external shelf:

a. Verify the system recognizes the SSD shelf. In the example below, the SSD shelf is enclosure 2.

```
# disk show state
Enclosure  Disk
Row(disk-id) 1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
-----  -----
1          .  .  .  .
2          U  U  U  U  U  U  U  U  -  -  -  -  -  -  -
3          .  .  .  .  .  .  .  .  .  .  .  .  .  .  v
4          .  .  .  .  .  .  .  .  .  .  .  .  .  .  v
5          v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
6          v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
7          v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
8          v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
9          v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
10         |-----|-----|-----|-----|
           | Pack 1 | Pack 2 | Pack 3 | Pack 4 |
           |-----|-----|-----|-----|
           E(49-60) |v  v  v |v  v  v |v  v  v |v  v  v |
           D(37-48) |v  v  v |v  v  v |v  v  v |v  v  v |
           C(25-36) |v  v  v |v  v  v |v  v  v |v  v  v |
           B(13-24) |v  v  v |v  v  v |v  v  v |v  v  v |
           A( 1-12) |v  v  v |v  v  v |v  v  v |v  v  v |
           |-----|-----|-----|-----|
11         v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
12         v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
13         v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
-----  -----

Legend    State                Count
-----  -----
```

```
.      In Use Disks      32
v      Available Disks  182
U      Unknown Disks    8
-      Not Installed Disks 7
-----
Total 222 disks
```

b. Identify the shelf ID of the SSD shelf. SSDs will display as SAS-SSD or SATA-SSD in the Type column.

```
# disk show hardware
```

Disk (enc/disk)	Slot	Manufacturer/Model	Firmware	Serial No.	Capacity	Type
1.1	0	TG32C10400GA3EMC	118000371	PRO6E344	FG009826	372.61 GiB SATA-SSD
1.2	1	TG32C10400GA3EMC	118000371	PRO6E344	FG0097VL	372.61 GiB SATA-SSD
1.3	2	TG32C10400GA3EMC	118000371	PRO6E344	FG009881	372.61 GiB SATA-SSD
1.4	3	TG32C10400GA3EMC	118000371	PRO6E344	FG00988X	372.61 GiB SATA-SSD
2.1	0	HITACHI HUSMR148_CLAR800	C29C	07V4P2AA	745.22 GiB	SAS-SSD
2.2	1	HITACHI HUSMR148_CLAR800	C29C	07V4P3LA	745.22 GiB	SAS-SSD
2.3	2	HITACHI HUSMR148_CLAR800	C29C	07V4P2XA	745.22 GiB	SAS-SSD
2.4	3	HITACHI HUSMR148_CLAR800	C29C	07V4TW4A	745.22 GiB	SAS-SSD
2.5	4	HITACHI HUSMR148_CLAR800	C29C	07V4ULYA	745.22 GiB	SAS-SSD
2.6	5	HITACHI HUSMR148_CLAR800	C29C	07V4P0BA	745.22 GiB	SAS-SSD
2.7	6	HITACHI HUSMR148_CLAR800	C29C	07V4UVBA	745.22 GiB	SAS-SSD
2.8	7	HITACHI HUSMR148_CLAR800	C29C	07V4UTNA	745.22 GiB	SAS-SSD

c. Add the SSD shelf to the cache tier

```
# storage add enclosure 2 tier cache
```

```
sysadmin@apolloplus-1# storage add enclosure 2 tier cache
Checking storage requirements...done
Adding enclosure 2 to the cache tier...Enclosure 2 successfully added to the cache tier.
Updating system information...done
Successfully added: 2 done
```

d. Verify the state of the newly added SSDs.

```
# disk show state
```

Enclosure	Disk	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Row(disk-id)		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	
2	
3		v
4		v
5		v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
6		v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
7		v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
8		v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
9		v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
10		-----			-----			-----			-----					
		Pack 1			Pack 2			Pack 3			Pack 4					
	E (49-60)	v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
	D (37-48)	v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
	C (25-36)	v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
	B (13-24)	v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
	A (1-12)	v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
		-----			-----			-----			-----					
11		v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
12		v	v	v	v	v	v	v	v	v	v	v	v	v	v	v
13		v	v	v	v	v	v	v	v	v	v	v	v	v	v	v

```
Legend  State      Count
-----  -
.      In Use Disks      32
v      Available Disks  182
U      Unknown Disks    8
-      Not Installed Disks 7
```

```
-----
Total 222 disks
```

To remove a controller-mounted SSD from the cache tier:

```
# storage remove disk 1.13
```

```
Removing disk 1.13...done
```

```
Updating system information...done
```

```
Disk 1.13 successfully removed.
```

To remove an SSD shelf from the system:

```
# storage remove enclosure 2
```

```
Removing enclosure 2...Enclosure 2 successfully removed.
```

```
Updating system information...done
```

```
Successfully removed: 2 done
```

SSD alerts

There are three alerts specific to the SSD cache tier.

The SSD cache tier alerts are:

- **Licensing**
If the file system is enabled and less physical cache capacity present than what the license permits is configured, an alert is generated with the current SSD capacity present, and the capacity license. This alert is classified as a warning alert. The absence of cache does not prevent file system operation, it only impacts file system performance. Additional cache can be added to a live system, without the need to disable and enable the file system.
- **Read only condition**
When the number of spare blocks remaining gets close to zero, the SSD enters a read only condition. When a read only condition occurs, DD OS treats the drive as read-only cache.
Alert `EVT-STORAGE-00001` displays when the SSD is in a read-only state and should be replaced.
- **SSD end of life**
When an SSD reaches the end of its lifespan, the system generates a hardware failure alert identifying the location of the SSD within the SSD shelf. This alert is classified as a critical alert.
Alert `EVT-STORAGE-00016` displays when the EOL counter reaches 98. The drive is failed proactively when the EOL counter reaches 99.

CHAPTER 13

SCSI Target

This chapter includes:

- [SCSI Target overview](#)..... 298
- [Fibre Channel view](#).....299
- [Differences in FC link monitoring among DD OS versions](#)..... 309

SCSI Target overview

SCSI (Small Computer System Interface) Target is a unified management daemon for all SCSI services and transports. SCSI Target supports DD VTL (Virtual Tape Library), DD Boost over FC (Fibre Channel), and vDisk/ProtectPoint Block Services, as well as anything that has a target LUN (logical unit number) on a DD system.

SCSI Target Services and Transports

The SCSI Target daemon starts when FC ports are present or DD VTL is licensed. It provides unified management for all SCSI Target *services* and *transports*.

- A *service* is anything that has a target LUN on a DD system that uses SCSI Target commands, such as DD VTL (tape drives and changers), DD Boost over FC (processor devices), or vDisk (Virtual Disk Device).
- 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, vDisk, or DD VTL, but not all three.
- *Devices* are visible on a SAN (storage area network) through physical ports. Host initiators communicate with the DD system through the SAN.
- *Access groups* manage access between devices and initiators.
- An *endpoint* is the logical target on a DD system to which an initiator connects. You can disable, enable, and rename endpoints. To delete endpoints, the associated transport hardware must no longer exist. Endpoints are automatically discovered and created when a new transport connection occurs. Endpoints have the following attributes: port topology, FCP2-RETRY status, WWPN, and WWNN.
- *NPIV* (N_port ID Virtualization) is an FC feature that lets multiple endpoints share a single physical port. NPIV eases hardware requirements and provides failover capabilities.
- In DD OS 6.0, users can specify the sequence of secondary system addresses for failover. For example, if the system specifies 0a, 0b, 1a, 1b and the user specifies 1b, 1a, 0a, 0b, the user-specified sequence is used for failover. The `scsitarget endpoint show detailed` command displays the user-specified sequence.

Note the following exceptions:

- DD Boost can service both FC and IP clients simultaneously; however, both transports cannot share the same initiator.
- Only one initiator should be present per access group. Each access group is assigned a type (DD VTL, vDisk/ProtectPoint Block Services, or DD Boost over FC).

SCSI Target Architectures - Supported and Unsupported

SCSI Target supports the following architectures:

- **DD VTL plus DD Boost over FC from different initiators:** Two different initiators (on the same or different clients) may access a DD system using DD VTL and DD Boost over FC, through the same or different DD system target endpoints.
- **DD VTL plus DD Boost over FC from one initiator to two different DD systems:** A single initiator may access two different DD systems using any service.

SCSI Target does not support the following architecture:

- **DD VTL plus DD Boost over FC from one initiator to the same DD system:** A single initiator may not access the same DD system through different services.

Thin Protocol

The thin protocol is a lightweight daemon for VDisk and DD VTL that responds to SCSI commands when the primary protocol can't. For Fibre Channel environments with multiple protocols, thin protocol:

- Prevents initiator hangs
- Prevents unnecessary initiator aborts
- Prevents initiator devices from disappearing
- Supports a standby mode
- Supports fast and early discoverable devices
- Enhances protocol HA behavior
- Doesn't require fast registry access

For More Information about DD Boost and the `scscitarget` Command (CLI)

For more information about using DD Boost through the DD System Manager, see the related chapter in this book. For other types of information about DD Boost, see the *Data Domain Boost for OpenStorage Administration Guide*.

This chapter focuses on using SCSI Target through the DD System Manager. After you have become familiar with basic tasks, the `scscitarget` command in the *Data Domain Operating System Command Reference Guide* provides more advanced management tasks.

When there is heavy DD VTL traffic, avoid running the `scsitarget group use` command, which switches the in-use endpoint lists for one or more SCSI Target or vdisk devices in a group between primary and secondary endpoint lists.

Fibre Channel view

The Fibre Channel view displays the current status of whether Fibre Channel and/or NPIV is enabled. It also displays two tabs: Resources and Access Groups. Resources include ports, endpoints, and initiators. An access group holds a collection of initiator WWPNs (worldwide port names) or aliases and the drives and changers they are allowed to access.

Enabling NPIV

NPIV (N_Port ID Virtualization), is a Fibre Channel feature in which multiple endpoints can share a single physical port. NPIV eases hardware requirements and provides endpoint failover/failback capabilities. NPIV is not configured by default; you must enable it.

Note

NPIV is enabled by default in HA configuration.

NPIV provides simplified multiple-system consolidation:

- NPIV is an ANSI T11 standard that allows a single HBA physical port to register with a Fibre Channel fabric using multiple WWPNs
- The virtual and physical ports have the same port properties and behave exactly the same.

- There may be m:1 relationships between the endpoints and the port, that is, multiple endpoints can share the same physical port.

Specifically, enabling NPIV enables the following features:

- Multiple endpoints are allowed per physical port, each using a virtual (NPIV) port. The base port is a placeholder for the physical port and is not associated with an endpoint.
- Endpoint failover/failback is automatically enabled when using NPIV.

Note

After NPIV is enabled, the "Secondary System Address" must be specified at each of the endpoints. If not, endpoint failover will not occur.

- Multiple DD systems can be consolidated into a single DD system, however, the number of HBAs remains the same on the single DD system.
- The endpoint failover is triggered when FC-SSM detects when a port goes from online to offline. In the case where the physical port is offline before scsitarget is enabled and the port is still offline after scsitarget is enabled, a endpoint failover is not possible because FC-SSM does not generate a port offline event. If the port comes back online and auto-failback is enabled, any failed over endpoints that use that port as a primary port will fail-back to the primary port.

The Data Domain HA features requires NPIV to move WWNs between the nodes of an HA pair during the failover process.

Note

Before enabling NPIV, the following conditions must be met:

- The DD system must be running DD OS 5.7.
- All ports must be connected to 4Gb, 8Gb, and 16 Gb Fibre Channel HBA and SLIC.
- The DD system ID must be valid, that is, it must not be 0.

In addition, port topologies and port names will be reviewed and may prevent NPIV from being enabled:

- NPIV is allowed if the topology for *all* ports is loop-preferred.
- NPIV is allowed if the topology for *some* of the ports is loop-preferred; however, NPIV must be disabled for ports that are loop-only, or you must reconfigure the topology to loop-preferred for proper functionality.
- NPIV is *not* allowed if *none* of the ports has a topology of loop-preferred.
- If port names are present in access groups, the port names are replaced with their associated endpoint names.

Procedure

1. Select **Hardware > Fibre Channel**.
2. Next to NPIV: Disabled, select **Enable**.
3. In the Enable NPIV dialog, you will be warned that all Fibre Channel ports must be disabled before NPIV can be enabled. If you are sure that you want to do this, select **Yes**.

CLI Equivalent

- a. Make sure (global) NPIV is enabled.

```
# scsitarget transport option show npiv
SCSI Target Transport Options
Option      Value
-----
npiv        disabled
-----
```

- b. If NPIV is disabled, then enable it. You must first disable all ports.

```
# scsitarget port disable all
All ports successfully disabled.
# scsitarget transport option set npiv enabled
Enabling FiberChannel NPIV mode may require SAN zoning to
be changed to configure both base port and NPIV WWPNs.
Any FiberChannel port names used in the access groups will
be converted to their corresponding endpoint names in order
to prevent ambiguity.
Do you want to continue? (yes|no) [no]:
```

- c. Re-enable the disabled ports.

```
# scsitarget port enable all
All ports successfully enabled.
```

- d. Make sure the physical ports have an NPIV setting of “auto”.

```
# scsitarget port show detailed 0a
System Address:      0a
Enabled:             Yes
Status:              Online
Transport:           FibreChannel
Operational Status: Normal
FC NPIV:             Enabled (auto)
.
.
.
```

- e. Create a new endpoint using the primary and secondary ports you have selected.

```
# scsitarget endpoint add test0a0b system-address 0a
primary-system-address 0a secondary-system-address 0b
```

Note that the endpoint is disabled by default, so enable it.

```
# scsitarget endpoint enable test0a0b
```

Then display the endpoint information.

```
# scsitarget endpoint show detailed test0a0b
Endpoint:            test0a0b
Current System Address: 0b
Primary System Address: 0a
Secondary System Address: 0b
Enabled:             Yes
Status:              Online
Transport:           FibreChannel
FC WWNN:             50:02:18:80:08:a0:00:91
FC WWPN:             50:02:18:84:08:b6:00:91
```

- f. Zone a host system to the auto-generated WWPN of the newly created endpoint.
- g. Create a DD VTL, vDisk, or DD Boost over Fibre Channel (DFC) device, and make this device available on the host system.
- h. Ensure that the DD device chosen can be accessed on the host (read and/or written).
- i. Test the endpoint failover by using the “secondary” option to move the endpoint to the SSA (secondary system address).

```
# scsitaraget endpoint use test0a0b secondary
```

- j. Ensure that the DD device chosen can still be accessed on the host (read and/or written). Test the failback by using the “primary” option to move the endpoint back to the PSA (primary system address).

```
# scsitaraget endpoint use test0a0b primary
```

- k. Ensure that the DD device chosen can still be accessed on the host (read and/or written).

Disabling NPIV

Before you can disable NPIV, you must not have any ports with multiple endpoints.

Note

NPIV is required for HA configuration. It is enabled by default and cannot be disabled.

Procedure

1. Select **Hardware > Fibre Channel**.
2. Next to NPIV: Enabled, select **Disable**.
3. In the Disable NPIV dialog, review any messages about correcting the configuration, and when ready, select **OK**.

Resources tab

The **Hardware > Fibre Channel > Resources** tab displays information about ports, endpoints, and initiators.

Table 118 Ports

Item	Description
System Address	System address for port
WWPN	Unique worldwide 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 (FC) port.
WWNN	Unique worldwide 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 FC node
Enabled	Port operational status; either Enabled or Disabled.
NPIV	NPIV status; either Enabled or Disabled.
Link Status	Link status: either Online or Offline; that is, whether or not the port is up and capable of handling traffic.
Operation Status	Operation status: either Normal or Marginal.
# of Endpoints	Number of endpoints associated with this port.

Table 119 Endpoints

Item	Description
Name	Name of endpoint.
WWPN	Unique worldwide 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 (FC) port.
WWNN	Unique worldwide 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 FC node
System Address	System address of endpoint.
Enabled	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.

Table 120 Initiators

Item	Description
Name	Name of initiator.
Service	Service support by the initiator, which is either DD VTL, DD Boost, or vDisk.
WWPN	Unique worldwide 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 (FC) port.
WWNN	Unique worldwide 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 FC node.
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.

Configuring a port

Ports are discovered, and a single endpoint is automatically created for each port, at startup.

The properties of the base port depend on whether NPIV is enabled:

- In non-NPIV mode, ports use the same properties as the endpoint, that is, the WWPN for the base port and the endpoint are the same.
- In NPIV mode, the base port properties are derived from default values, that is, a new WWPN is generated for the base port and is preserved to allow consistent switching between NPIV modes. Also, NPIV mode provides the ability to support multiple endpoints per port.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Under **Ports**, select an port, and then select **Modify** (pencil).

3. In the Configure Port dialog, select whether to automatically enable or disable NPIV for this port.
4. For Topology, select Loop Preferred, Loop Only, Point to Point, or Default.
5. For Speed, select 1, 2, 4 or 8 Gbps, or auto.
6. Select **OK**.

Enabling a port

Ports must be enabled before they can be used.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Select **More Tasks > Ports > Enable**. If all ports are already enabled, a message to that effect is displayed.
3. In the Enable Ports dialog, select one or more ports from the list, and select **Next**.
4. After the confirmation, select **Next** to complete the task.

Disabling a port

You can simply disable a port (or ports), or you can chose to failover all endpoints on the port (or ports) to another port.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Select **More Tasks > Ports > Disable**.
3. In the Disable Ports dialog, select one or more ports from the list, and select **Next**.
4. In the confirmation dialog, you can continue with simply disabling the port, or you can chose to failover all endpoints on the ports to another port.

Adding an endpoint

An endpoint is a virtual object that is mapped to a underlying virtual port. In non-NPIV mode (not available on HA configuration), only a single endpoint is allowed per physical port, and the base port is used to configure that endpoint to the fabric. When NPIV is enabled, multiple endpoints are allowed per physical port, each using a virtual (NPIV) port, and endpoint failover/failback is enabled.

Note

Non-NPIV mode is not available on HA configurations. NPIV is enabled by default and cannot be disabled.

Note

In NPIV mode, endpoints:

- have a primary system address.
 - may have zero or more secondary system addresses.
 - are all candidates for failover to an alternate system address on failure of a port; however, failover to a marginal port is not supported.
 - may be failed back to use their primary port when the port comes back up online.
-

Note

When using NPIV, it is recommended that you use only one protocol (that is, DD VTL Fibre Channel, DD Boost-over-Fibre Channel, or vDisk Fibre Channel) per endpoint. For failover configurations, secondary endpoints should also be configured to have the same protocol as the primary.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Under **Endpoints**, select **Add (+ sign)**.
3. In the Add Endpoint dialog, enter a Name for the endpoint (from 1 to 128 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 [(,)].
4. For Endpoint Status, select Enabled or Disabled.
5. If NPIV is enabled, for Primary system address, select from the drop-down list. The primary system address must be different from any secondary system address.
6. If NPIV is enabled, for Fails over to secondary system addresses, check the appropriate box next to the secondary system address.
7. Select **OK**.

Configuring an endpoint

After you have added an endpoint, you can modify it using the Configure Endpoint dialog.

Note

When using NPIV, it is recommended that you use only one protocol (that is, DD VTL Fibre Channel, DD Boost-over-Fibre Channel, or vDisk Fibre Channel) per endpoint. For failover configurations, secondary endpoints should also be configured to have the same protocol as the primary.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Under **Endpoints**, select an endpoint, and then select **Modify** (pencil).
3. In the Configure Endpoint dialog, enter a Name for the endpoint (from 1 to 128 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 [(,)].

4. For Endpoint Status, select Enabled or Disabled.
5. For Primary system address, select from the drop-down list. The primary system address must be different from any secondary system address.
6. For Fails over to secondary system addresses, check the appropriate box next to the secondary system address.
7. Select **OK**.

Modifying an endpoint's system address

You can modify the active system address for a SCSI Target endpoint using the `scsitarget endpoint modify` command option. This is 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 example, endpoint `ep-1` was assigned to system address `5a`, but this system address is no longer valid. A new controller HBA was added at system address `10a`. The SCSI Target subsystem automatically created a new endpoint, `ep-new`, for the newly discovered system address. Because only a single endpoint can be associated with a given system address, `ep-new` must be deleted, and then `ep-1` must be assigned to system address `10a`.

Note

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. You may also need to update SAN zoning to reflect the new 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
```

Enabling an endpoint

Enabling an endpoint enables the port only if it is currently disabled, that is, you are in non-NPIV mode.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.

2. Select **More Tasks > Endpoints > Enable**. If all endpoints are already enabled, a message to that effect is displayed.
3. In the Enable Endpoints dialog, select one or more endpoints from the list, and select **Next**.
4. After the confirmation, select **Next** to complete the task.

Disabling an endpoint

Disabling an endpoint does not disable the associated port, unless all endpoints using the port are disabled, that is, you are in non- NPIV mode.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Select **More Tasks > Endpoints > Disable**.
3. In the Disable Endpoints dialog, select one or more endpoints from the list, and select **Next**. If an endpoint is in use, you are warned that disabling it might disrupt the system.
4. Select **Next** to complete the task.

Deleting an endpoint

You may want to delete an endpoint if the underlying hardware is no longer available. However, if the underlying hardware is still present, or becomes available, a new endpoint for the hardware is discovered automatically and configured based on default values.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Select **More Tasks > Endpoints > Delete**.
3. In the Delete Endpoints dialog, select one or more endpoints from the list, and select **Next**. If an endpoint is in use, you are warned that deleting it might disrupt the system.
4. Select **Next** to complete the task.

Adding an initiator

Add initiators to provide backup clients to connect to the system to read and write data using the FC (Fibre Channel) protocol. A specific initiator can support DD Boost over FC, or DD VTL, but not both. A maximum of 1024 initiators can be configured for a DD system.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Under Initiators, select Add (+ sign)
3. In the Add Initiator dialog, 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**.

CLI Equivalent

```
# scsitarget group add My_Group initiator My_Initiator
```

Modifying or deleting an initiator

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. You must delete all initiators in an access group before you can delete the access group. If an initiator remains visible, it may be automatically rediscovered.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Under Initiators, select one of the initiators. If you want to delete it, select Delete (X). If you want to modify it, select Modify (pencil) to display the Modify Initiator dialog.
3. Change the initiator's Name and/or 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**.

Recommendation to Set Initiator Aliases - CLI only

It is strongly recommended that Initiator aliases be set to reduce confusion and human error during the configuration process.

```
# vtl initiator set alias NewAliasName wwpn 21:00:00:e0:8b:9d:0b:e8
# vtl initiator show
```

Initiator	Group	Status	WWNN	WWPN	Port
NewVTL	aussiel	Online	20:00:00:e0:8b:9d:0b:e8	21:00:00:e0:8b:9d:0b:e8	6a
		Offline	20:00:00:e0:8b:9d:0b:e8	21:00:00:e0:8b:9d:0b:e8	6b

```
Initiator Symbolic Port Name Address Method
-----
```

NewVTL			auto
--------	--	--	------

```
-----
```

Setting a hard address (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 from 0 to 125.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Select **More Tasks > Set Loop ID**.
3. In the Set Loop ID dialog, enter the loop ID (from 0 to 125), and select **OK**.

Setting failover options

You can set options for automatic failover and failback when NPIV is enabled.

Here is the expected behavior for Fibre Channel port failover, by application:

- DD Boost-over-Fibre Channel operation is expected to continue without user intervention when the Fibre Channel endpoints failover.
- DD VTL Fibre Channel operation is expected to be interrupted when the DD VTL Fibre Channel endpoints failover. You may need to perform discovery (that is, operating system discovery and configuration of DD VTL devices) on the initiators

using the affected Fibre Channel endpoint. You should expect to re-start active backup and restore operations.

- vDisk Fibre Channel operation is expected to continue without user intervention when the Fibre Channel endpoints failover.

Automatic failback is not guaranteed if all ports are disabled and then subsequently enabled (which could be triggered by the administrator), as the order in which ports get enabled is unspecified.

Procedure

1. Select **Hardware > Fibre Channel > Resources**.
2. Select **More Tasks > Set Failover Options**.
3. In the Set Failover Options dialog, enter the Failover and Failback Delay (in seconds) and whether to enable Automatic Failback, and select **OK**.

Access Groups tab

The **Hardware > Fibre Channel > Access Groups** tab provides information about DD Boost and DD VTL access groups. Selecting the link to *View DD Boost Groups* or *View VTL Groups* takes you to the DD Boost or DD VTL pages.

Table 121 Access Groups

Item	Description
Group Name	Name of access group.
Service	Service for this access group: either DD Boost or DD VTL.
Endpoints	Endpoints associated with this access group.
Initiators	Initiators associated with this access group.
Number of Devices	Number of devices associated with this access group.

Differences in FC link monitoring among DD OS versions

Different releases of DD OS handle FC (Fibre Channel) Link Monitoring in different ways.

DD OS 5.3 and later

Port monitoring detects an FC 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 port` 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 DD 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 that 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 DD VTL group to be monitored.

CHAPTER 14

Working with DD Boost

This chapter includes:

- [About Data Domain Boost](#)..... 312
- [Managing DD Boost with DD System Manager](#)..... 313
- [About interface groups](#)..... 327
- [Destroying DD Boost](#)..... 335
- [Configuring DD Boost-over-Fibre Channel](#)..... 335
- [Using DD Boost on HA systems](#)..... 340
- [About the DD Boost tabs](#)..... 340

About Data Domain Boost

Data Domain Boost (DD Boost) provides advanced integration with backup and enterprise applications for increased performance and ease of use. DD Boost distributes parts of the deduplication process to the backup server or application clients, enabling client-side deduplication for faster, more efficient backup and recovery.

DD Boost 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 Data Domain.

Note

A special license, BLOCK-SERVICES-PROTECTPOINT, is available to enable clients using ProtectPoint block services to have DD Boost functionality without a DD Boost license. If DD Boost is enabled for ProtectPoint clients only—that is, if only the BLOCK-SERVICES-PROTECTPOINT license is installed—the license status indicates that DD Boost is enabled for ProtectPoint only.

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 NetWorker backup application, Avamar backup application and other DDBoost partner backup applications, the component that runs on the backup server (DD Boost libraries) is integrated into the particular backup application.
- In the context of Symantec backup applications (NetBackup and Backup Exec) and the Oracle RMAN plug-in, 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.

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 the DD Boost protocol. For Symantec applications, storage units are viewed as disk pools. For NetWorker, storage units are viewed as logical storage units (LSUs). A storage unit is 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 *Data Domain Boost for OpenStorage Administration Guide*. For information on setting up DD Boost with any other application, see the application-specific documentation.

Additional information about configuring and managing DD Boost on the Data Domain system can also be found in the *Data Domain Boost for OpenStorage Administration Guide* (for NetBackup and Backup Exec) and the *Data Domain Boost for Partner Integration Administration Guide* (for other backup applications).

Managing DD Boost with DD System Manager

Access the DD Boost view in DD System Manager.

Procedure

1. Select **Data Management > File System**. Verify that the file system is enabled and running by checking its state.
2. Select **Protocols > DD Boost**.

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.

Note

A special license, BLOCK-SERVICES-PROTECTPOINT, is available to enable clients using ProtectPoint block services to have DD Boost functionality without a DD Boost license. If DD Boost is enabled for ProtectPoint clients only—that is, if only the BLOCK-SERVICES-PROTECTPOINT license is installed—the license status indicates that DD Boost is enabled for ProtectPoint only.

Use the DD Boost tabs—Settings, Active Connections, IP Network, Fibre Channel, and Storage Units—to manage DD Boost.

Specifying DD Boost user names

A DD Boost user is also a DD OS user. Specify a DD Boost user either by selecting an existing DD OS user name or by creating a new DD OS user name and making that name a DD Boost user.

Backup applications use the DD Boost user name and password to connect to the Data Domain system. You must configure these credentials on each backup server that connects to this system. The Data Domain system supports multiple DD Boost users. For complete information about setting up DD Boost with Symantec NetBackup and Backup Exec, see the *Data Domain Boost for OpenStorage Administration Guide*. For information on setting up DD Boost with other applications, see the *Data Domain Boost for Partner Integration Administration Guide* and the application-specific documentation.

Procedure

1. Select **Protocols > DD Boost**.
2. Select **Add (+)** above the Users with DD Boost Access list.

The Add User dialog appears.

3. To select an existing user, select the user name in the drop-down list.

If possible, select a user name with management role privileges set to *none*.

4. To create and select a new user, select **Create a new Local User** and do the following:
 - a. Enter the new user name in the User field.

The user must be configured in the backup application to connect to the Data Domain system.

- b. Enter the password twice in the appropriate fields.
5. Click **Add**.

Changing DD Boost user passwords

Change a DD Boost user password.

Procedure

1. Select **Protocols > DD Boost > Settings**.
2. Select a user in the Users with DD Boost Access list.
3. Click the **Edit** button (pencil icon) above the DD Boost user list.
The Change Password dialog appears.
4. Enter the password twice in the appropriate boxes.
5. Click **Change**.

Removing a DD Boost user name

Remove a user from the DD Boost access list.

Procedure

1. Select **Protocols > DD Boost > Settings**.
2. Select the user in the Users with DD Boost Access list that needs to be removed.
3. Click **Remove (X)** above the DD Boost user list.
The Remove User dialog appears.
4. Click **Remove**.
After removal, the user remains in the DD OS access list.

Enabling DD Boost

Use the DD Boost Settings tab to enable DD Boost and to select or add a DD Boost user.

Procedure

1. Select **Protocols > DD Boost > Settings**.
2. Click **Enable** in the DD Boost Status area.
The Enable DD Boost dialog box is displayed.
3. Select an existing user name from the menu, or add a new user by supplying the name, password, and role.

Configuring Kerberos

You can configure Kerberos by using the DD Boost Settings tab.

Procedure

1. Select **Protocols > DD Boost > Settings**.
2. Click **Configure** in the Kerberos Mode status area.
The Authentication tab under **Administration > Access** appears.

Note

You can also enable Kerberos by navigating directly to Authentication under **Administration > Access** in System Manager.

- Under Active Directory/Kerberos Authentication, click **Configure**.

The Active Directory/Kerberos Authentication dialog box appears. Choose the type of Kerberos Key Distribution Center (KDC) you want to use:

- **Disabled**
If you select **Disabled**, NFS clients will not use Kerberos authentication. CIFS clients will use Workgroup authentication.
- **Windows/Active Directory**
Enter the Realm Name, Under Name, and Password for Active Directory authentication.
- **Unix**
 - Enter the Realm Name, the IP Address/Host Names of one to three KDC servers.
 - Upload the keytab file from one of the KDC servers.

Disabling 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.

Before you begin

Ensure there are no jobs running from your backup application before disabling.

Note

File replication started by DD Boost between two Data Domain restores is not canceled.

Procedure

- Select **Protocols > DD Boost > Settings**.
- Click **Disable** in the DD Boost Status area.
- Click **OK** in the Disable DD Boost confirmation dialog box.

Viewing DD Boost storage units

Access the Storage Units tab to view and manage DD Boost storage units.

The DD Boost Storage Unit tab:

- Lists the storage units and provides the following information for each storage unit:

Table 122 Storage unit information

Item	Description
Storage Unit	The name of the storage unit.
User	The DD Boost user owning the storage unit.

Table 122 Storage unit information (continued)

Item	Description
Quota Hard Limit	The percentage of hard limit quota used.
Last 24 hr Pre-Comp	The amount of raw data from the backup application that has been written in the last 24 hours.
Last 24 hr Post-Comp	The 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	The average amount of compressed storage used in the last five weeks.
Last Week Post-Comp	The 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, modify, and delete storage units.
- Displays four related tabs for a storage unit selected from the list: Storage Unit, Space Usage, Daily Written, and Data Movement.

Note

The Data Movement tab is available only if the optional Data Domain Extended Retention (formerly DD Archiver) license is installed.

- Takes you to **Replication > On-Demand > File Replication** 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.

Creating a storage unit

You must create at least one storage unit on the Data Domain system, and a DD Boost user must be assigned to that storage unit. Use the Storage Units tab to create a storage unit.

Each storage unit is a top-level subdirectory of the `/data/col1` directory; there is no hierarchy among storage units.

Procedure

1. Select **Protocols > DD Boost > Storage Units**.
2. Click **Create (+)**.
The Create Storage Unit dialog box is displayed.
3. Enter the storage unit name in the Name box.

Each storage unit name must be unique. Storage unit 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

Note

The storage-unit name must be enclosed in double quotes (") if the name has an embedded space.

- comma (,)
 - period (.), as long as it does not precede the name
 - exclamation 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 (')
 - minus sign (-)
 - underscore (_)
4. To select an existing username that will have access to this storage unit, select the user name in the dropdown list.
If possible, select a username with management role privileges set to *none*.
 5. To create and select a new username that will have access to this storage unit, select **Create a new Local User** and:
 - a. Enter the new user name in the User box.
The user must be configured in the backup application to connect to the Data Domain system.
 - b. Enter the password twice in the appropriate boxes.
 6. 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.

7. Click **Create**.
8. Repeat the above steps for each Data Domain Boost-enabled system.

Viewing storage unit information

From the DD Boost Storage Units tab, you can select a storage unit and access the Storage Unit, Space Usage, Daily Written, and Data Movement tabs for the selected storage unit.

Storage Unit tab

The Storage Unit tab shows detailed information for a selected storage unit in its Summary and Quota panels. The Snapshot panel shows snapshot details, allows you to create new snapshots and schedules, and provides a link to the **Data Management > Snapshots** tab.

- The Summary panel shows summarized information for the selected storage unit.

Table 123 Summary panel

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.

- The Quota panel shows quota information for the selected storage unit.

Table 124 Quota panel

Quota item	Description
Quota Enforcement	Enabled or disable. Clicking Quota takes you to the Data Management > Quota 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 panel.
 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**.
- **Snapshots**
The Snapshots panel shows information about the storage unit's snapshots.

Table 125 Snapshots panel

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.

Using the Snapshots panel, you can:

- 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

The snapshot name can be composed only of letters, numbers, `_`, `-`, `%d` (numeric day of the month: 01-31), `%a` (abbreviated weekday name), `%m` (numeric month of the year: 01-12), `%b` (abbreviated month name), `%Y` (year, two digits), `%Y` (year, four digits), `%H` (hour: 00-23), and `%M` (minute: 00-59), 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.

- Click the Snapshots link to go 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.

- 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 displayed: Logical Space Used (Pre-Compression) and Physical Capacity Used (Post-Compression).

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.

Data Movement tab

A graph in the same format as the Daily Written graph that shows the amount of disk space moved to the DD Extended Retention storage area (if the DD Extended Retention license is enabled).

Modifying a storage unit

Use the Modify Storage Unit dialog to rename a storage unit, select a different existing user, create and select a new user, and edit quota settings.

Procedure

1. Select **Protocols > DD Boost > Storage Units**.
2. In the Storage Unit list, select the storage unit to modify.
3. Click the pencil icon.

The Modify Storage Unit dialog appears.

4. To rename the storage unit, edit the text in the **Name** field.
5. To select a different existing user, select the user name in the drop-down list.
If possible, select a username with management role privileges set to *none*.
6. To create and select a new user, select **Create a new Local User** and do the following:
 - a. Enter the new user name in the User box.

The user must be configured in the backup application to connect to the Data Domain system.
 - b. Enter the password twice in the appropriate boxes.
7. Edit the Quota Settings as needed.

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.

8. Click **Modify**.

Renaming a storage unit

Use the Modify Storage Unit dialog to rename a storage unit.

Renaming a storage unit changes the name of the storage unit while retaining its:

- Username ownership
- Stream limit configuration
- Capacity quota configuration and physical reported size
- AIR association on the local Data Domain system

Procedure

1. Go to **Protocols > DD Boost > Storage Units**.
2. In the Storage Unit list, select the storage unit to rename.
3. Click the pencil icon.

The Modify Storage Unit dialog appears.

4. Edit the text in the **Name** field.
5. Click **Modify**.

Deleting a storage unit

Use the Storage Units tab to delete a storage unit from your Data Domain system. Deleting a storage unit removes the storage unit, as well as any images contained in the storage unit, from your Data Domain system.

Procedure

1. Select **Protocols > DD Boost > Storage Units**.
2. Select the storage unit to be deleted from the list.
3. Click **Delete (X)**.
4. Click **OK**.

Results

The storage unit is removed from your Data Domain system. You must also manually remove the corresponding backup application catalog entries.

Undeleting a storage unit

Use the Storage Units tab to undelete a storage unit.

Undeleting a storage unit recovers a previously deleted storage unit, including its:

- Username ownership
- Stream limit configuration
- Capacity quota configuration and physical reported size
- AIR association on the local Data Domain system

Note

Deleted storage units are available until the next `filesys clean` command is run.

Procedure

1. Select **Protocols > DD Boost > Storage Units > More Tasks > Undelete Storage Unit....**
2. In the Undelete Storage Units dialog box, select the storage unit(s) that you want to undelete.
3. Click **OK**.

Selecting DD Boost options

Use the Set DD Boost Options dialog to specify settings for distributed segment processing, virtual synthetics, low bandwidth optimization for file replication, file replication encryption, and file replication network preference (IPv4 or IPv6).

Procedure

1. To display the DD Boost option settings, select **Protocols > DD Boost > Settings > Advanced Options**.
2. To change the settings, select **More Tasks > Set Options**.

The Set DD Boost Options dialog appears.

3. Select any option to be enabled.
4. Deselect any option to be disabled.
To deselect a File Replication Network Preference option, select the other option.
5. Set the DD Boost security options.
 - a. Select the **Authentication Mode**:
 - None
 - Two-way
 - Two-way Password
 - b. Select the **Encryption Strength**:
 - None
 - Medium
 - High

The Data Domain system compares the global authentication mode and encryption strength against the per-client authentication mode and encryption strength to calculate the effective authentication mode and authentication encryption strength. The system does not use the highest authentication mode from one entry, and the highest encryption settings from a different entry. The effective authentication mode and encryption strength come from the single entry that provides the highest authentication mode.

6. Click **OK**.

Note

You can also manage distributed segment processing via the `ddboost option` commands, which are described in detail in the *Data Domain Operating System Command Reference Guide*.

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.

You can manage distributed segment processing via the `ddboost option` commands, which are described in detail in the *Data Domain Operating System Command Reference Guide*.

Note

Distributed segment processing is enabled by default with Data Domain Extended Retention (formerly Data Domain Archiver) configurations and cannot be disabled.

Virtual synthetics

A virtual synthetic full backup is the combination of the last full (synthetic or full) backup and all subsequent incremental backups. Virtual synthetics are enabled by default.

Low-bandwidth optimization

If you use 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 *Data Domain Operating System Command Reference Guide*.

File replication encryption

You can encrypt the data replication stream by enabling the DD Boost file replication encryption option.

Note

If DD Boost file replication encryption is used on systems without the Data at Rest option, 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, use the same global listen port on both the source and target Data Domain systems. To set the listen port, use the `replication option` command as described in the *Data Domain Operating System Command Reference Guide*.

File replication network preference

Use this option to set the preferred network type for DD Boost file replication to either IPv4 or IPv6.

Managing certificates for DD Boost

A host certificate allows DD Boost client programs to verify the identity of the system when establishing a connection. CA certificates identify certificate authorities that should be trusted by the system. The topics in this section describe how to manage host and CA certificates for DD Boost.

Adding a host certificate for DD Boost

Add a host certificate to your system. DD OS supports one host certificate for DD Boost.

Procedure

1. If you have not yet requested a host certificate, request one from a trusted CA.
2. When you have received a host certificate, copy or move it to the computer from which you run DD Service Manager.
3. Start DD System Manager on the system to which you want to add a host certificate.

Note

DD System Manager supports certificate management only on the management system (which is the system running DD System Manager).

4. Select **Protocols > DD Boost > More Tasks > Manage Certificates....**
-

Note

If you try to remotely manage certificates on a managed system, DD System Manager displays an information message at the top of the certificate management dialog. To manage certificates for a system, you must start DD System Manager on that system.

5. In the Host Certificate area, click **Add**.
6. To add a host certificate enclosed in a .p12 file, do the following:
 - a. Select **I want to upload the certificate as a .p12 file**.
 - b. Type the password in the **Password** box.
 - c. Click **Browse** and select the host certificate file to upload to the system.
 - d. Click **Add**.
7. To add a host certificate enclosed in a .pem file, do the following:
 - a. Select **I want to upload the public key as a .pem file and use a generated private key**.
 - b. Click **Browse** and select the host certificate file to upload to the system.
 - c. Click **Add**.

Adding CA certificates for DD Boost

Add a certificate for a trusted CA to your system. DD OS supports multiple certificates for trusted CAs.

Procedure

1. Obtain a certificate for the trusted CA.
 2. Copy or move the trusted CA certificate to the computer from which you run DD Service Manager.
 3. Start DD System Manager on the system to which you want to add the CA certificate.
-

Note

DD System Manager supports certificate management only on the management system (which is the system running DD System Manager).

4. Select **Protocols > DD Boost > More Tasks > Manage Certificates....**

Note

If you try to remotely manage certificates on a managed system, DD System Manager displays an information message at the top of the certificate management dialog. To manage certificates for a system, you must start DD System Manager on that system.

5. In the CA Certificates area, click **Add**.
The Add CA Certificate for DD Boost dialog appears.
6. To add a CA certificate enclosed in a .pem file, do the following:
 - a. Select **I want to upload the certificate as a .pem file**.
 - b. Click **Browse**, select the host certificate file to upload to the system, and click **Open**.
 - c. Click **Add**.
7. To add a CA certificate using copy and paste, do the following:
 - a. Copy the certificate text to the clipboard using the controls in your operating system.
 - b. Select **I want to copy and paste the certificate text**.
 - c. Paste the certificate text in the box below the copy and paste selection.
 - d. Click **Add**.

Managing DD Boost client access and encryption

Use the DD Boost Settings tab to configure which specific clients, or set of clients, can establish a DD Boost connection with the Data Domain System and whether or not the client will use encryption. By default, the system is configured to allow all clients to have access, with no encryption.

Note

Enabling in-flight encryption will impact system performance.

Note

DD Boost offers global authentication and encryption options to defend your system against man-in-the-middle (MITM) attacks. You specify authentication and encryption settings using the GUI, or CLI commands on the Data Domain system. For details, see the *Data Domain Boost for OpenStorage 3.4 Administration Guide*, and [Adding a DD Boost client](#) on page 326 or the *Data Domain 6.1 Command Reference Guide*.

Adding a DD Boost client

Create an allowed DD Boost client and specify whether the client will use encryption.

Procedure

1. Select **Protocols > DD Boost > Settings**.
2. In the Allowed Clients section, click **Create (+)**.

The Add Allowed Client dialog appears.

3. Enter the hostname of the client.

This can be a fully-qualified domain name (e.g. host1.emc.com) or a hostname with a wildcard (e.g. *.emc.com).

4. Select the Encryption Strength.

The options are None (no encryption), Medium (AES128-SHA1), or High (AES256-SHA1).

5. Select the Authentication Mode.

The options are One Way, Two Way, Two Way Password, or Anonymous.

6. Click **OK**.

Modifying a DD Boost client

Change the name, encryption strength, and authentication mode of an allowed DD Boost client.

Procedure

1. Select **Protocols > DD Boost > Settings**.
2. In the Allowed Clients list, select the client to modify.
3. Click the **Edit** button, which displays a pencil icon.
The Modify Allowed Client dialog appears.
4. To change the name of a client, edit the Client text.
5. To change the Encryption Strength, select the option.
The options are None (no encryption), Medium (AES128-SHA1), or High (AES256-SHA1).
6. To change the Authentication Mode, select the option.
The options are One Way, Two Way, or Anonymous.
7. Click **OK**.

Removing a DD Boost client

Delete an allowed DD Boost client.

Procedure

1. Select **Protocols > DD Boost > Settings**.
2. Select the client from the list.
3. Click **Delete (X)**.
The Delete Allowed Clients dialog appears.
4. Confirm and select the client name. Click **OK**.

About interface groups

This feature lets you combine multiple Ethernet links into a group and register only one interface on the Data Domain system with the backup application. The DD Boost

Library negotiates with the Data Domain system to obtain the best interface to send data. Load balancing provides higher physical throughput to the Data Domain system.

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, and the group interface uses load balancing 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.

Should an interface in the group fail, all the in-flight jobs to that interface are automatically resumed on healthy operational links (unbeknownst 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.

Consider the following information when managing interface groups.

- The IP address must be configured on the Data Domain system, and its interface enabled. To check the interface configuration, select **Hardware > Ethernet > Interfaces** page, and check for free ports. See the `net` chapter of the *Data Domain Operating System Command Reference Guide* or the *Data Domain Operating System Initial Configuration Guide* for information about configuring an IP address for an interface.
- You can use the `ifgroup` commands to manage interface groups; these commands are described in detail in the *Data Domain Operating System Command Reference Guide*.
- Interface groups provide full support for static IPv6 addresses, providing the same capabilities for IPv6 as for IPv4. Concurrent IPv4 and IPv6 client connections are allowed. A client connected with IPv6 sees IPv6 ifgroup interfaces only. A client connected with IPv4 sees IPv4 ifgroup interfaces only. Individual ifgroups include all IPv4 addresses or all IPv6 addresses. For details, see the *Data Domain Boost for Partner Integration Administration Guide* or the *Data Domain Boost for OpenStorage Administration Guide*.
- Configured interfaces are listed in Active Connections, on the lower portion of the Activities page.

Note

See [Using DD Boost on HA systems](#) on page 340 for important information about using interface groups with HA systems.

The topics that follow describe how to manage interface groups.

Interfaces

IFGROUP supports physical and virtual interfaces.

An IFGROUP interface is a member of a single IFGROUP *<group-name>* and may consist of:

- Physical interface such as `eth0a`
- Virtual interface, created for link failover or link aggregation, such as `veth1`
- Virtual alias interface such as `eth0a:2` or `veth1:2`
- Virtual VLAN interface such as `eth0a.1` or `veth1.1`
- Within an IFGROUP *<group-name>*, all interfaces must be on unique interfaces (Ethernet, virtual Ethernet) to ensure failover in the event of network error.

IFGROUP provides full support for static IPv6 addresses, providing the same capabilities for IPv6 as for IPv4. Concurrent IPv4 and IPv6 client connections are allowed. A client connected with IPv6 sees IPv6 IFGROUP interfaces only. A client connected with IPv4 sees IPv4 IFGROUP interfaces only. Individual IFGROUPs include all IPv4 addresses or all IPv6 addresses.

For more information, see the *Data Domain Boost for Partner Integration Administration Guide* or the *Data Domain Boost for OpenStorage Administration Guide*.

Interface enforcement

IFGROUP lets you enforce private network connectivity, ensuring that a failed job does not reconnect on the public network after network errors.

When interface enforcement is enabled, a failed job can only retry on an alternative private network IP address. Interface enforcement is only available for clients that use IFGROUP interfaces.

Interface enforcement is off (FALSE) by default. To enable interface enforcement, you must add the following setting to the system registry:

```
system.ENFORCE_IFGROUP_RW=TRUE
```

After you've made this entry in the registry, you must do a `filesys restart` for the setting to take effect.

For more information, see the *Data Domain Boost for Partner Integration Administration Guide* or the *Data Domain Boost for OpenStorage Administration Guide*.

Clients

IFGROUP supports various naming formats for clients. Client selection is based on a specified order of precedence.

An IFGROUP client is a member of a single ifgroup *<group-name>* and may consist of:

- A fully qualified domain name (FQDN) such as `ddbboost.datadomain.com`
- A partial host, allowing search on the first *n* characters of the hostname. For example, when *n*=3, valid formats are `rtp_.*emc.com` and `dur_.*emc.com`. Five different values of *n* (1-5) are supported.
- Wild cards such as `*.datadomain.com` or `**`
- A short name for the client, such as `ddbboost`
- Client public IP range, such as `128.5.20.0/24`

Prior to write or read processing, the client requests an IFGROUP IP address from the server. To select the client IFGROUP association, the client information is evaluated according to the following order of precedence.

1. IP address of the connected Data Domain system. If there is already an active connection between the client and the Data Domain system, and the connection exists on the interface in the IFGROUP, then the IFGROUP interfaces are made available for the client.
2. Connected client IP range. An IP mask check is done against the client source IP; if the client's source IP address matches the mask in the IFGROUP clients list, then the IFGROUP interfaces are made available for the client.

- For IPv4, you can select five different range masks, based on network.
- For IPv6, fixed masks /64, /112, and /128 are available.

This host-range check is useful for separate VLANs with many clients where there isn't a unique partial hostname (domain).

3. Client Name: `abc-11.dl.com`
4. Client Domain Name: `*.dl.com`
5. All Clients: `*`

For more information, see the *Data Domain Boost for Partner Integration Administration Guide*.

Creating interface groups

Use the IP Network tab to create interface groups and to add interfaces and clients to the groups.

Multiple interface groups improve the efficiency of DD Boost by allowing you to:

- 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.

Create interface groups first, and then add clients (as new media servers become available) to an interface group.

Procedure

1. Select **Protocols > DD Boost > IP Network**.
2. In the Interface Groups section, click Add (+).
3. Enter the interface group name.
4. 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.

5. Click **OK**.
6. In the Configured Clients section, click Add (+).
7. 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.

8. Select a previously configured interface group, and click **OK**.

Enabling and disabling interface groups

Use the IP Network tab to enable and disable interface groups.

Procedure

1. Select **Protocols > DD Boost > IP Network**.
2. In the Interface Groups section, 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.

3. Click **Edit** (pencil).
4. Click **Enabled** to enable the interface group; clear the checkbox to disable.
5. Click **OK**.

Modifying an interface group's name and interfaces

Use the IP Network tab to change an interface group's name and the interfaces associated with the group.

Procedure

1. Select **Protocols > DD Boost > IP Network**.
2. In the Interface Groups section, select the interface group in the list.
3. Click **Edit** (pencil).
4. 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.”

5. Select or deselect client interfaces in the Interfaces list.

Note

If you remove all interfaces from the group, it will be automatically disabled.

6. Click **OK**.

Deleting an interface group

Use the IP Network tab to delete an interface group. Deleting an interface group deletes all interfaces and clients associated with the group.

Procedure

1. Select **Protocols > DD Boost > IP Network**.
2. In the Interface Groups section, select the interface group in the list. The default group cannot be deleted.
3. Click Delete (X).
4. Confirm the deletion.

Adding a client to an interface group

Use the IP Network tab to add clients to interface groups.

Procedure

1. Select **Protocols > DD Boost > IP Network**.
2. In the Configured Clients section, click Add (+).
3. Enter a name for the client.

Client names must be unique and may consist of:

- FQDN
- *.domain
- Client public IP range:
 - For IPv4, `xx.xx.xx.0/24` provides a 24-bit mask against the connecting IP. The /24 represents what bits are masked when the client's source IP address is evaluated for access to the IFGROUP.
 - For IPv6, `xxxxx::0/112` provides a 112-bit mask against the connecting IP. The /112 represents what bits are masked when the client's source IP address is evaluated for access to the IFGROUP.

Client names have a maximum length of 128 characters.

4. Select a previously configured interface group, and click **OK**.

Modifying a client's name or interface group

Use the IP Network tab to change a client's name or interface group.

Procedure

1. Select **Protocols > DD Boost > IP Network**.
2. In the Configured Clients section, select the client.
3. Click **Edit** (pencil).
4. Type a new client name.

Client names must be unique and may consist of:

- FQDN
- *.domain
- Client public IP range:

- For IPv4, `xx.xx.xx.0/24` provides a 24-bit mask against the connecting IP. The `/24` represents what bits are masked when the client's source IP address is evaluated for access to the IFGROUP.
- For IPv6, `xxxx: :0/112` provides a 112-bit mask against the connecting IP. The `/112` represents what bits are masked when the client's source IP address is evaluated for access to the IFGROUP.

Client names have a maximum length of 128 characters.

5. Select a new interface group from the menu.

Note

The old interface group is disabled if it has no clients.

6. Click **OK**.

Deleting a client from the interface group

Use the IP Network tab to delete a client from an interface group.

Procedure

1. Select **Protocols > DD Boost > IP Network**.
2. In the Configured Clients section, select the client.
3. Click Delete (**X**).

Note

If the interface group to which the client belongs has no other clients, the interface group is disabled.

4. Confirm the deletion.

Using interface groups for Managed File Replication (MFR)

Interface groups can be used to control the interfaces used for DD Boost MFR, to direct the replication connection over a specific network, and to use multiple network interfaces with high bandwidth and reliability for failover conditions. All Data Domain IP types are supported—IPv4 or IPv6, Alias IP/VLAN IP, and LACP/failover aggregation.

Note

Interface groups used for replication are different from the interface groups previously explained and are supported for DD Boost Managed File Replication (MFR) only. For detailed information about using interface groups for MFR, see the *Data Domain Boost for Partner Integration Administration Guide* or the *Data Domain Boost for OpenStorage Administration Guide*.

Without the use of interface groups, configuration for replication requires several steps:

1. Adding an entry in the `/etc/hosts` file on the source Data Domain system for the target Data Domain system and hard coding one of the private LAN network interfaces as the destination IP address.

2. Adding a route on the source Data Domain system to the target Data Domain system specifying a physical or virtual port on the source Data Domain system to the remote destination IP address.
3. Configuring LACP through the network on all switches between the Data Domain systems for load balancing and failover.
4. Requiring different applications to use different names for the target Data Domain system to avoid naming conflicts in the `/etc/hosts` file.

Using interface groups for replication simplifies this configuration through the use of the DD OS System Manager or DD OS CLI commands. Using interface groups to configure the replication path lets you:

- Redirect a hostname-resolved IP address away from the public network, using another private Data Domain system IP address.
- Identify an interface group based on configured selection criteria, providing a single interface group where all the interfaces are reachable from the target Data Domain system.
- Select a private network interface from a list of interfaces belonging to a group, ensuring that the interface is healthy.
- Provide load balancing across multiple Data Domain interfaces within the same private network.
- Provide a failover interface for recovery for the interfaces of the interface group.
- Provide host failover if configured on the source Data Domain system.
- Use Network Address Translation (NAT)

The selection order for determining an interface group match for file replication is:

1. Local MTree (storage-unit) path and a specific remote Data Domain hostname
2. Local MTree (storage-unit) path with any remote Data Domain hostname
3. Any MTree (storage-unit) path with a specific Data Domain hostname

The same MTree can appear in multiple interface groups only if it has a different Data Domain hostname. The same Data Domain hostname can appear in multiple interface groups only if it has a different MTree path. The remote hostname is expected to be an FQDN, such as `dd890-1.emc.com`.

The interface group selection is performed locally on both the source Data Domain system and the target Data Domain system, independent of each other. For a WAN replication network, only the remote interface group needs to be configured since the source IP address corresponds to the gateway for the remote IP address.

Adding a replication path to an interface group

Use the IP Network tab to add replication paths to interface groups.

Procedure

1. Select **Protocols > DD Boost > IP Network**.
2. In the Configured Replication Paths section, click Add (+).
3. Enter values for **MTree** and/or **Remote Host**.
4. Select a previously configured interface group, and click **OK**.

Modifying a replication path for an interface group

Use the IP Network tab to modify replication paths for interface groups.

Procedure

1. Select **Protocols > DD Boost > IP Network**.
2. In the Configured Replication Paths section, select the replication path.
3. Click **Edit** (pencil).
4. Modify any or all values for **MTree**, **Remote Host**, or **Interface Group**.
5. Click **OK**.

Deleting a replication path for an interface group

Use the IP Network tab to delete replication paths for interface groups.

Procedure

1. Select **Protocols > DD Boost > IP Network**.
2. In the Configured Replication Paths section, select the replication path.
3. Click **Delete (X)**.
4. In the Delete Replication Path(s) dialog, click **OK**.

Destroying DD Boost

Use this option to permanently remove 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. Select **Protocols > DD Boost > More Tasks > Destroy DD Boost....**
3. Enter your administrative credentials when prompted.
4. Click **OK**.

Configuring DD Boost-over-Fibre Channel

In earlier versions of DD OS, all communication between the DD Boost Library and any Data Domain system was performed using IP networking. DD OS now offers Fibre Channel as an alternative transport mechanism for communication between the DD Boost Library and the Data Domain system.

Note

Windows, Linux, HP-UX (64-bit Itanium architecture), AIX, and Solaris client environments are supported.

Enabling DD Boost users

Before you can configure the DD Boost-over-FC service on a Data Domain system, you must add one or more DD Boost users and enable DD Boost.

Before you begin

- Log in to DD System Manager. For instructions, see “Logging In and Out of DD System Manager.”

CLI equivalent

```
login as: sysadmin
Data Domain OS 5.7.x.x-12345
Using keyboard-interactive authentication.
Password:
```

- If you are using the CLI, ensure that the SCSI target daemon is enabled:

```
# scsitararget enable
Please wait ...
SCSI Target subsystem is enabled.
```

Note

If you are using DD System Manager, the SCSI target daemon is automatically enabled when you enable the DD Boost-over-FC service (later in this procedure).

- Verify that the DD Boost license is installed. In DD System Manager, select **Protocols > DD Boost > Settings**. If the Status indicates that DD Boost is not licensed, click **Add License** and enter a valid license in the Add License Key dialog box.

CLI equivalents

```
# license show

# license add license-code
```

Procedure

1. Select **Protocols > DD Boost > Settings**.
2. In the Users with DD Boost Access section, specify one or more DD Boost user names.

A DD Boost user is also a DD OS user. When specifying a DD Boost user name, you can select an existing DD OS user name, or you can create a new DD OS user name and make that name a DD Boost user. This release supports multiple DD Boost users. For detailed instructions, see “Specifying DD Boost User Names.”

CLI equivalents

```
# user add username [password password]

# ddbboost set user-name exampleuser
```

3. Click **Enable** to enable DD Boost.

CLI equivalent

```
# ddbboost enable
Starting DDBOOST, please wait.....
DDBOOST is enabled.
```


Results

You are now ready to configure the DD Boost-over-FC service on the Data Domain system.

Configuring DD Boost

After you have added user(s) and enabled DD Boost, you need to enable the Fibre Channel option and specify the DD Boost Fibre Channel server name. Depending on your application, you may also need to create one or more storage units and install the DD Boost API/plugin on media servers that will access the Data Domain system.

Procedure

1. Select **Protocols > DD Boost > Fibre Channel**.
2. Click **Enable** to enable Fibre Channel transport.

CLI equivalent

```
# ddbboost option set fc enabled
Please wait...
DD Boost option "FC" set to enabled.
```

3. To change the DD Boost Fibre Channel server name from the default (hostname), click **Edit**, enter a new server name, and click **OK**.

CLI equivalent

```
# ddbboost fc dfc-server-name set DFC-ddbeta2
DDBoost dfc-server-name is set to "DFC-ddbeta2" for DDBoost FC.
Configure clients to use "DFC-DFC-ddbeta2" for DDBoost FC.
```

4. Select **Protocols > DD Boost > Storage Units** to create a storage unit (if not already created by the application).

You must create at least one storage unit on the Data Domain system, and a DD Boost user must be assigned to that storage unit. For detailed instructions, see “Creating a Storage Unit.”

CLI equivalent

```
# ddbboost storage-unit create storage_unit_name-su
```

5. Install the DD Boost API/plugin (if necessary, based on the application).

The DD Boost OpenStorage plug-in software must be installed on NetBackup media servers that need to access the Data Domain system. This plug-in includes the required DD Boost Library that integrates with the Data Domain system. For detailed installation and configuration instructions, see the *Data Domain Boost for Partner Integration Administration Guide* or the *Data Domain Boost for OpenStorage Administration Guide*.

Results

You are now ready to verify connectivity and create access groups.

Verifying connectivity and creating access groups

Go to **Hardware > Fibre Channel > Resources** to manage initiators and endpoints for access points. Go to **Protocols > DD Boost > Fibre Channel** to create and manage DD Boost-over-FC 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.

Procedure

1. Select **Hardware > Fibre Channel > Resources > Initiators** to verify that initiators are present.

It is recommended that you assign aliases to initiators to reduce confusion during the configuration process.

CLI equivalent

```
# scsitarget initiator show list
Initiator          System Address          Group      Service
-----
initiator-1        21:00:00:24:ff:31:b7:16  n/a       n/a
initiator-2        21:00:00:24:ff:31:b8:32  n/a       n/a
initiator-3        25:00:00:21:88:00:73:ee  n/a       n/a
initiator-4        50:06:01:6d:3c:e0:68:14  n/a       n/a
initiator-5        50:06:01:6a:46:e0:55:9a  n/a       n/a
initiator-6        21:00:00:24:ff:31:b7:17  n/a       n/a
initiator-7        21:00:00:24:ff:31:b8:33  n/a       n/a
initiator-8        25:10:00:21:88:00:73:ee  n/a       n/a
initiator-9        50:06:01:6c:3c:e0:68:14  n/a       n/a
initiator-10       50:06:01:6b:46:e0:55:9a  n/a       n/a
tsm6_p23           21:00:00:24:ff:31:ce:f8  SetUp_Test VTL
```

2. To assign an alias to an initiator, select one of the initiators and click the pencil (edit) icon. In the Name field of the Modify Initiator dialog, enter the alias and click OK.

CLI equivalents

```
# scsitarget initiator rename initiator-1 initiator-renamed
Initiator 'initiator-1' successfully renamed.
```

```
# scsitarget initiator show list
Initiator          System Address          Group
Service
-----
initiator-2        21:00:00:24:ff:31:b8:32  n/a
n/a
initiator-renamed  21:00:00:24:ff:31:b7:16  n/a
n/a
```

3. On the Resources tab, verify that endpoints are present and enabled.

CLI equivalent

```
# scsitarget endpoint show list
-----
endpoint-fc-0      5a                      FibreChannel  Yes          Online
```

endpoint-fc-1	5b	FibreChannel	Yes	Online
-----	-----	-----	-----	-----

4. Go to **Protocols > DD Boost > Fibre Channel**.
5. In the DD Boost Access Groups area, click the + icon to add an access group.
6. Enter a unique name for the access group. Duplicate names are not supported.

CLI equivalent

```
# ddbboost fc group create test-dfc-group
DDBoost FC Group "test-dfc-group" successfully created.
```

7. Select one or more initiators. Optionally, replace the initiator name by entering a new one. Click **Next**.

CLI equivalent

```
# ddbboost fc group add test-dfc-group initiator initiator-5
Initiator(s) "initiator-5" added to group "test-dfc-group".
```

An initiator is a port on an HBA attached to a backup client that connects to the system for the purpose of reading and writing data using the Fibre Channel protocol. The WWPN is the unique World-Wide Port Name of the Fibre Channel port in the media server.

8. Specify the number of DD Boost devices to be used by the group. This number determines which devices the initiator can discover and, therefore, the number of I/O paths to the Data Domain system. The default is one, the minimum is one, and the maximum is 64.

CLI equivalent

```
# ddbboost fc group modify Test device-set count 5
Added 3 devices.
```

See the *Data Domain Boost for OpenStorage Administration Guide* for the recommended value for different clients.

9. Indicate which endpoints to include in the group: all, none, or select from the list of endpoints. Click **Next**.

CLI equivalents

```
# scsitarget group add Test device ddbboost-dev8 primary-
endpoint all
secondary-endpoint all
Device 'ddbboost-dev8' successfully added to group.
```

```
# scsitarget group add Test device ddbboost-dev8 primary-endpoint
endpoint-fc-1 secondary-endpoint fc-port-0
Device 'ddbboost-dev8' is already in group 'Test'.
```

When presenting LUNs via attached FC ports on HBAs, ports can be designated as primary, secondary or none. A primary port for a set of LUNs is the port that is currently advertizing those LUNs to a fabric. A secondary port is a port that will broadcast a set of LUNs in the event of primary path failure (this requires manual intervention). A setting of none is used in the case where you do not wish to advertize selected LUNs. The presentation of LUNs is dependent upon the SAN topology.

10. Review the Summary and make any modifications. Click **Finish** to create the access group, which is displayed in the DD Boost Access Groups list.

CLI equivalent

```
# scsitarget group show detailed
```

Note

To change settings for an existing access group, select it from the list and click the pencil icon (Modify).

Deleting access groups

Use the Fibre Channel tab to delete access groups.

Procedure

1. Select **Protocols > DD Boost > Fibre Channel**.
 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).

Using DD Boost on HA systems

HA provides seamless failover of any application using DD Boost—that is, any backup or restore operation continues with no manual intervention required. All other DD Boost user scenarios are supported on HA systems as well, including managed file replication (MFR), distributed segment processing (DSP), filecopy, and dynamic interface groups (DIG).

Note these special considerations for using DD Boost on HA systems:

- On HA-enabled Data Domain systems, failovers of the DD server occur in less than 10 minutes. However, recovery of DD Boost applications may take longer than this, because Boost application recovery cannot begin until the DD server failover is complete. In addition, Boost application recovery cannot start until the application invokes the Boost library.
- DD Boost on HA systems requires that the Boost applications be using Boost HA libraries; applications using non-HA Boost libraries do not see seamless failover.
- MFR will fail over seamlessly when both the source and destination systems are HA-enabled. MFR is also supported on partial HA configurations (that is, when either the source or destination system is enabled, but not both) when the failure occurs on the HA-enabled system. For more information, see the *DD Boost for OpenStorage Administration Guide* or the *DD Boost for Partner Integration Administration Guide*.
- Dynamic interface groups should not include IP addresses associated with the direct interconnection between the active and standby Data Domain systems.
- DD Boost clients must be configured to use floating IP addresses.

About the DD Boost tabs

Learn to use the DD Boost tabs in DD System Manager.

Settings

Use the Settings tab to enable or disable DD Boost, select clients and users, and specify advanced options.

The Settings tab shows the DD Boost status (Enabled or Disabled). Use the **Status** button to switch between **Enabled** or **Disabled**.

Under **Allowed Clients**, select the clients that are to have access to the system. Use the **Add**, **Modify**, and **Delete** buttons to manage the list of clients.

Under **Users with DD Boost Access**, select the users that are to have DD Boost access. Use the **Add**, **Change Password**, and **Remove** buttons to manage the list of users.

Expand **Advanced Options** to see which advanced options are enabled. Go to **More Tasks > Set Options** to reset these options.

Active Connections

Use the Active Connections tab to see information about clients, interfaces, and outbound files.

Table 126 Connected client information

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.
Encrypted	Whether the connection is encrypted (Yes) or not (No).
DSP	Whether or not the connection is using Distributed Segment Processing (DSP) or not.
Transport	Type of transport being used, such as IPv4, IPv6 or DFC (Fibre Channel).

Table 127 Configured interface connection information

Item	Description
Interface	The IP address of the interface.
Interface Group	One of the following: <ul style="list-style-type: none"> The name of the interface group. None, if not a member of one.
Backup	The number of active backup connections.

Table 127 Configured interface connection information (continued)

Item	Description
Restore	The number of active restore connections.
Replication	The number of active replication connections.
Synthetic	The number of synthetic backups.
Total	The total number of connections for the interface.

Table 128 Outbound file replication information

Outbound files item	Description
File Name	The name of the outgoing 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.

IP Network

The IP Network tab lists configured interface groups. Details include whether or not a group is enabled and any configured client interfaces. Administrators can use the Interface Group menu to view which clients are associated with an interface group.

Fibre Channel

The Fibre Channel tab lists configured DD Boost access groups. Use the Fibre Channel tab to create and delete access groups and to configure initiators, devices, and endpoints for DD Boost access groups.

Storage Units

Use the Storage Unit tab to create, modify, and delete storage units. To see detailed information about a listed storage unit, select its name.

Table 129 Storage unit: Detailed information

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.

Table 129 Storage unit: Detailed information (continued)

Item	Description
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"> • D—Deleted • RO—Read-only • RW—Read/write • RD—Replication destination • RLE—DD Retention lock enabled • RLD—DD Retention lock disabled
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.
Locally Compressed Size	Total size after local compression of the files in the storage unit when they were written.

CHAPTER 15

DD Virtual Tape Library

This chapter includes:

- [DD Virtual Tape Library overview](#)..... 346
- [Planning a DD VTL](#)..... 346
- [Managing a DD VTL](#)..... 353
- [Working with libraries](#)..... 357
- [Working with a selected library](#)..... 360
- [Viewing changer information](#)..... 368
- [Working with drives](#)..... 368
- [Working with a selected drive](#)..... 371
- [Working with tapes](#)..... 371
- [Working with the vault](#)..... 373
- [Working with the cloud-based vault](#)..... 373
- [Working with access groups](#)..... 380
- [Working with a selected access group](#)..... 384
- [Working with resources](#)..... 386
- [Working with pools](#)..... 391
- [Working with a selected pool](#)..... 393

DD Virtual Tape Library overview

Data Domain Virtual Tape Library (DD VTL) is a disk-based backup system that emulates the use of physical tapes. It enables backup applications to connect to and manage DD system storage using functionality almost identical to a physical tape library.

Virtual tape drives are accessible to backup software in the same way as physical tape drives. After you create these drives in a DD VTL, they appear to the backup software as SCSI tape drives. The DD VTL, itself, appears to the backup software as a SCSI robotic device accessed through standard driver interfaces. However, the backup software (not the DD system that is configured as a DD VTL) manages the movement of the media changer and backup images.

The following terms have special meaning when used with DD VTL:

- *Library*: A library emulates a physical tape library with drives, changer, CAPs (cartridge access ports), and slots (cartridge slots).
- *Tape*: A tape is represented as a file. 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 maps to a directory on the file system. Pools are used to replicate tapes to a destination. You can convert directory-based pools to MTree-based pools 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.

DD VTL has been tested with, and is supported by, specific backup software and hardware configurations. For more information, see the appropriate *Backup Compatibility Guide* on the Online Support Site.

DD VTL supports simultaneous use of the tape library and file system (NFS/CIFS/DD Boost) interfaces.

When DR (disaster recovery) is needed, pools and tapes can be replicated to a remote DD system using the DD Replicator.

To protect data on tapes from modification, tapes can be locked using DD Retention Lock Governance software.

Note

At present, for 16 Gb/s, Data Domain supports fabric and point-to-point topologies. Other topologies will present issues.

Planning a DD VTL

The DD VTL (Virtual Tape Library) feature has very specific requirements, such as proper licensing, interface cards, user permissions, etc. These requirements are listed here, complete with details and recommendations.

- An appropriate DD VTL license.
 - DD VTL is a licensed feature, and you must use NDMP (Network Data Management Protocol) over IP (Internet Protocol) or DD VTL directly over FC (Fibre Channel).

- An additional license is required for IBM i systems – the I/OS license.
- Adding a DD VTL license through the DD System Manager automatically disables and enables the DD VTL feature.
- An installed FC interface card or DD VTL configured to use NDMP.
 - If the DD VTL communication between a backup server and a DD system is through an FC interface, the DD system must have an FC interface card installed. Notice that whenever an FC interface card is removed from (or changed within) a DD system, any DD VTL configuration associated with that card must be updated.
 - If the DD VTL communication between a backup server and a DD system is through NDMP, no FC interface card is required. However, you must configure the TapeServer access group. Also, when using NDMP, all initiator and port functionality does not apply.
 - The net filter must be configured to allow the NDMP client to send information to the DD system. Run the `net filter add operation allow clients <client-IP-address>` command to allow access for the NDMP client.
 - For added security, run the `net filter add operation allow clients <client-IP-address> interfaces <DD-interface-IP-address>` command.
 - Add the `seq-id 1` option to the command to enforce this rule before any other net filter rules.
- A backup software minimum record (block) size.
 - If possible, set backup software to use a minimum record (block) size of 64 KiB or larger. Larger sizes usually give faster performance and better data compression.
 - 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 DD VTL services and perform other configuration tasks, a sysadmin login is required.

DD VTL limits

Before setting up or using a DD VTL, review these limits on size, slots, etc.

- I/O Size – The maximum supported I/O size for any DD system using DD VTL is 1 MB.
- Libraries – DD VTL supports a maximum of 64 libraries per DD system (that is, 64 DD VTL instances on each DD system).
- Initiators – DD VTL supports a maximum of 1024 initiators or WWPNs (world-wide port names) per DD system.
- Tape Drives – Information about tape drives is presented in the next section.
- Data Streams – Information about data streams is presented in the following table.

Table 130 Data streams sent to a Data Domain system

Model	RAM/ NVRAM	Backup write streams	Backup read streams	Repl ^a source streams	Repl ^a dest streams	Mixed
DD140, DD160, 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
DD620, DD630, DD640	8 GB / 0.5 GB or 1 GB	20	16	20	20	w<=20; r<=16; ReplSrc<=30; ReplDest<=20; ReplDest+w<=20; Total<=30
DD640, DD670	16 GB or 20 GB / 1 GB	90	30	60	90	w<=90; r<=30; ReplSrc<=60; 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<=90
DD860	72 GB ^b / 1 GB	90	50	90	90	w<=90; r<=50; ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; Total<=90
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 ^b / 4 GB	540	150	270	540	w<=540; r<=150; ReplSrc<=270; ReplDest<=540; ReplDest +w<=540; Total<=540
DD2200	8 GB	35	6	18	20	w<=35; r<=6; ReplSrc<=18; ReplDest<=20; ReplDest+w<=35; Total<=35
DD2200	16 GB	60	16	30	60	w<=60; r<=16; ReplSrc<=30; ReplDest<=60; ReplDest+w<=60; Total<=60
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 ^b / 4 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest +w<=270; Total<=270
DD4500	192 GB ^b / 4 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest +w<=270; Total<=270
DD7200	128 or 256 GB ^b / 4 GB	540	150	270	540	w<=540; r<=150; ReplSrc<=270; ReplDest<=540; ReplDest +w<=540; Total<=540
DD9500	256/512 GB	1885	300	540	1080	w<=1885; r<=300; ReplSrc<=540; ReplDest<=1080; ReplDest +w<=1080; Total<=1885

Table 130 Data streams sent to a Data Domain system (continued)

Model	RAM/ NVRAM	Backup write streams	Backup read streams	Repl ^a source streams	Repl ^a dest streams	Mixed
DD9800	256/768 GB	1885	300	540	1080	w<=1885; r<=300; ReplSrc<=540; ReplDest<=1080; ReplDest +w<=1080; Total<=1885
DD6300	48/96 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest +w<=270; Total<=270
DD6800	192 GB	400	110	220	400	w<=400; r<=110; ReplSrc<=220; ReplDest<=400; ReplDest +w<=400; Total<=400
DD9300	192/384 GB	800	220	440	800	w<=800; r<=220; ReplSrc<=440; ReplDest<=800; ReplDest +w<=800; Total<=800
DD VE 8 TB	8 GB / 512 MB	20	16	20	20	w<= 20 ; r<= 16 ReplSrc<=20; ReplDest<=20; ReplDest+w<=20; w+r+ReplSrc <=20;Total<=20
DD VE 16 TB	16 GB / 512 MB or 24 GB / 1 GB	45	30	45	45	w<= 45 ; r<= 30 ReplSrc<=45; ReplDest<=45; ReplDest+w<=45; w+r+ReplSrc <=45;Total<=45
DD VE 32 TB	24 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 48 TB	36 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 64 TB	48 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 96 TB	64 GB / 2 GB	180	50	90	180	w<= 180 ; r<= 50 ReplSrc<=90; ReplDest<=180; ReplDest +w<=180; w+r+ReplSrc <=180;Total<=180
DD3300 4 TB	12 GB (virtual memory) / 512 MB	20	16	30	20	w<= 20 ; r<= 16 ReplSrc<=30; ReplDest<=20; ReplDest+w<=20; w+r+ReplSrc <=30;Total<=30
DD3300 8 TB	32 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD3300 16 TB	32 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD3300 32 TB	46 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=140

Table 130 Data streams sent to a Data Domain system (continued)

- a. DirRepl, OptDup, MTreeRepl streams
- b. The Data Domain Extended Retention software option is available only for these devices with extended (maximum) memory

- Slots – DD VTL supports a maximum of:
 - 32,000 slots per library
 - 64,000 slots per DD system

The DD system automatically adds slots to keep the number of slots equal to, or greater than, the number of drives.

Note

Some device drivers (for example, IBM AIX atape device drivers) limit library configurations to specific drive/slot limits, which may be less than what the DD system supports. Backup applications, and drives used by those applications, may be affected by this limitation.

- CAPs (cartridge access ports) – DD VTL supports a maximum of:
 - 100 CAPs per library
 - 1000 CAPs per DD system

Number of drives supported by a DD VTL

The maximum number of drives supported by a DD VTL depends on the number of CPU cores and the amount of memory installed (both RAM and NVRAM, if applicable) on a DD 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 – not on the particular model, itself.

Table 131 Number of drives supported by a DD VTL

Number of CPU cores	RAM (in GB)	NVRAM (in GB)	Maximum number of supported 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

Table 131 Number of drives supported by a DD VTL (continued)

Number of CPU cores	RAM (in GB)	NVRAM (in GB)	Maximum number of supported drives
40 to 59	NA	NA	540
60 or more	NA	NA	1080

Tape barcodes

When you create a tape, you must assign a unique *barcode* (never duplicate barcodes as this can cause unpredictable behavior). Each barcode consists of eight characters: the first six are numbers or uppercase letters (0-9, A-Z), and the last two are the tape code for the supported tape type, as shown in the following table.

Note

Although a DD VTL barcode consists of eight characters, either six or eight characters may be transmitted to a backup application, depending on the changer type.

Table 132 Tape Codes by Tape Type

Tape Type	Default Capacity (unless noted)	Tape Code
LTO-1	100 GiB	L1
LTO-1	50 GiB (non-default)	LA ^a
LTO-1	30 GiB (non-default)	LB
LTO-1	10 GiB (non-default)	LC
LTO-2	200 GiB	L2
LTO-3	400 GiB	L3
LTO-4	800 GiB	L4
LTO-5 (default)	1.5 TiB	L5

a. For TSM, use the L2 tape code if the LA code is ignored.

For multiple tape libraries, barcodes are automatically incremented, if the sixth character (just before the "L") is a number. If an overflow occurs (9 to 0), numbering moves one position to the left. If the next character to increment is a letter, incrementation stops. Here are a few sample barcodes and how each will be incremented:

- 000000L1 creates tapes of 100 GiB capacity and can accept a count of up to 100,000 tapes (from 000000 to 999999).
- 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.

LTO tape drive compatibility

You may have different generations of LTO (Linear Tape-Open) technology in your setup; the compatibility between these generations is presented in tabular form.

In this table:

- RW = read and write compatible
- R = read-only compatible
- — = not compatible

Table 133 LTO tape drive compatibility

tape format	LTO-5	LTO-4	LTO-3	LTO-2	LTO-1
LTO-5	RW				
LTO-4	RW	RW	—	—	—
LTO-3	R	RW	RW	—	—
LTO-2		R	RW	RW	—
LTO-1		—	R	RW	RW

Setting up a DD VTL

To set up a simple DD VTL, use the Configuration Wizard, which is described in the *Getting Started* chapter.

Similar documentation is available in the *Data Domain Operating System Initial Configuration Guide*.

Then, continue with the following topics to enable the DD VTL, create libraries, and create and import tapes.

HA systems and DD VTL

HA systems are compatible with DD VTL; however, if a DD VTL job is in progress during a failover, the job will need to be restarted manually after the failover is complete.

The *Data Domain Operating System Backup Compatibility Guide* provides additional details about the HBA, switch, firmware, and driver requirements for using DD VTL in an HA environment.

DD VTL tape out to cloud

DD VTL supports storing the VTL vault on DD Cloud Tier storage. To use this functionality, the Data Domain system must be a supported Cloud Tier configuration, and have a Cloud Tier license in addition to the VTL license.

Configure and license the DD Cloud Tier storage before configuring DD VTL to use cloud storage for the vault. [DD Cloud Tier](#) on page 457 provides additional information about the requirements for DD Cloud Tier, and how to configure DD Cloud Tier.

The FC and network interface requirements for VTL are the same for both cloud-based and local vault storage. DD VTL does not require special configuration to use cloud storage for the vault. When configuring the DD VTL, select the cloud storage as the vault location. However, when working with a cloud-based vault, there are some data management options that are unique to the cloud-based vault. [Working with the cloud-based vault](#) on page 373 provides more information.

Managing a DD VTL

You can manage a DD VTL using the Data Domain System Manager (DD System Manager) or the Data Domain Operating System (DD OS) Command Line Interface (CLI). After you login, you can check the status of your DD VTL process, check your license information, and review and configure options.

Logging In

To use a graphical user interface (GUI) to manage your DD Virtual Tape Library (DD VTL), log in to the DD System Manager.

CLI Equivalent

You can also log in at the CLI:

```
login as: sysadmin
Data Domain OS
Using keyboard-interactive authentication.
Password:
```

Enabling SCSI Target Daemon (CLI only)

If you do log in from the CLI, you must enable the `scsitarget` daemon (the Fibre Channel service). This daemon is enabled during the DD VTL or DD Boost-FC enable selections in DD System Manager. In the CLI, these processes need to be enabled separately.

```
# scsitarget enable
Please wait ...
SCSI Target subsystem is enabled.
```

Accessing DD VTL

From the menu at the left of the DD System Manager, select **Protocols > VTL**.

Status

In the **Virtual Tape Libraries > VTL Service** area, you can see the status of your DD VTL process is displayed at the top, for example, **Enabled: Running**. The first part of the status will be **Enabled** (on) or **Disabled** (off). The second part will be one of the following process states.

Table 134 DD VTL process states

State	Description
Running	DD VTL process is enabled and active (shown in green).
Starting	DD VTL process is starting.
Stopping	DD VTL process is being shut down.
Stopped	DD VTL process is disabled (shown in red).
Timing out	DD VTL process crashed and is attempting an automatic restart.
Stuck	After several failed automatic restarts, the DD VTL process is unable to shut down normally, so an attempt is being made to kill it.

DD VTL License

The VTL License line tells you whether your DD VTL license has been applied. If it says Unlicensed, select **Add License**. Enter your license key in the Add License Key dialog. Select **Next** and **OK**.

Note

All license information should have been populated as part of the factory configuration process; however, if DD VTL was purchased later, the DD VTL license key may not have been available at that time.

CLI Equivalent

You can also verify that the DD VTL license has been installed at the CLI:

```
# license show
## License Key                               Feature
-----
1      DEFA-EFCD-FCDE-CDEF                    Replication
2      EFCD-FCDE-CDEF-DEFA                    VTL
-----
```

If the license is not present, each unit comes with documentation – a quick install card – which will show the licenses that have been purchased. Enter the following command to populate the license key.

```
# license add license-code
```

I/OS License (for IBM i users)

For customers of IBM i, the I/OS License line tells you whether your I/OS license has been applied. If it says Unlicensed, select **Add License**. You must enter a valid I/OS license in either of these formats: `xxxx-xxxx-xxxx-xxxx` or `xxxx-xxxx-xxxx-xxxx-xxxx`. Your I/OS license must be installed before creating a library and drives to be used on an IBM i system. Select **Next** and **OK**.

Enabling DD VTL

Enabling DD VTL broadcasts the WWN of the Data Domain HBA to customer fabric and enables all libraries and library drives. If a forwarding plan is required in the form of change control processes, this process should be enabled to facilitate zoning.

Procedure

1. Make sure that you have a DD VTL license and that the file system is enabled.

2. Select **Virtual Tape Libraries > VTL Service**.
3. To the right of the Status area, select **Enable**.
4. In the Enable Service dialog, select **OK**.
5. After DD VTL has been enabled, note that Status will change to **Enabled: Running** in green. Also note that the configured DD VTL options are displayed in the Option Defaults area.

CLI Equivalent

```
# vtl enable
Starting VTL, please wait ...
VTL is enabled.
```

Disabling DD VTL

Disabling DD VTL closes all libraries and shuts down the DD VTL process.

Procedure

1. Select **Virtual Tape Libraries > VTL Service**.
2. To the right of the Status area, select **Disable**.
3. In the Disable Service dialog, select **OK**.
4. After DD VTL has been disabled, notice that the Status has changed to **Disabled: Stopped** in red.

CLI Equivalent

```
# vtl disable
```

DD VTL option defaults

The Option Default area of the VTL Service page displays the current settings for default DD VTL options (auto-eject, auto-offline, and barcode-length) that you can configure.

In the **Virtual Tape Libraries > VTL Service** area, the current default options for your DD VTL are displayed. Select **Configure** to change any of these values.

Table 135 Option Defaults

Item	Description
Property	Lists the configured options: <ul style="list-style-type: none"> • auto-eject • auto-offline • barcode-length
Value	Provides the value for each configured option: <ul style="list-style-type: none"> • auto-eject: default (disabled), enabled, or disabled • auto-offline: default (disabled), enabled, or disabled • barcode-length: default (8), 6, or 8

Configuring DD VTL default options

You can configure DD VTL default options when you add a license, create a library, or any time thereafter.

Note

DD VTLs are assigned global options, by default, and those options are updated whenever global options change, unless you change them manually using this method.

Procedure

1. Select **Virtual Tape Libraries > VTL Service**.
2. In the Option Defaults area, select **Configure**. In the Configure Default Options dialog, change any or all of the default options.

Table 136 DD VTL default options

Option	Values	Notes
auto-eject	default (disabled), enable, or disable	Enabling auto-eject causes any tape put into a CAP (cartridge access port) to automatically move to the virtual vault, unless: <ul style="list-style-type: none"> • the tape came from the vault, in which case the tape stays in the CAP. • an <code>ALLOW_MEDIUM_REMOVAL</code> 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.
auto-offline	default (disabled), enable, or disable	Enabling auto-offline takes a drive offline automatically before a tape move operation is performed.
barcode-length	default (8), 6 or 8 [automatically set to 6 for L180, RESTORER-L180, and DDVTL changer models]	Although a DD VTL barcode consists of 8 characters, either 6 or 8 characters may be transmitted to a backup application, depending on the changer type.

3. Select **OK**.
4. Or to disable all of these service options, select **Reset to Factory**, and the values will be immediately reset to factory defaults.

Working with libraries

A library emulates a physical tape library with drives, changer, CAPs (cartridge access ports), and slots (cartridge slots). Selecting **Virtual Tape Libraries > VTL Service > Libraries** displays detailed information for all configured libraries.

Table 137 Library information

Item	Description
Name	The name of a configured library.
Drives	The number of drives configured in the library.
Slots	The number of slots configured in the library.
CAPs	The number of CAPs (cartridge access ports) configured in the library.

From the More Tasks menu, you can create and delete libraries, as well as search for tapes.

Creating libraries

DD VTL supports a maximum of 64 libraries per system, that is, 64 concurrently active virtual tape library instances on each DD system.

Procedure

1. Select **Virtual Tape Libraries > VTL Service > Libraries**.
2. Select **More Tasks > Library > Create**
3. In the Create Library dialog, enter the following information:

Table 138 Create Library dialog

Field	User input
Library Name	Enter a name of from 1 to 32 alphanumeric characters.
Number of Drives	Enter the number of drives (from 1 to 98 (see Note). The number of drives to be created will correspond to the number of data streams that will write to a library.
	<p>Note</p> <p>The maximum number of drives supported by a DD VTL depends on the number of CPU cores and the amount of memory installed (both RAM and NVRAM, if applicable) on a DD system.</p>
Drive Model	<p>Select the desired model from the drop-down list:</p> <ul style="list-style-type: none"> • IBM-LTO-1 • IBM-LTO-2 • IBM-LTO-3 • IBM-LTO-4 • IBM-LTO-5 (default)

Table 138 Create Library dialog (continued)

Field	User input
	<ul style="list-style-type: none"> • HP-LTO-3 • HP-LTO-4 <p>Do not mix drive types, or media types, in the same library. This can cause unexpected results and/or errors in the backup operation.</p>
Number of Slots	<p>Enter the number of slots in the library. Here are some things to consider:</p> <ul style="list-style-type: none"> • The number of slots must be equal to or greater than the number of drives. • You can have up to 32,000 slots per individual library • You can have up to 64,000 slots per system. • Try to have enough slots so tapes remain in the DD VTL and never need to be exported to a vault – to avoid reconfiguring the DD VTL and to ease management overhead. • Consider any applications that are licensed by the number of slots. <p>As an example, for a standard 100-GB cartridge on a DD580, you might configure 5000 slots. This would be enough to hold up to 500 TB (assuming reasonably compressible data).</p>
Number of CAPs	<p>(Optional) Enter the number of cartridge access ports (CAPs).</p> <ul style="list-style-type: none"> • You can have up to 100 CAPs per library. • You can have up to 1000 CAPs per system. <p>Check your particular backup software application documentation on the Online Support Site for guidance.</p>
Changer Model Name	<p>Select the desired model from the drop-down list:</p> <ul style="list-style-type: none"> • L180 (default) • RESTORER-L180 • TS3500 (which should be used for IBMi deployments) • I2000 • I6000 • DDVTL <p>Check your particular backup software application documentation on the Online Support Site for guidance. Also refer to the DD VTL support matrix to see the compatibility of emulated libraries to supported software.</p>
Options	
auto-eject	default (disabled), enable, disable
auto-offline	default (disabled), enable, disable
barcode-length	default (8), 6, 8 [automatically set to 6 for L180, RESTORER-L180, and DDVTL changer models]

4. 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 that access to VTLs and drives is managed with Access Groups.

CLI Equivalent

```
# vt1 add NewVTL model L180 slots 50 caps 5
This adds the VTL library, NewVTL. Use 'vt1 show config NewVTL'
to view it.

# vt1 drive add NewVTL count 4 model IBM-LTO-3
This adds 4 IBM-LTO-3 drives to the VTL library, NewVTL.
```

Deleting libraries

When a tape is in a drive within a library, and that library is deleted, the tape is moved to the vault. However, the tape's pool does not change.

Procedure

1. Select **Virtual Tape Libraries > VTL Service > Libraries**.
2. Select **More Tasks > Library > Delete**.
3. In the Delete Libraries dialog, select or confirm the checkbox of the items to delete:
 - The name of each library, or
 - Library Names, to delete all libraries
4. Select **Next**.
5. Verify the libraries to delete, and select **Submit** in the confirmation dialogs.
6. After the Delete Libraries Status dialog shows **Completed**, select **Close**. The selected libraries are deleted from the DD VTL.

CLI Equivalent

```
# vt1 del OldVTL
```

Searching for tapes

You can use a variety of criteria – location, pool, and/or barcode – to search for a tape.

Procedure

1. Select **Virtual Tape Libraries** or **Pools**.
2. Select the area to search (library, vault, pool).
3. Select **More Tasks > Tapes > Search**.
4. In the Search Tapes dialog, enter information about the tape(s) you want to find.

Table 139 Search Tapes dialog

Field	User input
Location	Specify a location, or leave the default (All).
Pool	Select the name of the pool in which to search for the tape. If no pools have been created, use the Default pool.
Barcode	Specify a unique barcode. or leave the default (*) to return a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.

5. Select **Search**.

Working with a selected library

Selecting **Virtual Tape Libraries > VTL Service > Libraries > library** displays detailed information for a selected library.

Table 140 Devices

Item	Description
Device	The elements in the library, such a drives, slots, and CAPs (cartridge access ports).
Loaded	The number of devices with media loaded.
Empty	The number of devices with no media loaded.
Total	The total number of loaded and empty devices.

Table 141 Options

Property	Value
auto-eject	enabled or disabled
auto-offline	enabled or disabled
barcode-length	6 or 8

Table 142 Tapes

Item	Description
Pool	The name of the pool where the tapes are located.
Tape Count	The number of tapes in that pool.
Capacity	The total configured data capacity of the tapes in that pool, in GiB (Gibibytes, the base-2 equivalent of GB, Gigabytes).
Used	The amount of space used on the virtual tapes in that pool.

Table 142 Tapes (continued)

Item	Description
Average Compression	The average amount of compression achieved on the data on the tapes in that pool.

From the More Tasks menu, you can delete, rename, or set options for a library; create, delete, import, export, or move tapes; and add or delete slots and CAPs.

Creating tapes

You can create tapes in either a library or a pool. If initiated from a pool, the system first creates the tapes, then imports them to the library.

Procedure

1. Select **Virtual Tape Libraries > VTL Service > Libraries > *library* or Vault or Pools > Pools > *pool***.
2. Select **More Tasks > Tapes > Create**.
3. In the Create Tapes dialog, enter the following information about the tape:

Table 143 Create Tapes dialog

Field	User input
Library (if initiated from a library)	If a drop-down menu is enabled, select the library or leave the default selection.
Pool Name	Select the name of the pool in which the tape will reside, from the drop-down list. If no pools have been created, use the Default pool.
Number of Tapes	For a library, select from 1 to 20. For a pool, select from 1 to 100,000, or leave the default (20). [Although the number of supported tapes is unlimited, you can create no more than 100,000 tapes at a time.]
Starting Barcode	Enter the initial barcode number (using the format A99000LA).
Tape Capacity	(optional) Specify the number of GiBs from 1 to 4000 for each tape (this setting overrides the barcode capacity setting). For efficient use of disk space, use 100 GiB or fewer.

4. Select **OK** and **Close**.

CLI Equivalent

```
# vt1 tape add A00000L1 capacity 100 count 5 pool VTL_Pool
... added 5 tape(s)...
```

Note

You must auto-increment tape volume names in base10 format.

Deleting tapes

You can delete tapes from either a library or a pool. If initiated from a library, the system first exports the tapes, then deletes them. The tapes must be in the vault, not in a library. On a Replication destination DD system, deleting a tape is not permitted.

Procedure

1. Select **Virtual Tape Libraries > VTL Service > Libraries > *library*** or **Vault or Pools > Pools > *pool***.
2. Select **More Tasks > Tapes > Delete**.
3. In the Delete Tapes dialog, enter search information about the tapes to delete, and select **Search**:

Table 144 Delete Tapes dialog

Field	User input
Location	If there is a drop-down list, select a library, or leave the default Vault 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.
Barcode	Specify a unique barcode, or leave the default (*) to search for a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.
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 Select All Pages checkbox to select all tapes returned by the search query.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

4. Select the checkbox of the tape that should be deleted or the checkbox on the heading column to delete all tapes, and select **Next**.
5. Select **Submit** in the confirmation window, and select **Close**.

Note

After a tape is removed, the physical disk space used for the tape is not reclaimed until after a file system cleaning operation.

CLI Equivalent

```
# vt1 tape del barcode [count count] [pool pool]
```

For example:

```
# vt1 tape del A00000L1
```

Note

You can act on ranges; however, if there is a missing tape in the range, the action will stop.

Importing tapes

Importing a tape means that an existing tape will be moved 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, that is, 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, a slot is reserved.
- If a tape is in a drive, and the tape origin is unknown (slot or CAP), a slot is reserved.
- If a tape is in a drive, and the tape origin is known to be a CAP, a slot is not reserved. (The tape returns to the CAP when removed from the drive.)
- To move a tape to a drive, see the section on moving tapes, which follows.

Procedure

1. You can import tapes using either step a. or step b.
 - a. Select **Virtual Tape Libraries > VTL Service > Libraries > library**. Then, select **More Tasks > Tapes > Import**. In the Import Tapes dialog, enter search information about the tapes to import, and select **Search**:

Table 145 Import Tapes dialog

Field	User input
Location	If there is a drop-down list, select the location of the tape, or leave the default of Vault .
Pool	Select the name of the pool in which to search for the tape. If no pools have been created, use the Default pool.
Barcode	Specify a unique barcode. or leave the default (*) to return a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

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 from which to choose.

b. Select **Virtual Tape Libraries > VTL Service > Libraries > library > Changer > Drives > drive > Tapes**. Select tapes to import by selecting the checkbox next to:

- An individual tape, or
- 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.

Only tapes showing Vault in the Location can be imported.

Select **Import from Vault**. This button is disabled by default and enabled only if all of the 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.

CLI Equivalent

```
# vtl tape show pool VTL_Pool
Processing tapes....
Barcode Pool Location State Size Used (%) Comp ModTime
-----
A00000L3 VTL_Pool vault RW 100 GiB 0.0 GiB (0.00%) 0x 2010/07/16 09:50:41
A00001L3 VTL_Pool vault RW 100 GiB 0.0 GiB (0.00%) 0x 2010/07/16 09:50:41
A00002L3 VTL_Pool vault RW 100 GiB 0.0 GiB (0.00%) 0x 2010/07/16 09:50:41
A00003L3 VTL_Pool vault RW 100 GiB 0.0 GiB (0.00%) 0x 2010/07/16 09:50:41
A00004L3 VTL_Pool vault RW 100 GiB 0.0 GiB (0.00%) 0x 2010/07/16 09:50:41
-----
VTL Tape Summary
-----
Total number of tapes: 5
Total pools: 1
Total size of tapes: 500 GiB
Total space used by tapes: 0.0 GiB
Average Compression: 0.0x

# vtl import NewVTL barcode A00000L3 count 5 pool VTL_Pool
... imported 5 tape(s)...

# vtl tape show pool VTL_Pool
Processing tapes....

VTL Tape Summary
-----
Total number of tapes: 5
Total pools: 1
Total size of tapes: 500 GiB
Total space used by tapes: 0.0 GiB
Average Compression: 0.0x
```

Exporting tapes

Exporting a tape removes that tape from a slot, drive, or cartridge-access port (CAP) and sends it to the vault.

Procedure

1. You can export tapes using either step a. or step b.
 - a. Select **Virtual Tape Libraries > VTL Service > Libraries > library**. Then, select **More Tasks > Tapes > Export**. In the Export Tapes dialog, enter search information about the tapes to export, and select **Search**:

Table 146 Export Tapes dialog

Field	User input
Location	If there is a drop-down list, select the name of the library where the tape is located, or leave the selected library.
Pool	Select the name of the pool in which to search for the tape. If no pools have been created, use the Default pool.
Barcode	Specify a unique barcode. or leave the default (*) to return a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.
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 Select All Pages checkbox to select all tapes returned by the search query.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

b. Select **Virtual Tape Libraries > VTL Service > Libraries > library > Changer > Drives > drive > Tapes**. Select tapes to export by selecting the checkbox next to:

- An individual tape, or
- 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.

Only tapes with a library name in the Location column can be exported.

Select **Export from Library**. This button is disabled by default and enabled only if all of the selected tapes have a library name in the Location column.

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 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. Select **Virtual Tape Libraries > VTL Service > Libraries > library**.

Note that when started from a library, the Tapes panel allows tapes to be moved only between devices.

2. Select **More Tasks > Tapes > Move**.

Note that when started from a library, the Tapes panel allows tapes to be moved only between devices.

3. In the Move Tape dialog, enter search information about the tapes to move, and select **Search**:

Table 147 Move Tape dialog

Field	User input
Location	Location cannot be changed.
Pool	N/A
Barcode	Specify a unique barcode. or leave the default (*) to return a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

4. From the search results list, select the tape or tapes to move.
5. 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. 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.
6. Select **Next**.
7. In the Move Tape dialog, verify the summary information and the tape listing, and select **Submit**.
8. Select **Close** in the status window.

Adding slots

You can add slots from a configured library to change the number of storage elements.

Note

Some backup applications do not automatically recognize that slots have been added to a DD VTL. See your application documentation for information on how to configure the application to recognize this type of change.

Procedure

1. Select **Virtual Tape Libraries > VTL Service > Libraries > library**.
2. Select **More Tasks > Slots > Add**.

3. In the Add Slots dialog, enter the Number of Slots to add. The total number of slots in a library, or in all libraries on a system, cannot exceed 32,000 for a library and 64,000 for a system.
4. Select **OK** and **Close** when the status shows *Completed*.

Deleting slots

You can delete slots from a configured library to change the number of storage elements.

Note

Some backup applications do not automatically recognize that slots have been deleted from a DD VTL. See your application documentation for information on how to configure the application to recognize this type of change.

Procedure

1. If the slot that you want to delete contains cartridges, move those cartridges to the vault. The system will delete only empty, uncommitted slots.
2. Select **Virtual Tape Libraries > VTL Service > Libraries > *library***.
3. Select **More Tasks > Slots > Delete**.
4. In the Delete Slots dialog, enter the Number of Slots to delete.
5. Select **OK** and **Close** when the status shows *Completed*.

Adding CAPs

You can add CAPs (cartridge access ports) from a configured library to change the number of storage elements.

Note

CAPs are used by a limited number of backup applications. See your application documentation to ensure that CAPs are supported.

Procedure

1. Select **Virtual Tape Libraries > VTL Service > Libraries > *library***.
2. Select **More Tasks > CAPs > Add**.
3. In the Add CAPs dialog, enter the Number of CAPs to add. You can add from 1 to 100 CAPs per library and from 1 to 1,000 CAPs per system.
4. Select **OK** and **Close** when the status shows *Completed*.

Deleting CAPs

You can delete CAPs (cartridge access ports) from a configured library to change the number of storage elements.

Note

Some backup applications do not automatically recognize that CAPs have been deleted from a DD VTL. See your application documentation for information on how to configure the application to recognize this type of change.

Procedure

1. If the CAP that you want to delete contains cartridges, move those cartridges to the vault, or this will be done automatically.
2. Select **Virtual Tape Libraries > VTL Service > Libraries > library**.
3. Select **More Tasks > CAPs > Delete**.
4. In the Delete CAPs dialog, enter the Number of CAPs to delete. You can delete a maximum of 100 CAPs per library or 1000 CAPs per system.
5. Select **OK** and **Close** when the status shows *Completed*.

Viewing changer information

There can be only one changer per DD VTL. The changer model you select depends on your specific configuration.

Procedure

1. Select **Virtual Tape Libraries > VTL Service > Libraries**.
2. Select a specific library.
3. If not expanded, select the plus sign (+) on the left to open the library, and select a Changer element to display the Changer information panel, which provides the following information.

Table 148 Changer information panel

Item	Description
Vendor	The name of the vendor who manufactured the changer
Product	The model name
Revision	The revision level
Serial Number	The changer serial number

Working with drives

Selecting **Virtual Tape Libraries > VTL Service > Libraries > library > Drives** displays detailed information for all drives for a selected library.

Table 149 Drives information panel

Column	Description
Drive	The 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	The manufacturer or vendor of the drive, for example, IBM.
Product	The product name of the drive, for example, ULTRIUM-TD5.
Revision	The revision number of the drive product.
Serial Number	The serial number of the drive product.

Table 149 Drives information panel (continued)

Column	Description
Status	Whether the drive is Empty, Open, Locked, or Loaded. A tape must be present for the drive to be locked or loaded.
Tape	The barcode of the tape in the drive (if any).
Pool	The pool of the tape in the drive (if any).

Tape and library drivers – To work with 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, IBM LTO-4, IBM LTO-5 (default), HP-LTO-3, or HP-LTO-4 drives and the StorageTek L180 (default), RESTORER-L180, IBM TS3500, I2000, I6000, or DDVTL libraries. For more information, see the *Application Compatibility Matrices and Integration Guides* for your vendors. When configuring drives, also keep in mind the limits on backup data streams, which are determined by the platform in use.

LTO drive capacities – Because the DD system treats 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 GiB
- LTO-2 drive: 200 GiB
- LTO-3 drive: 400 GiB
- LTO-4 drive: 800 GiB
- LTO-5 drive: 1.5 TiB

Migrating LTO-1 tapes – 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 LTO tape migration guide specific to your application. To find the appropriate guide, go to the Online Support Site, and in the search text box, type in **LTO Tape Migration for VTLs**.

Tape full: Early warning – You will 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.

From the More Tasks menu, you can create or delete a drive.

Creating drives

See the *Number of drives supported by a DD VTL* section to determine the maximum number of drives supported for your particular DD VTL.

Procedure

1. Select **Virtual Tape Libraries > VTL Service > Libraries > *library* > Changer > Drives**.
2. Select **More Tasks > Drives > Create**.
3. In the Create Drive dialog, enter the following information:

Table 150 Create Drive dialog

Field	User input
Location	Select a library name, or leave the name selected.
Number of Drives	See the table in the <i>Number of Drives Supported by a DD VTL</i> section, earlier in this chapter.
Model Name	Select the model from the drop-down list. If another drive already exists, this option is inactive, and the existing drive type must be used. You cannot mix drive types in the same library. <ul style="list-style-type: none"> • IBM-LTO-1 • IBM-LTO-2 • IBM-LTO-3 • IBM-LTO-4 • IBM-LTO-5 (default) • HP-LTO-3 • HP-LTO-4

4. Select **OK**, and when the status shows *Completed*, select **OK**.

The added drive appears in the Drives list.

Deleting drives

A drive must be empty before it can be deleted.

Procedure

1. If there is a tape in the drive that you want to delete, remove the tape.
2. Select **Virtual Tape Libraries > VTL Service > Libraries > library > Changer > Drives**.
3. Select **More Tasks > Drives > Delete**.
4. In the Delete Drives dialog, select the checkboxes of the drives to delete, or select the **Drive** checkbox to delete all drives.
5. Select **Next**, and after verifying that the correct drive(s) has been selected for deletion, select **Submit**.
6. When the Delete Drive Status dialog shows *Completed*, select **Close**.

The drive will have been removed from the Drives list.

Working with a selected drive

Selecting **Virtual Tape Libraries > VTL Service > Libraries > library > Drives > drive** displays detailed information for a selected drive.

Table 151 Drive Tab

Column	Description
Drive	The 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	The manufacturer or vendor of the drive, for example, IBM.
Product	The product name of the drive, for example, ULTRIUM-TD5.
Revision	The revision number of the drive product.
Serial Number	The serial number of the drive product.
Status	Whether the drive is Empty, Open, Locked, or Loaded. A tape must be present for the drive to be locked or loaded.
Tape	The barcode of the tape in the drive (if any).
Pool	The pool of the tape in the drive (if any).

Table 152 Statistics Tab

Column	Description
Endpoint	The specific name of the endpoint.
Ops/s	The operations per second.
Read KiB/s	The speed of reads in KiB per second.
Write KiB/s	The speed of writes in KiB per second.

From the More Tasks menu, you can delete the drive or perform a refresh.

Working with tapes

A tape is represented as a file. 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 (cartridge slots), and CAPs (cartridge access ports).

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, or removed.

Selecting **Virtual Tape Libraries > VTL Service > Libraries > library > Tapes** displays detailed information for all tapes for a selected library.

Table 153 Tape description

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"> • RW – Read-writable • RL – Retention-locked • RO – Readable only • WP – Write-protected • RD – Replication destination
Capacity	The total capacity of the tape.
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. 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 DD System Manager.
Locked Until	If a DD Retention Lock deadline has been set, the time set is shown. If no retention lock exists, this value is <code>Not specified</code> .

From the information panel, you can import a tape from the vault, export a tape to the library, set a tape's state, create a tape, or delete a tape.

From the More Tasks menu, you can move a tape.

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. DD 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**.
2. 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.

Working with the vault

The vault holds tapes not being used by any library. Tapes reside in either a library or the vault.

Selecting **Virtual Tape Libraries > VTL Service > Vault** displays detailed information for the Default pool and any other existing pools in the vault.

Systems with DD Cloud Tier and DD VTL provide the option of storing the vault on cloud storage.

Table 154 Pool Summary

Item	Description
Pool Count	The number of VTL pools.
Tape Count	The number of tapes in the pools.
Size	The total amount of space in the pools.
Logical Used	The amount of space used in the pools.
Compression	The average amount of compression in the pools.

The **Protection Distribution** pane displays the following information.

Table 155 Protection Distribution

Item	Description
Storage type	Vault or Cloud.
Cloud provider	For systems with tapes in DD Cloud Tier, there is a column for each cloud provider.
Logical Used	The amount of space used in the pools.
Pool Count	The number of VTL pools.
Tape Count	The number of tapes in the pools.

From the More Tasks menu, you can create, delete, and search for tapes in the vault.

Working with the cloud-based vault

DD VTL supports several parameters that are unique to configurations where the vault is stored on DD Cloud Tier storage.

The following operations are available for working with cloud-based vault storage.

- Configure the data movement policy and cloud unit information for the specified VTL pool. Run the `vtl pool modify <pool-name> data-movement-policy {user-managed | age-threshold <days> | none} to-tier {cloud} cloud-unit <cloud-unit-name>` command.

The available data movement policies are:

- **User-managed:** The administrator can set this policy on a pool, to manually select tapes from the pool for migration to the cloud tier. The tapes migrate to

the cloud tier on the first data movement operation after the tapes are selected.

- **Age-threshold:** The administrator can set this policy on a pool, to allow the DD VTL to automatically select tapes from the pool for migration to the cloud tier based on the age of the tape. The tapes are selected for migration within six hours after they meet the age threshold, and are migrated on the first data movement operation after the tapes are selected.
- Select a specified tape for migration to the cloud tier. Run the `vtl tape select-for-move barcode <barcode> [count <count>] pool <pool> to-tier {cloud}` command.
- Deselect a specified tape for migration to the cloud tier. Run the `vtl tape deselect-for-move barcode <barcode> [count <count>] pool <pool> to-tier {cloud}` command.
- Recall a tape from the cloud tier. Run the `vtl tape recall start barcode <barcode> [count <count>] pool <pool>` command. After the recall, the tape resides in a local DD VTL vault and must be imported to the library for access.

Note

Run the `vtl tape show` command at any time to check the current location of a tape. The tape location updates within one hour of the tape moving to or from the cloud tier.

Prepare the VTL pool for data movement

Set the data movement policy on the VTL pool to manage migration of VTL data from the local vault to DD Cloud Tier.

Data movement for VTL occurs at the tape volume level. Individual tape volumes or collections of tape volumes can be moved to the cloud tier but only from the vault location. Tapes in other elements of a VTL cannot be moved.

Note

The default VTL pool and vault, `/data/coll/backup` directories or legacy library configurations cannot be used for Tape out to Cloud.

Procedure

1. Select **Protocols > DD VTL**.
2. Expand the list of pools, and select a pool on which to enable migration to DD Cloud Tier.
3. In the **Cloud Data Movement** pane, click **Create** under **Cloud Data Movement Policy**.
4. In the **Policy** drop-down list, select a data movement policy:
 - **Age of tapes in days**
 - **Manual selection**
5. Set the data movement policy details.
 - For **Age of tapes in days**, select an age threshold after which tapes are migrated to DD Cloud Tier, and specify a destination cloud unit.

- For **Manual selection**, specify a destination cloud unit.
6. Click **Create**.

Note

After creating the data movement policy, the **Edit** and **Clear** buttons can be used to modify or delete the data movement policy.

CLI equivalent

Procedure

1. Set the data movement policy to user-managed or age-threshold

Note

VTL pool and cloud unit names are case sensitive and commands will fail if the case is not correct.

- To set the data movement policy to user-managed, run the following command:

```
vtl pool modify cloud-vtl-pool data-movement-policy
user-managed to-tier cloud cloud-unit ecs-unit1
```

** Any tapes that are already selected will be migrated on the next data-movement run.
VTL data-movement policy is set to "user-managed" for VTL pool "cloud-vtl-pool".

- To set the data movement policy to age-threshold, run the following command:

Note

The minimum is 14 days, and the maximum is 182,250 days.

```
vtl pool modify cloud-vtl-pool data-movement-policy age-
threshold 14 to-tier cloud cloud-unit ecs-unit1
```

** Any tapes that are already selected will be migrated on the next data-movement run.
VTL data-movement policy "age-threshold" is set to 14 days for the VTL pool "cloud-vtl-pool".

2. Verify the data movement policy for the VTL pool.

Run the following command:

```
vtl pool show all
```

```
VTL Pools
Pool          Status  Tapes  Size (GiB)  Used (GiB)  Comp  Cloud Unit
Cloud Policy
-----
cloud-vtl-pool  RW      50     250         41         45x   ecs-unit1
user-managed
Default
none          RW       0       0           0           0x    -
-----
8080 tapes in 5 pools

RO  : Read Only
RD  : Replication Destination
BCM : Backwards-Compatibility
```

3. Verify the policy for the VTL pool MTree is `app-managed`.

Run the following command:

```
data-movement policy show all
```

Mtree	Target (Tier/Unit Name)	Policy	Value
/data/coll/cloud-vtl-pool	Cloud/ecs-unit1	app-managed	enabled

Remove tapes from the backup application inventory

Use the backup application verify the tape volumes that will move to the cloud are marked and inventoried according to the backup application requirements.

Select tape volumes for data movement

Manually select tapes for migration to DD Cloud Tier (immediately or at the next scheduled data migration), or manually remove tapes from the migration schedule.

Before you begin

Verify the backup application is aware of status changes for volumes moved to cloud storage. Complete the necessary steps for the backup application to refresh its inventory to reflect the latest volume status.

If the tape is not in the vault, it cannot be migrated to DD Cloud Tier.

Procedure

1. Select **Protocols > DD VTL**.
2. Expand the list of pools, and select the pool which is configured to migrate tapes to DD Cloud Tier.
3. In the pool pane, click the **Tape** tab.
4. Select tapes for migration to DD Cloud Tier.
5. Click **Select for Cloud Move** to migrate the tape at the next scheduled migration, or **Move to Cloud Now** to immediately migrate the tape.

Note

If the data movement policy is based on tape ages, the **Select for Cloud Move** is not available, as the Data Domain system automatically selects tapes for migration.

6. Click **Yes** at the confirmation dialog.

Unselect tape volumes for data movement

Tapes selected for migration to DD Cloud Tier can be removed from the migration schedule.

Procedure

1. Select **Protocols > DD VTL**.
2. Expand the list of pools, and select the pool which is configured to migrate tapes to DD Cloud Tier.
3. In the pool pane, click the **Tape** tab.
4. Select tapes for migration to DD Cloud Tier.

5. Click **Unselect Cloud Move** to remove the tape from the migration schedule.
6. Click **Yes** at the confirmation dialog.

CLI equivalent

Procedure

1. Identify the slot location of the tape volume to move.

Run the following command:

```
vtl tape show cloud-vtl
```

```
Processing tapes....
Barcode      Pool          Location      State   Size      Used (%)
Comp  Modification Time
-----  -
T00001L3    cloud-vtl-pool  cloud-vtl slot 1  RW     5 GiB    5.0 GiB (99.07%)
205x      2017/05/05 10:43:43
T00002L3    cloud-vtl-pool  cloud-vtl slot 2  RW     5 GiB    5.0 GiB (99.07%)
36x      2017/05/05 10:45:10
T00003L3    cloud-vtl-pool  cloud-vtl slot 3  RW     5 GiB    5.0 GiB (99.07%)
73x      2017/05/05 10:45:26
```

2. Specify the numeric slot value to export the tape from the DD VTL.

Run the following command:

```
vtl export cloud-vtl-pool slot 1 count 1
```

3. Verify the tape is in the vault.

Run the following command:

```
vtl tape show vault
```

4. Select the tape for data movement.

Run the following command:

```
vtl tape select-for-move barcode T00001L3 count 1 pool
cloud-vtl-pool to-tier cloud
```

Note

If the data movement policy is age-threshold, data movement occurs automatically after 15-20 minutes.

5. View the list of tapes scheduled to move to cloud storage during the next data movement operation. The tapes selected for movement display an (S) in the location column.

Run the following command:

```
vtl tape show vault
```

```
Processing tapes.....
Barcode      Pool          Location      State   Size      Used (%)  Comp
Modification Time
-----  -
T00003L3    cloud-vtl-pool  vault (S)    RW     5 GiB    5.0 GiB (99.07%)  63x
2017/05/05 10:43:43
T00006L3    cloud-vtl-pool  ecs-unit1   n/a     5 GiB    5.0 GiB (99.07%)  62x
2017/05/05 10:45:49
-----  -
```

* RD : Replication Destination

(S) Tape selected for migration to cloud. Selected tapes will move to cloud on the next data-movement run.

(R) Recall operation is in progress for the tape.

VTL Tape Summary

```
-----
Total number of tapes:      4024
Total pools:                3
Total size of tapes:        40175 GiB
Total space used by tapes:  39.6 GiB
Average Compression:        9.7x
```

6. If the data movement policy is user-managed, initiate the data movement operation.

Run the following command:

```
data-movement start
```

7. Observe the status of the data movement operation.

Run the following command:

```
data-movement watch
```

8. Verify the tape volumes successfully move to cloud storage.

Run the following command:

```
vtl tape show all cloud-unit ecs-unit1
```

```
Processing tapes.....
```

Barcode	Pool	Location	State	Size	Used (%)	Comp	Modification Time
T00001L3	cloud-vtl-pool	ecs-unit1	n/a	5 GiB	5.0 GiB (99.07%)	89x	2017/05/05 10:41:41
T00006L3	cloud-vtl-pool	ecs-unit1	n/a	5 GiB	5.0 GiB (99.07%)	62x	2017/05/05 10:45:49

```
(S) Tape selected for migration to cloud. Selected tapes will move to cloud on the next data-movement run.
```

```
(R) Recall operation is in progress for the tape.
```

VTL Tape Summary

```
-----
Total number of tapes:      4
Total pools:                2
Total size of tapes:        16 GiB
Total space used by tapes:  14.9 GiB
Average Compression:        59.5x
```

Restore data held in the cloud

When a client requests data for restore from the backup application server, the backup application should generate an alert or message requesting the required volumes from the cloud unit.

The volume must be recalled from the cloud and checked into the Data Domain VTL library before the backup application must be notified of the presence of the volumes.

Note

Verify the backup application is aware of status changes for volumes moved to cloud storage. Complete the necessary steps for the backup application to refresh its inventory to reflect the latest volume status.

Manually recall a tape volume from cloud storage

Recall a tape from DD Cloud Tier to the local VTL vault.

Procedure

1. Select **Protocols > DD VTL**.

2. Expand the list of pools, and select the pool which is configured to migrate tapes to DD Cloud Tier.
3. In the pool pane, click the **Tape** tab.
4. Select one or more tapes that are located in a cloud unit.
5. Click **Recall Cloud Tapes** to recall tapes from DD Cloud Tier.

Results

After the next scheduled data migration, the tapes are recalled from the cloud unit to the vault. From the vault, the tapes can be returned to a library.

CLI equivalent

Procedure

1. Identify the volume required to restore data.
2. Recall the tape volume from the vault.

Run the following command:

```
vtl tape recall start barcode T00001L3 count 1 pool cloud-
vtl-pool
```

3. Verify the recall operation started.

Run the following command:

```
data-movement status
```

4. Verify the recall operation completed successfully.

Run the following command:

```
vtl tape show all barcode T00001L3
```

```
Processing tapes....
Barcode      Pool          Location          State    Size    Used (%)
Comp  Modification Time
-----  -
T00001L3    cloud-vtl-pool  cloud-vtl slot 1  RW      5 GiB  5.0 GiB (99.07%)
239x      2017/05/05 10:41:41
-----  -

(S) Tape selected for migration to cloud. Selected tapes will move to cloud on the next
data-movement run.
(R) Recall operation is in progress for the tape.

VTL Tape Summary
-----
Total number of tapes:      1
Total pools:                1
Total size of tapes:       5 GiB
Total space used by tapes:  5.0 GiB
Average Compression:       239.1x
```

5. Validate the file location.

Run the following command:

```
filesys report generate file-location path /data/coll/
cloud-vtl-pool
```

```
filesys report generate file-location path /data/coll/cloud-vtl-pool
-----
File Name                               Location(Unit Name)
-----
/data/coll/cloud-vtl-pool/.vtl_pool      Active
/data/coll/cloud-vtl-pool/.vtc/T00001L3 Active
-----
```

6. Import the recalled tape to the DD VTL.

Run the following command:

```
vtl import cloud-vtl barcode T00001L3 count 1 pool cloud-
vtl-pool element slot
```

```
imported 1 tape(s)...sysadmin@ddbета70# vtl tape show cloud-vtlProcessing tapes.....
```

7. Check the volume into the backup application inventory.
8. Restore data through the backup application.
9. When restore is completed check the tape volume out of the backup application inventory.
10. Export the tape volume from the Data Domain VTL to the Data Domain Vault.
11. Move the tape back to the cloud unit.

Working with access groups

Access groups hold a collection of initiator WWPNs (worldwide port names) or aliases and the drives and changers they are allowed to access. A DD VTL default group named *TapeServer* lets you add devices that will support NDMP (Network Data Management Protocol)-based backup applications.

Access group configuration allows initiators (in general backup applications) to read and write data to devices in the same access group.

Access groups let clients 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.

Avoid making access group changes on a DD 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.

Selecting **Access Groups > Groups** displays the following information for all access groups.

Table 156 Access group information

Item	Description
Group Name	Name of group.
Initiators	Number of initiators in group.
Devices	Number of devices in group.

If you select **View All Access Groups**, you are taken to the Fibre Channel view.

From the More Tasks menu, you can create or delete a group.

Creating an access group

Access groups manage access between devices and initiators. Do not use the default *TapeServer* access group unless you are using NDMP.

Procedure

1. Select **Access Groups > Groups**.
2. Select **More Tasks > Group > Create**
3. In the Create Access Group dialog, enter a name, from 1 to 128 characters, and select **Next**.

4. Add devices, and select **Next**.
5. Review the summary, and select **Finish** or **Back**, as appropriate.

CLI Equivalent

```
# scsitarget group create My_Group service My_Service
```

Adding an access group device

Access group configuration allows initiators (in general backup applications) to read and write data to devices in the same access group.

Procedure

1. Select **Access Groups > Groups**. You can also select a specific *group*.
2. Select **More Tasks > Group > Create** or **Group > Configure**.
3. In the Create or Modify Access Group dialog, enter or modify the **Group Name** if desired. (This field is required.)
4. To configure initiators to the access group, check the box next to the initiator. You can add initiators to the group later.
5. Select **Next**.
6. In the Devices display, select Add (+) to display the Add Devices dialog.
 - a. Verify that 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 DD 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.

When presenting LUNs via attached FC ports on FC HBA/SLIC, ports can be designated as primary, secondary, or none. A Primary port for a set of LUNs is the port that is currently advertizing those LUNs to a fabric. A secondary port is a port that will broadcast a set of LUNs in the event of primary path failure (this requires manual intervention). A setting of none is used in the case where you do not wish to advertize selected LUNs. The presentation of LUNs depends on the SAN topology in question.

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.

Some initiators (clients) have specific rules for 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 DD VTL target port.

Check your initiator documentation for special rules, and if necessary, alter the device LUNs on the DD VTL target port to follow the rules. For example, if an initiator requires LUN 0 to be assigned on the DD 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.

The switchover to a secondary port is not an automatic operation. You must manually switch the DD VTL device to the secondary ports if the primary ports become unavailable.

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.

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 is displayed.

CLI Equivalent

```
# vtl group add VTL_Group vtl NewVTL changer lun 0 primary-port all secondary-port all
# vtl group add VTL_Group vtl NewVTL drive 1 lun 1 primary-port all secondary-port all
# vtl group add Setup_Test vtl Setup_Test drive 3 lun 3 primary-port endpoint-fc-0
secondary-port endpoint-fc-1
```

```
# vtl group show Setup_Test
Group: Setup_Test
```

```
Initiators:
Initiator Alias      Initiator WWPN
-----
tsm6_p23             21:00:00:24:ff:31:ce:f8
-----
```

```
Devices:
Device Name          LUN   Primary Ports   Secondary Ports   In-use Ports
-----
Setup_Test changer   0     all             all              all
Setup_Test drive 1   1     all             all              all
Setup_Test drive 2   2     5a              5b              5a
Setup_Test drive 3   3     endpoint-fc-0   endpoint-fc-1    endpoint-fc-0
-----
```

Modifying or deleting an access group device

You may need to modify or delete a device from an access group.

Procedure

1. Select **Protocols > VTL > Access Groups > Groups > group**.
2. Select **More Tasks > Group > Configure**.
3. In the Modify Access Group dialog, enter or modify the **Group Name**. (This field is required.)
4. To configure initiators to the access group, check the box next to the initiator. You can add initiators to the group later.
5. Select **Next**.
6. Select a device, and select the edit (pencil) icon to display the Modify Devices dialog. Then, follow steps a-e. If you simply want to delete the device, select the delete (X) icon, and skip to step e.

- a. Verify that the correct library is selected in the Library drop-down list, or select another library.
- b. In the Devices 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 LUN Start Address box.

This is the LUN that the DD 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.

Some initiators (clients) have specific rules for 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 DD VTL target port.

Check your initiator documentation for special rules, and if necessary, alter the device LUNs on the DD VTL target port to follow the rules. For example, if an initiator requires LUN 0 to be assigned on the DD 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 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.

The switchover to a secondary port is not an automatic operation. You must manually switch the DD VTL device to the secondary ports if the primary ports become unavailable.

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.

A drive appears with the same LUN on all ports that you have configured.

e. Select **OK**.

Deleting an access group

Before you can delete an access group, you must remove all of its initiators and LUNs.

Procedure

1. Remove all of the initiators and LUNs from the group.
2. Select **Access Groups > Groups**.
3. Select **More Tasks > Group > Delete**.
4. In the Delete Group dialog, select the checkbox of the group to be removed, and select **Next**.
5. In the groups confirmation dialog, verify the deletion, and select **Submit**.
6. Select **Close** when the Delete Groups Status displays *Completed*.

CLI Equivalent

```
# scsitarget group destroy My_Group
```

Working with a selected access group

Selecting **Access Groups > Groups > group** displays the following information for a selected access group.

Table 157 LUNs tab

Item	Description
LUN	Device address – maximum number is 16383. A LUN can be used only once within a group, but can be used again within another group. DD 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.

Table 157 LUNs tab (continued)

Item	Description
Primary Endpoints	Initial (or default) endpoint used by 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 primary endpoint fails.

Table 158 Initiators tab

Item	Description
Name	Name of initiator, which is either the WWPN or the alias assigned to the initiator.
WWPN	Unique worldwide 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.

From the More Tasks menu, with a group selected, you can configure that group, or set endpoints in use.

Selecting endpoints for a device

Since endpoints connect a device to an initiator, use this process to set up the endpoints before you connect the device.

Procedure

1. Select **Access Groups > Groups > *group***.
2. Select **More Tasks > Endpoints > Set In-Use**.
3. In the Set in-Use Endpoints dialog, select only specific devices, or select **Devices** to select all devices in the list.
4. Indicate whether the endpoints are primary or secondary.
5. Select **OK**.

Configuring the NDMP device TapeServer group

The DD VTL TapeServer group holds tape drives that interface with NDMP (Network Data Management Protocol)-based backup applications and that send control information and data streams over IP (Internet Protocol) instead of Fibre Channel (FC). A device used by the NDMP TapeServer must be in the DD VTL group TapeServer and is available *only* to the NDMP TapeServer.

Procedure

1. Add tape drives to a new or existing library (in this example, named “dd990-16”).
2. Create slots and CAPs for the library.
3. Add the created devices in a library (in this example, “dd990-16”) to the TapeServer access group.
4. Enable the NDMP daemon by entering at the command line:

```
# ndmpd enable
Starting NDMP daemon, please wait.....
NDMP daemon is enabled.
```

5. Ensure that the NDMP daemon sees the devices in the TapeServer group:

```
# ndmpd show devicenames
NDMP Device          Virtual Name          Vendor   Product          Serial Number
-----
/dev/dd_ch_c0t010    dd990-16 changer     STK      L180              6290820000
/dev/dd_st_c0t110    dd990-16 drive 1     IBM      ULTRIUM-TD3       6290820001
/dev/dd_st_c0t210    dd990-16 drive 2     IBM      ULTRIUM-TD3       6290820002
/dev/dd_st_c0t310    dd990-16 drive 3     IBM      ULTRIUM-TD3       6290820003
/dev/dd_st_c0t410    dd990-16 drive 4     IBM      ULTRIUM-TD3       6290820004
-----
```

6. Add an NDMP user (ndmp in this example) with the following command:

```
# ndmpd user add ndmp
Enter password:
Verify password:
```

7. Verify that user ndmp is added correctly:

```
# ndmpd user show
ndmp
```

8. Display the NDMP configuration:

```
# 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 changed from text to md5):

```
# ndmpd option set authentication md5
# 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 *Data Domain Operating System Command Reference Guide* for the complete command set and options.

Working with resources

Selecting **Resources** > **Resources** displays information about initiators and endpoints. An *initiator* is a backup client that connects to a system to read and write data using the Fibre Channel (FC) protocol. A specific initiator can support DD Boost over FC or DD VTL, but not both. An *endpoint* is the logical target on a DD system to which the initiator connects.

Table 159 Initiators tab

Item	Description
Name	Name of initiator, which is either the WWPN or the alias assigned to the initiator.
WWPN	Unique worldwide 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 (FC) port.
WWNN	Unique worldwide 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 FC node.
Online Endpoints	Group name where ports are seen by initiator. Displays <i>None</i> or <i>Offline</i> if the initiator is unavailable.

Table 160 Endpoints tab

Item	Description
Name	Specific name of endpoint.
WWPN	Unique worldwide 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 (FC) port.
WWNN	Unique worldwide 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 FC node.
System Address	System address for the endpoint.
Enabled	HBA (host bus adapter) port operational state, which is either <i>Yes</i> (enabled) or <i>No</i> (not enabled).
Status	DD VTL link status, which is either <i>Online</i> (capable of handling traffic) or <i>Offline</i> .

Configure Resources

Selecting **Configure Resources** takes you to the Fibre Channel area, where you can configure endpoints and initiators.

Working with initiators

Selecting **Resources > Resources > Initiators** displays information about initiators. An *initiator* is a client system FC HBA (fibre channel host bus adapter) WWPN (worldwide port name) with which the DD system interfaces. An *initiator name* is an alias for the client's WWPN, for ease of use.

While a client is mapped as an initiator – but before an access group has been added – the client cannot access any data on a DD system.

After adding an access group for the initiator or client, the client can access only the devices in that access group. A client can have access groups for multiple devices.

An access group may contain multiple initiators, but an initiator can exist in only one access group.

Note

A maximum of 1024 initiators can be configured for a DD system.

Table 161 Initiator information

Item	Description
Name	Name of initiator.
Group	Group associated with initiator.
Online Endpoints	Endpoints seen by initiator. Displays <code>none</code> or <code>offline</code> if initiator is unavailable.
WWPN	Unique worldwide 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 (FC) port.
WWNN	Unique worldwide 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 FC node.
Vendor Name	Name of vendor for initiator.

Selecting **Configure Initiators** takes you to the Fibre Channel area, where you can configure endpoints and initiators.

CLI Equivalent

```
# vtl initiator show
Initiator  Group      Status  WWNN                               WWPN                               Port
-----
tsm6_p1   tsm3500_a  Online  20:00:00:24:ff:31:ce:f8           21:00:00:24:ff:31:ce:f8           10b

Initiator  Symbolic Port Name                Address Method
-----
tsm6_p1   QLE2562 FW:v5.06.03 DVR:v8.03.07.15.05.09-k  auto
```

Working with endpoints

Selecting **Resources > Resources > Endpoints** provides information about endpoint hardware and connectivity.

Table 162 Hardware Tab

Item	Description
System Address	System address of endpoint.
WWPN	Unique worldwide 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 (FC) port.
WWNN	Unique worldwide 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 FC node.
Enabled	HBA (host bus adapter) port operational state, which is either <code>Yes</code> (enabled) or <code>No</code> (not enabled).

Table 162 Hardware Tab (continued)

Item	Description
NPIV	NPIV status of this endpoint: either Enabled or Disabled.
Link Status	Link status of this endpoint: either Online or Offline.
Operation Status	Operation status of this endpoint: either Normal or Marginal.
# of Endpoints	Number of endpoints associated with this endpoint.

Table 163 Endpoints Tab

Item	Description
Name	Specific name of endpoint.
WWPN	Unique worldwide 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 (FC) port.
WWNN	Unique worldwide 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 FC node.
System Address	System address of endpoint.
Enabled	HBA (host bus adapter) port operational state, which is either Yes (enabled) or No (not enabled).
Link Status	Link status of this endpoint: either Online or Offline.

Configure Endpoints

Selecting **Configure Endpoints** takes you to the Fibre Channel area, where you can change any of the above information for the endpoint.

CLI Equivalent

```
# scsitarget endpoint show list
Endpoint      System Address  Transport      Enabled  Status
-----
endpoint-fc-0 5a              FibreChannel   Yes      Online
endpoint-fc-1 5b              FibreChannel   Yes      Online
```

Working with a selected endpoint

Selecting **Resources > Resources > Endpoints > endpoint** provides information about the endpoint's hardware, connectivity, and statistics.

Table 164 Hardware tab

Item	Description
System Address	System address of endpoint.
WWPN	Unique worldwide 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.

Table 164 Hardware tab (continued)

Item	Description
WWNN	Unique worldwide 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 FC node.
Enabled	HBA (host bus adapter) port operational state, which is either <i>Yes</i> (enabled) or <i>No</i> (not enabled).
NPIV	NPIV status of this endpoint: either Enabled or Disabled.
Link Status	Link status of this endpoint: either Online or Offline.
Operation Status	Operation status of this endpoint: either Normal or Marginal.
# of Endpoints	Number of endpoints associated with this endpoint.

Table 165 Summary tab

Item	Description
Name	Specific name of endpoint.
WWPN	Unique worldwide 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.
WWNN	Unique worldwide 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 FC node.
System Address	System address of endpoint.
Enabled	HBA (host bus adapter) port operational state, which is either <i>Yes</i> (enabled) or <i>No</i> (not enabled).
Link Status	Link status of this endpoint: either Online or Offline.

Table 166 Statistics tab

Item	Description
Endpoint	Specific name of endpoint.
Library	Name of library containing the 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 167 Detailed Statistics tab

Item	Description
Endpoint	Specific name of endpoint.

Table 167 Detailed Statistics tab (continued)

Item	Description
# 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.

Working with pools

Selecting **Pools > Pools** displays detailed information for the Default pool and any other existing pools. A *pool* is a collection of tapes that maps to a directory on the file system. Pools are used to replicate tapes to a destination. You can convert directory-based pools to MTree-based pools to take advantage of the greater functionality of MTrees.

Note the following about pools:

- Pools can be of two types: MTree (recommended), or Directory, which is backward-compatible.
- A 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 pool to another.
- Pools are not accessible by backup software.
- No DD VTL configuration or license is needed on a replication destination when replicating pools.
- You must create tapes with unique barcodes. Duplicate barcodes may cause unpredictable behavior in backup applications and can be confusing to users.
- Two tapes in two different pools on a DD system may have the same name, and in this case, neither tape can be moved to the other tape's pool. Likewise, a pool sent to a replication destination must have a name that is unique on the destination.

Table 168 Pools tab

Item	Description
Location	The location of the pool.

Table 168 Pools tab (continued)

Item	Description
Type	Whether it is a Directory or MTree pool.
Tape Count	The number of tapes in the pool.
Capacity	The total configured data capacity of tapes in the pool, in GiB (Gibibytes base-2 equivalent of GB, Gigabytes).
Used	The amount of space used on virtual tapes in the pool.
Average Compression	The average amount of compression achieved for data on tapes in the pool.

Table 169 Replication tab

Item	Description
Name	The name of the pool.
Configured	Whether replication is configured for the pool: yes or no.
Remote Source	Contains an entry only if the pool is replicated from another DD system.
Remote Destination	Contains an entry only if the pool replicates to another DD system.

From the More Tasks menu, you can create and delete pools, as well as search for tapes.

Creating pools

You can create backward-compatible pools, if necessary for your setup, for example, for replication with a pre-5.2 DD OS system.

Procedure

1. Select **Pools > Pools**.
2. Select **More Tasks > Pool > Create**.
3. In the Create Pool dialog, enter a Pool Name, noting that a pool name:
 - cannot be “all,” “vault,” or “summary.”
 - cannot have a space or period at its beginning or end.
 - is case-sensitive.
4. If you want to create a directory pool (which is backward compatible with the previous version of DD System Manager), select the option “Create a directory backwards compatibility mode pool.” However, be aware that the advantages of using an MTree 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.
5. Select **OK** to display the Create Pool Status dialog.
 6. When the Create Pool Status dialog shows **Completed**, select **Close**. The pool is added to the Pools subtree, and you can now add virtual tapes to it.

CLI Equivalent

```
# vtl pool add VTL_Pool
A VTL pool named VTL_Pool is added.
```

Deleting pools

Before a pool can be deleted, you must have deleted 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.

Procedure

1. Select **Pools > Pools > *pool***.
2. Select **More Tasks > Pool > Delete**.
3. In the Delete Pools dialog, select the checkbox of items to delete:
 - The name of each pool, or
 - **Pool Names**, to delete all pools.
4. Select **Submit** in the confirmation dialogs.
5. When the Delete Pool Status dialog shows **Completed**, select **Close**.

The pool will have been removed from the Pools subtree.

Working with a selected pool

Both **Virtual Tape Libraries > VTL Service > Vault > *pool*** and **Pools > Pools > *pool*** display detailed information for a selected pool. Notice that pool “Default” always exists.

Pool tab

Table 170 Summary

Item	Description
Convert to MTree Pool	Select this button to convert a Directory pool to an MTree pool.
Type	Whether it is a Directory or MTree pool.
Tape Count	The number of tapes in the pool.
Capacity	The total configured data capacity of tapes in the pool, in GiB (Gibibytes, base-2 equivalent of GB, Gigabytes).

Table 170 Summary (continued)

Item	Description
Logical Used	The amount of space used on virtual tapes in the pool.
Compression	The average amount of compression achieved for data on tapes in the pool.

Table 171 Pool Tab: Cloud Data Movement - Protection Distribution

Item	Description
Pool type (%)	VTL Pool and Cloud (if applicable), with the current percentage of data in parentheses.
Name	Name of the local VTL pool, or cloud provider.
Logical Used	The amount of space used on virtual tapes in the pool.
Tape Count	The number of tapes in the pool.

Table 172 Pool Tab: Cloud Data Movement - Cloud Data Movement Policy

Item	Description
Policy	Age of tapes in days, or manual selection.
Older Than	Age threshold for an age-based data movement policy.
Cloud Unit	Destination cloud unit.

Tape tab**Table 173** Tape controls

Item	Description
Create	Create a new tape.
Delete	Delete the selected tapes.
Copy	Make a copy of a tape.
Move between Pool	Move the selected tapes to a different pool.
Select for Cloud Move	Schedule the selected tapes for migration to DD Cloud Tier.
Unselect from Cloud Move	Remove the selected tapes from the schedule for migration to DD Cloud Tier.
Recall Cloud Tapes	Recall the selected tapes from DD Cloud Tier.
Move to Cloud Now	Migrate the selected tapes to DD Cloud Tier without waiting for the next scheduled migration.

Table 174 Tape information

Item	Description
Barcode	Tape barcode.

Table 174 Tape information (continued)

Item	Description
Size	Maximum size of the tape.
Physical Used	Physical storage capacity used by the tape.
Compression	Compression ratio on the tape.
Location	Location of the tape.
Modification Time	Last time the tape was modified.
Recall Time	Last time the tape was recalled.

Replication tab**Table 175** Replication

Item	Description
Name	The name of the pool.
Configured	Whether replication is configured for this pool: yes or no.
Remote Source	Contains an entry only if the pool is replicated from another DD system.
Remote Destination	Contains an entry only if the pool replicates to another DD system.

You can also select the **Replication Detail** button, at the top right, to go directly to the Replication information panel for the selected pool.

From either the Virtual Tape Libraries or Pools area, from the More Tasks menu, you can create, delete, move, copy, or search for a tape in the pool.

From the Pools area, from the More Tasks menu, you can rename or delete a pool.

Converting a directory pool to an MTree pool

MTree pools have many advantages over directory pools. See the *Creating pools* section for more information.

Procedure

1. Make sure the following prerequisites have been met:
 - The source and destination pools must have been synchronized, so that the number of tapes, and the data on each side, remains intact.
 - The directory pool must not be a replication source or destination.
 - The file system must not be full.
 - The file system must not have reached the maximum number of MTrees allowed (100).
 - There must not already be an MTree with the same name.
 - If the directory pool is being replicated on multiple systems, those replicating systems must be known to the managing system.
 - If the directory pool is being replicated to an older DD OS (for example, from DD OS 5.5 to DD OS 5.4), it cannot be converted. As a workaround:

- Replicate the directory pool to a second DD system.
 - Replicate the directory pool from the second DD system to a third DD system.
 - Remove the second and third DD systems from the managing DD system's Data Domain network.
 - On any of the systems running DD OS 5.5, from the Pools submenu, select **Pools** and a directory pool. In the Pools tab, select **Convert to MTree Pool**.
2. With the directory pool you wish to convert highlighted, choose **Convert to MTree Pool**.
 3. Select **OK** in the Convert to MTree Pool dialog.
 4. Be aware that conversion affects replication in the following ways:
 - DD VTL is temporarily disabled on the replicated systems during 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.
 - DD Retention Lock cannot be enabled on systems involved in MTree pool conversion.

Moving tapes between pools

If they reside in the vault, 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.

Note

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.
-

Procedure

1. With a pool highlighted, select **More Tasks > Tapes > Move**.
Note that when started from a pool, the Tapes Panel allows tapes to be moved only between pools.
2. In the Move Tapes dialog, enter information to search for the tapes to move, and select **Search**:

Table 176 Move Tapes dialog

Field	User input
Location	Location cannot be changed.
Pool	Select the name of the pool where the tapes reside. If no pools have been created, use the Default pool.
Barcode	Specify a unique barcode. or leave the default (*) to import a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

- From the search results list, select the tapes to move.
- 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.
- Select **Next**.
- From the Move Tapes view, verify the summary information and tape list, and select **Submit**.
- Select **Close** in the status window.

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 started from the (named) Pool view.

Procedure

- With a pool highlighted, select **More Tasks > Tapes > Copy**.
- In the Copy Tapes Between Pools dialog. select the checkboxes of tapes to copy, or enter information to search for the tapes to copy, and select **Search**:

Table 177 Copy Tapes Between Pools dialog

Field	User input
Location	Select either a library or the Vault for locating the tape. 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.
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 Default pool.

Table 177 Copy Tapes Between Pools dialog (continued)

Field	User input
Barcode	Specify a unique barcode. or leave the default (*) to import a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be imported. If you leave this blank, the barcode default (*) is used.
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.
Items Selected	Shows 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. 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.

Renaming pools

A pool can be renamed only if none of its tapes is in a library.

Procedure

1. Select **Pools > Pools > pool**.
2. Select **More Tasks > Pool > Rename**.
3. In the Rename Pool dialog, enter the new Pool Name, with the caveat that this name:
 - cannot be “all,” “vault,” or “summary.”
 - cannot have a space or period at its beginning or end.
 - is case-sensitive.
4. Select **OK** to display the Rename Pool status dialog.
5. After the Rename Pool status dialog shows **Completed**, select **OK**.

The pool will have been renamed in the Pools subtree in both the Pools and the Virtual Tape Libraries areas.

CHAPTER 16

DD Replicator

This chapter includes:

- [DD Replicator overview](#)400
- [Prerequisites for replication configuration](#) 401
- [Replication version compatibility](#) 403
- [Replication types](#)407
- [Using DD Encryption with DD Replicator](#) 412
- [Replication topologies](#)413
- [Managing replication](#)418
- [Monitoring replication](#) 433
- [Replication with HA](#)434
- [Replicating a system with quotas to one without](#) 435
- [Replication Scaling Context](#) 435
- [Directory-to-MTree replication migration](#) 436
- [Using collection replication for disaster recovery with SMT](#)440

DD Replicator overview

Data Domain Replicator (DD Replicator) provides automated, policy-based, network-efficient, and encrypted replication for DR (disaster recovery) and multi-site backup and archive consolidation. DD Replicator asynchronously replicates only compressed, deduplicated data over a WAN (wide area network).

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. Cross-site deduplication further reduces bandwidth requirements when multiple sites are replicating to the same destination system. With cross-site deduplication, any redundant segment that is 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 DD system. For the highest level of security, DD Replicator can encrypt data being replicated between DD systems using the standard SSL (Secure Socket Layer) protocol.

DD Replicator scales performance and supported fan-in ratios to support large enterprise environments.

Before getting started with DD Replicator, note the following general requirements:

- DD Replicator is a licensed product. See your Data Domain sales representative to purchase licenses.
- You can usually replicate only between machines that are within two releases of each other, for example, from 6.0 to 6.2. However, there may be exceptions to this (as a result of atypical release numbering), so review the tables in the *Replication version compatibility* section, or check with your Data Domain representative.
- If you are unable to manage and monitor DD Replicator from the current version of the DD System Manager, use the `replication` commands that are described in the *Data Domain Operating System Command Reference Guide*.

Prerequisites for replication configuration

Before configuring a replication, review the following prerequisites to minimize initial data transfer time, prevent overwriting of data, etc.

- **Contexts** – Determine the maximum number of contexts for your DD systems by reviewing the replication streams numbers in the following table.

Table 178 Data streams sent to a Data Domain system

Model	RAM/ NVRAM	Backup write streams	Backup read streams	Repl ^a source streams	Repl ^a dest streams	Mixed
DD140, DD160, 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
DD620, DD630, DD640	8 GB / 0.5 GB or 1 GB	20	16	20	20	w<=20; r<=16; ReplSrc<=30; ReplDest<=20; ReplDest+w<=20; Total<=30
DD640, DD670	16 GB or 20 GB / 1 GB	90	30	60	90	w<=90; r<=30; ReplSrc<=60; 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<=90
DD860	72 GB ^b / 1 GB	90	50	90	90	w<=90; r<=50; ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; Total<=90
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 ^b / 4 GB	540	150	270	540	w<=540; r<=150; ReplSrc<=270; ReplDest<=540; ReplDest +w<=540; Total<=540
DD2200	8 GB	35	6	18	20	w<=35; r<=6; ReplSrc<=18; ReplDest<=20; ReplDest+w<=35; Total<=35
DD2200	16 GB	60	16	30	60	w<=60; r<=16; ReplSrc<=30; ReplDest<=60; ReplDest+w<=60; Total<=60
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 ^b / 4 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest +w<=270; Total<=270

Table 178 Data streams sent to a Data Domain system (continued)

Model	RAM/ NVRAM	Backup write streams	Backup read streams	Repl ^a source streams	Repl ^a dest streams	Mixed
DD4500	192 GB ^b / 4 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest +w<=270; Total<=270
DD7200	128 or 256 GB ^b / 4 GB	540	150	270	540	w<=540; r<=150; ReplSrc<=270; ReplDest<=540; ReplDest +w<=540; Total<=540
DD9500	256/512 GB	1885	300	540	1080	w<=1885; r<=300; ReplSrc<=540; ReplDest<=1080; ReplDest +w<=1080; Total<=1885
DD9800	256/768 GB	1885	300	540	1080	w<=1885; r<=300; ReplSrc<=540; ReplDest<=1080; ReplDest +w<=1080; Total<=1885
DD6300	48/96 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest +w<=270; Total<=270
DD6800	192 GB	400	110	220	400	w<=400; r<=110; ReplSrc<=220; ReplDest<=400; ReplDest +w<=400; Total<=400
DD9300	192/384 GB	800	220	440	800	w<=800; r<=220; ReplSrc<=440; ReplDest<=800; ReplDest +w<=800; Total<=800
DD VE 8 TB	8 GB / 512 MB	20	16	20	20	w<= 20 ; r<= 16 ReplSrc<=20; ReplDest<=20; ReplDest+w<=20; w+r+ReplSrc <=20;Total<=20
DD VE 16 TB	16 GB / 512 MB or 24 GB / 1 GB	45	30	45	45	w<= 45 ; r<= 30 ReplSrc<=45; ReplDest<=45; ReplDest+w<=45; w+r+ReplSrc <=45;Total<=45
DD VE 32 TB	24 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 48 TB	36 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 64 TB	48 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 96 TB	64 GB / 2 GB	180	50	90	180	w<= 180 ; r<= 50 ReplSrc<=90; ReplDest<=180; ReplDest +w<=180; w+r+ReplSrc <=180;Total<=180

Table 178 Data streams sent to a Data Domain system (continued)

Model	RAM/ NVRAM	Backup write streams	Backup read streams	Repl ^a source streams	Repl ^a dest streams	Mixed
DD3300 4 TB	12 GB (virtual memory) / 512 MB	20	16	30	20	w<= 20 ; r<= 16 ReplSrc<=30; ReplDest<=20; ReplDest+w<=20; w+r+ReplSrc <=30;Total<=30
DD3300 8 TB	32 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD3300 16 TB	32 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD3300 32 TB	46 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=140

a. DirRepl, OptDup, MTreeRepl streams

b. The Data Domain Extended Retention software option is available only for these devices with extended (maximum) memory

- **Compatibility** – If you are using DD systems running different versions of DD OS, review the next section on Replication Version Compatibility.
- **Initial Replication** – If the source holds a lot of data, the initial replication operation can take many hours. Consider putting both DD systems in the same location with a high-speed, low-latency link. After the first replication, you can move the systems to their intended locations because only new data will be sent.
- **Bandwidth Delay Settings** – Both the source and destination must have the same bandwidth delay settings. These tuning controls benefit replication performance over higher latency links by controlling the TCP (transmission control protocol) buffer size. The source system can then send enough data to the destination while waiting for an acknowledgment.
- **Only One Context for Directories/Subdirectories** – A directory (and its subdirectories) can be in only one context at a time, so be sure that a subdirectory under a source directory is not used in another directory replication context.
- **Adequate Storage** – At a minimum, the destination must have the *same amount of space* as the source.
- **Destination Empty for Directory Replication** – The destination directory must be empty for directory replication, or its contents no longer needed, because it will be overwritten.
- **Security** – DD OS requires that port 3009 be open in order to configure secure replication over an Ethernet connection.

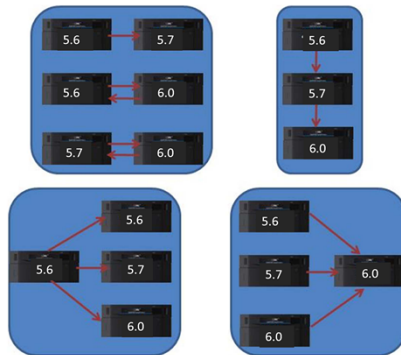
Replication version compatibility

To use DD systems running different versions of DD OS for a source or destination, the following tables provide compatibility information for single-node, DD Extended Retention, DD Retention Lock, MTree, directory, collection, delta (low bandwidth optimization), and cascaded replication.

In general:

- For DD Boost or OST, see “Optimized Duplication Version Compatibility” in the *Data Domain Boost for Partner Integration Administration Guide* or the *Data Domain Boost for OpenStorage Administration Guide* for supported configurations.
- MTree and directory replication cannot be used simultaneously for replicating the same data.
- The recovery procedure is valid for all supported replication configurations.
- File migration is supported whenever collection replication is supported.
- MTree replication between a source DD system running DD OS 5.2.x and a destination DD system running DD OS 5.4.x or DD OS 5.5.x is not supported when DD Retention Lock governance is enabled on the source MTree.
- For MTree replication from a source DD system running DD OS 6.0 to a target DD system running an earlier version of DD OS, the replication process behaves according to the older version of DD OS on the destination DD system. If a restore operation or cascade replication is performed from the destination DD system, no virtual synthetics are applied.
- For cascaded configurations, the maximum number of hops is two, that is, three DD systems.
Directory-to-MTree migration supports backward compatibility up to two previous releases. See [Directory-to-MTree replication migration](#) on page 436 for more information about directory-to-Mtree-migration.
- One-to-many, many-to-one, and cascaded replication support up to three consecutive DD OS release families, as seen in these figures.

Figure 9 Valid replication configurations



In these tables:

- Each DD OS release includes all releases in that family, for example, DD OS 5.7 includes 5.7.1, 5.7.x, 6.0, etc.
- c = collection replication
- dir = directory replication
- m = MTree replication
- del = delta (low bandwidth optimization) replication
- dest = destination
- src = source
- NA = not applicable

Table 179 Configuration: single-node to single-node

	5.0 (dest)	5.1 (dest)	5.2 (dest)	5.3 (dest)	5.4 (dest)	5.5 (dest)	5.6 (dest)	5.7 (dest)	6.0 (dest)	6.1 (dest)
5.0 (src)	c, dir, del	dir, del	dir, del	NA	NA	NA	NA	NA	NA	NA
5.1 (src)	dir, del	c, dir, del, m ^a	dir, del, m ^a	dir, del, m ^a	dir, del, m ^a	NA	NA	NA	NA	NA
5.2 (src)	dir, del	dir, del, m ^a	c, dir, del, m ^b	dir, del, m	dir, del, m	dir, del, m	NA	NA	NA	NA
5.3 (src)	NA	dir, del, m ^a	dir, del, m	c, dir, del, m	dir, del, m	dir, del, m	NA	NA	NA	NA
5.4 (src)	NA	dir, del, m ^a	dir, del, m	dir, del, m	c, dir, del, m	dir, del, m	dir, del, m	NA	NA	NA
5.5 (src)	NA	NA	dir, del, m	dir, del, m	dir, del, m	c, dir, del, m	dir, del, m	dir, del, m	NA	NA
5.6 (src)	NA	NA	NA	NA	dir, del, m	dir, del, m	c, dir, del, m	dir, del, m	dir, del, m	NA
5.7 (src)	NA	NA	NA	NA	NA	dir, del, m	dir, del, m	c, dir, del, m	dir, del, m	dir, del, m
6.0 (src)	NA	NA	NA	NA	NA	NA	dir, del, m	dir, del, m	c, dir, del, m	dir, del, m
6.1 (src)	NA	NA	NA	NA	NA	NA	NA	dir, del, m	dir, del, m	c, dir, del, m

- a. MTree replication is unsupported for DD VTL.
b. Collection replication is supported only for compliance data.

Table 180 Configuration: DD Extended Retention to DD Extended Retention

	5.0 (dest)	5.1 (dest)	5.2 (dest)	5.3 (dest)	5.4 (dest)	5.5 (dest)	5.6 (dest)	5.7 (dest)	6.0 (dest)	6.1 (dest)
5.0 (src)	c	NA	NA	NA	NA	NA	NA	NA	NA	NA
5.1 (src)	NA	c	m ^a	m ^b	m ^b	NA	NA	NA	NA	NA
5.2 (src)	NA	m ^a	c, m ^a	m ^a	m ^a	m ^a	NA	NA	NA	NA
5.3 (src)	NA	m ^c	m ^c	c, m	m	m	NA	NA	NA	
5.4 (src)	NA	m ^c	m ^c	m	c, m	m	m	NA	NA	NA
5.5 (src)	NA	NA	m ^c	m	m	c, m	m	m	NA	NA
5.6 (src)	NA	NA	NA	NA	m	m	c, m	m	m	
5.7 (src)	NA	NA	NA	NA	NA	m	m	c, m	m	m
6.0 (src)	NA	NA	NA	NA	NA	NA	m	m	c, m	m
6.1 (src)	NA	NA	NA	NA	NA	NA	NA	m	m	c, m

- a. File migration is not supported with MTree replication on either the source or destination in this configuration.

Table 180 Configuration: DD Extended Retention to DD Extended Retention (continued)

- b. File migration is not supported with MTree replication on the source in this configuration.
- c. File migration is not supported with MTree replication on the destination in this configuration.

Table 181 Configuration: single-node to DD Extended Retention

	5.0 (dest)	5.1 (dest)	5.2 (dest)	5.3 (dest)	5.4 (dest)	5.5 (dest)	5.6 (dest)	5.7 (dest)	6.0 (dest)	6.1 (dest)
5.0 (src)	dir	dir	NA	NA	NA	NA	NA	NA	NA	NA
5.1 (src)	dir	dir, m ^a	dir, m ^a	dir, m	dir, m	NA	NA	NA	NA	NA
5.2 (src)	dir	dir, m ^a	dir, m ^a	dir, m	dir, m	dir, m	NA	NA	NA	NA
5.3 (src)	NA	dir, m ^a	dir, m ^a	dir, m	dir, m	dir, m	NA	NA	NA	NA
5.4 (src)	NA	dir, m ^a	dir, m ^a	dir, m	dir, m	dir, m	dir, m	NA	NA	NA
5.5 (src)	NA	NA	dir, m ^a	dir, m	dir, m	dir, m	dir, m	dir, m	NA	NA
5.6 (src)	NA	NA	NA	NA	dir, m	dir, m	dir, m	dir, m	dir, m	NA
5.7 (src)	NA	NA	NA	NA	NA	dir, m	dir, m	dir, m	dir, m	dir, m
6.0 (src)	NA	NA	NA	NA	NA	NA	dir, m	dir, m	dir, m	dir, m
6.1 (src)	NA	NA	NA	NA	NA	NA	NA	dir, m	dir, m	dir, m

a. File migration is not supported for this configuration.

Replication types

Replication typically consists of a *source* DD system (which receives data from a backup system) and one or more *destination* DD systems. Each DD system can be the source and/or the destination for replication contexts. During replication, each DD system can perform normal backup and restore operations.

Each replication type establishes a *context* associated with an existing directory or MTree on the source. The replicated context is created on the destination when a context is established. The context establishes a replication pair, which is always active, and any data landing in the source will be copied to the destination at the earliest opportunity. Paths configured in replication contexts are absolute references and do not change based on the system in which they are configured.

A Data Domain system can be set up for directory, collection, or MTree replication.

- *Directory replication* provides replication at the level of individual directories.
- *Collection replication* duplicates the entire data store on the source and transfers that to the destination, and the replicated volume is read-only.
- *MTree replication* replicates entire MTrees (that is, a virtual file structure that enables advanced management). Media pools can also be replicated, and by default (as of DD OS 5.3), an MTree is created that will be replicated. (A media pool can also be created in backward-compatibility mode that, when replicated, will be a directory replication context.)

For any replication type, note the following requirements:

- A destination Data Domain system must have available storage capacity that is at least the size of the expected maximum size of the source directory. Be sure that the destination Data Domain system has enough network bandwidth and disk space to handle all traffic from replication sources.
- The file system must be enabled or, based on the replication type, will be enabled as part of the replication initialization.

- The source must exist.
- The destination must not exist.
- The destination will be created when a context is built and initialized.
- After replication is initialized, ownership and permissions of the destination are always identical to those of the source.
- In the replication command options, a specific replication pair is always identified by the destination.
- Both systems must have an active, visible route through the IP network so that each system can resolve its partner's host name.

The choice of replication type depends on your specific needs. The next sections provide descriptions and features of these three types, plus a brief introduction to Managed File Replication, which is used by DD Boost.

Managed file replication

Managed file replication, which is used by DD Boost, is a type of replication that is managed and controlled by backup software.

With managed file replication, backup images are directly transferred from one DD 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.

Here are some additional points to consider about managed file replication:

- Replication contexts do not need to be configured.
- Lifecycle policies control replication of information with no intervention from the user.
- DD Boost will build and tear down contexts as needed on the fly.

For more information, see the `ddboost file-replication` commands in the *Data Domain Operating System Command Reference Guide*.

Directory replication

Directory replication transfers deduplicated data within a DD file system directory configured as a replication source to a directory configured as a replication destination on a different system.

With directory replication, a DD system can simultaneously be the source of some replication contexts and the destination of other contexts. And that DD system can also receive data from backup and archive applications while it is replicating data.

Directory replication has the same flexible network deployment topologies and cross-site deduplication effects as managed file replication (the type used by DD Boost).

Here are some additional points to consider when using directory replication:

- Do not mix CIFS and NFS data within the same directory. A single destination DD system can receive backups from both CIFS clients and NFS clients as long as separate directories are used for CIFS and NFS.
- Any directory can be in only one context at a time. A parent directory may not be used in a replication context if a child directory of that parent is already being replicated.

- 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 DD 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.

Folder Creation Recommendations

Directory replication replicates data at the level of individual subdirectories under /data/col1/backup.

To provide a granular separation of data you must create, from a host system, other directories (DirA, DirB, etc.) within the /backup Mtree. Each directory should be based on your environment and the desire to replicate those directories to another location. You will not replicate the entire /backup MTree, but instead would set up replication contexts on each subdirectory underneath /data/col1/backup/ (ex. /data/col1/backup/DirC). The purpose of this threefold:

- It allows control of the destination locations as DirA may go to one site and DirB may go to another.
- This level of granularity allows management, monitoring, and fault isolation. Each replication context can be paused, stopped, destroyed, or reported on.
- Performance is limited on a single context. The creation of multiple contexts can improve aggregate replication performance.
- As a general recommendation, approximately 5 - 10 contexts may be required to distribute replication load across multiple replication streams. This must be validated against the site design and the volume and composition of the data at the location.

Note

Recommending a number of contexts is a design-dependent issue, and in some cases, significant implications are attached to the choices made about segregating data for the purposes of optimizing replication. Data is usually optimized for the manner in which it will rest – not in manner with which it will replicate. Keep this in mind when altering a backup environment.

MTree replication

MTree replication is used to replicate MTrees between DD 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.

With MTree replication, a DD system can be simultaneously the source of some replication contexts and the destination of other contexts. And that DD system can also receive data from backup and archive applications while it is replicating data.

MTree replication has the same flexible network deployment topologies and cross-site deduplication effects as managed file replication (the type used by DD Boost).

Here are some additional points to consider when using MTree replication:

- When replication is initialized, a destination read-only MTree is created automatically.
- Data can be logically segregated into multiple MTrees to promote greater replication performance.
- Snapshots must be created on source contexts.
- Snapshots cannot be created on a replication destination.
- Snapshots are replicated with a fixed retention of one year; however, the retention is adjustable on the destination and must be adjusted there.
- Replication contexts must be configured on both the source and the destination.
- Replicating DD VTL tape cartridges (or pools) simply means replicating MTrees or directories that contain DD VTL tape cartridges. Media pools are replicated by MTree replication, as a default. A media pool can be created in backward-compatibility mode and can then be replicated via directory-based replication. You cannot use the pool:// syntax to create replication contexts using the command line. When specifying pool-based replication in DD System Manager, either directory or MTree replication will be created, based on the media pool type.
- Replicating directories under an MTree is not permitted.
- A destination DD system must have available storage capacity of at least the post-compressed size of the expected maximum post-compressed size of the source MTree.
- 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.
- MTree replication is supported from DD Extended Retention systems to non-DD Extended Retention systems if both are running DD OS 5.5 or later.
- DD Retention Lock Compliance is supported with MTree replication, by default. If DD Retention Lock is licensed on a source, the destination must also have a DD Retention Lock license, or replication will fail. (To avoid this situation, you must disable DD Retention Lock.) If DD Retention Lock is enabled on a replication context, a replicated destination context will always contain data that is retention locked.

Automatic Multi-Streaming (AMS)

Automatic Multi-Streaming (AMS) improves MTree replication performance. It uses multiple streams to replicate a single large file (32 GB or larger) to improve network bandwidth utilization during replication. By increasing the replication speed for individual files, AMS also improves the pipeline efficiency of the replication queue, and provides improved replication throughput and reduced replication lag.

When the workload presents multiple optimization choices, AMS automatically selects the best option for the workload. For example, if the workload is a large file with fastcopy attributes, the replication operation uses fastcopy optimization to avoid the

overhead of scanning the file to identify unique segments between the replication pair. If the workload uses synthetics, replication uses synthetic replication on top of AMS to leverage local operations on the destination system for each replication stream to generate the file.

AMS is always enabled, and cannot be disabled.

Collection replication

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 DD file system.

Collection replication does not have the flexibility of the other types, but it can provide higher throughput and support more objects with less overhead, which may work better for high-scale enterprise cases.

Collection replication replicates the entire `/data/coll` area from a source DD system to a destination DD system.

Note

Collection replication is not supported for cloud-tier enabled systems.

Here are some additional points to consider when using collection replication:

- No granular replication control is possible. All data is copied from the source to the destination producing a read-only copy.
- Collection replication requires that the storage capacity of the destination system be equal to, or greater than, the capacity of the source system. If the destination capacity is less than the source capacity, the available capacity on the source is reduced to the capacity of the destination.
- The DD 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.
- With collection replication, all user accounts and passwords are replicated from the source to the destination. However, as of DD OS 5.5.1.0, other elements of configuration and user settings of the DD system are not replicated to the destination; you must explicitly reconfigure them after recovery.
- Collection replication is supported with DD Secure Multitenancy (SMT). Core SMT information, contained in the registry namespace, including the tenant and tenant-unit definitions with matching UUIDs is automatically transferred during replication operation. However, the following SMT information is not automatically included for replication, and must be configured manually on the destination system:
 - Alert notification lists for each tenant-unit
 - All users assigned to the DD Boost protocol for use by SMT tenants, if DD Boost is configured on the system
 - The default-tenant-unit associated with each DD Boost user, if any, if DD Boost is configured on the system

[Using collection replication for disaster recovery with SMT](#) on page 440 describes how to manually configure these items on the replication destination.

- DD Retention Lock Compliance supports collection replication.
- Collection replication is not supported in cloud tier-enabled systems.
- With collection replication, data in a replication context on the source system that has not been replicated cannot be processed for file system cleaning. If file system

cleaning cannot complete because the source and destination systems are out of sync, the system reports the cleaning operation status as `partial`, and only limited system statistics are available for the cleaning operation. If collection replication is disabled, the amount of data that cannot be processed for file system cleaning increases because the replication source and destination systems remain out of sync. The KB article *Data Domain: An overview of Data Domain File System (DDFS) clean/garbage collection (GC) phases*, available from the Online Support site at <https://support.emc.com>, provides additional information.

- To enhance throughput in a high bandwidth environment, run the `replication modify <destination> crepl-gc-gw-optim` command to disable collection replication bandwidth optimization.

Using DD Encryption with DD Replicator

DD Replicator can be used with the optional *DD Encryption* feature, enabling encrypted data to be replicated using collection, directory, or MTree replication

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 type works uniquely with encryption and offers the same level of security.

- *Collection replication* requires the source and destination to have the same encryption configuration, because the destination data is expected to be an exact replica of the source data. In particular, the encryption feature must be turned on or off at both the source and destination, 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 transmits the data in encrypted form, and also transmits the encryption keys to the destination. The data can be recovered at the destination because the destination has the same passphrase and the same system encryption key.

Note

Collection replication is not supported for cloud-tier enabled systems.

- *MTree or directory replication* does not require encryption configuration to be the same at both the source and destination. Instead, the source and destination securely exchange the destination's encryption key during the replication association phase, and the data is re-encrypted at the source using the destination's encryption key before transmission to the destination. If the destination has a different encryption configuration, the data transmitted is prepared appropriately. For example, if the feature is turned off at the destination, the source decrypts the data, and it is sent to the destination 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 destination, it uses the same encryption keys and encrypted data as its source. If the last system is an MTree or directory replication destination, it uses its own key, and the data is encrypted at its source. The encryption key for the destination at each link is used for encryption. Encryption for systems in the chain works as in a replication pair.

Replication topologies

DD Replicator supports five replication topologies (one-to-one, one-to-one bidirectional, one-to-many, many-to-one, and cascaded). The tables in this section show (1) how these topologies work with three types of replication (MTree, directory, and collection) and two types of DD systems [single node (SN) and DD Extended Retention] and (2) how mixed topologies are supported with cascaded replication.

In general:

- Single node (SN) systems support all replication topologies.
- Single node-to-single node (SN -> SN) can be used for all replication types.
- DD Extended Retention systems cannot be the source for directory replication.
- Collection replication cannot be configured from either a single node (SN) system to a DD Extended Retention-enabled system, nor from a DD Extended Retention-enabled system to an SN system.
- Collection replication cannot be configured from either an SN system to a DD high availability-enabled system, nor from a DD high availability-enabled system to an SN system.
- For MTree and Directory replication, DD high availability systems are treated like SN systems.
- Collection replication cannot be configured if any or both systems have Cloud Tier enabled.

In this table:

- SN = single node DD system (no DD Extended Retention)
- ER = DD Extended Retention system

Table 182 Topology Support by Replication Type and DD System Type

Topologies	MTree Replication	Directory Replication	Collection Replication
one-to-one	{SN ER} -> {SN ER} ER->SN [supported starting with 5.5 release; prior to 5.5, it is recovery only]	SN -> SN SN -> ER	SN -> SN ER -> ER
one-to-one bidirectional	{SN ER} -> {SN ER}	SN -> SN	not supported
one-to-many	{SN ER} -> {SN ER}	SN -> SN SN -> ER	not supported
many-to-one	{SN ER} -> {SN ER}	SN -> SN SN -> ER	not supported
cascaded	{SN ER} -> {SN ER} -> {SN ER}	SN -> SN -> SN SN -> SN -> ER	ER -> ER -> ER SN -> SN -> SN

Cascaded replication supports mixed topologies where the second leg in a cascaded connection is different from the first type in a connection (for example, A -> B is directory replication, and B -> C is collection replication).

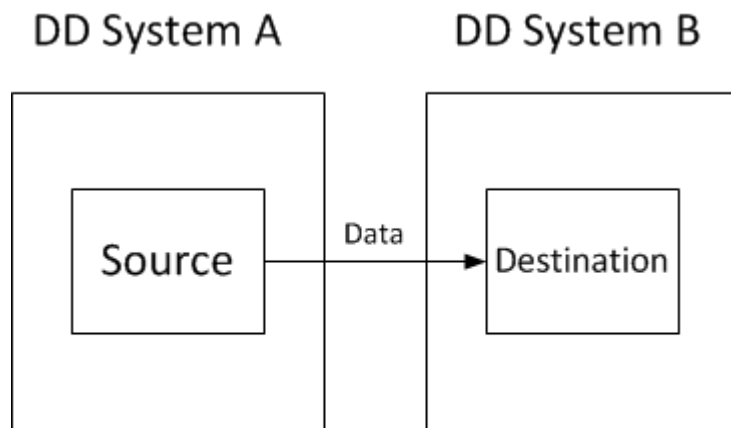
Table 183 Mixed Topologies Supported with Cascaded Replication

Mixed Topologies	
SN – Dir Repl -> ER – MTree Repl -> ER – MTree Repl	SN – Dir Repl -> ER – Col Repl -> ER – Col Repl
SN – MTree Repl -> SN – Col Repl -> SN – Col Repl	SN – MTree Repl -> ER – Col Repl -> ER – Col Repl

One-to-one replication

The simplest type of replication is from a DD source system to a DD destination system, otherwise known as a *one-to-one* replication pair. This replication topology can be configured with directory, MTree, or collection replication types.

Figure 10 One-to-one replication pair

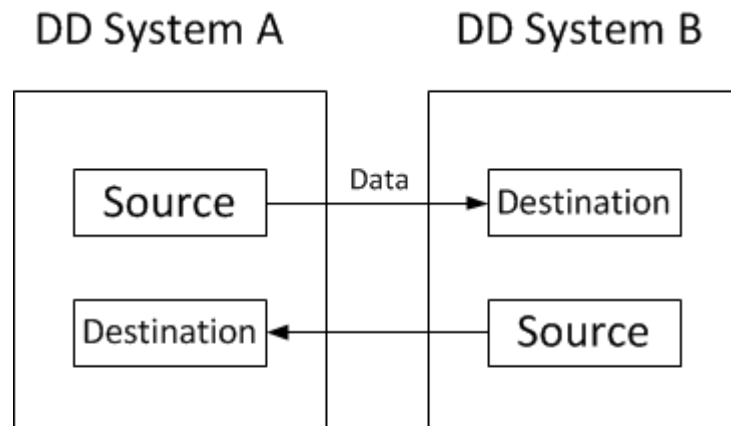


Data flows from the source to the destination system.

Bi-directional replication

In a bi-directional replication pair, data from a directory or MTree on DD system A is replicated to DD system B, and from another directory or MTree on DD system B to DD system A.

Figure 11 Bi-directional replication



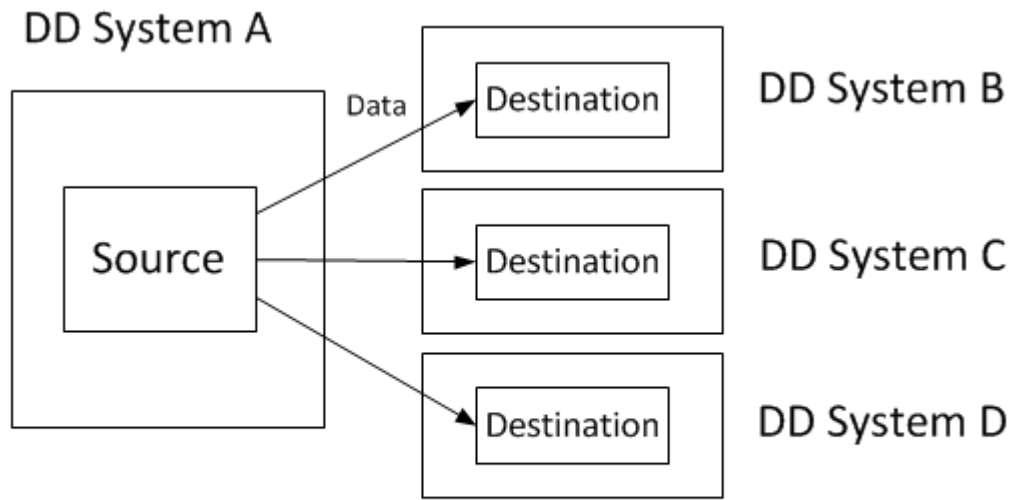
Data flows in both directions
between two systems.

One-to-many replication

In one-to-many replication, data flows from a source directory or MTree on one DD system to several destination DD 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.

Figure 12 One-to-many replication

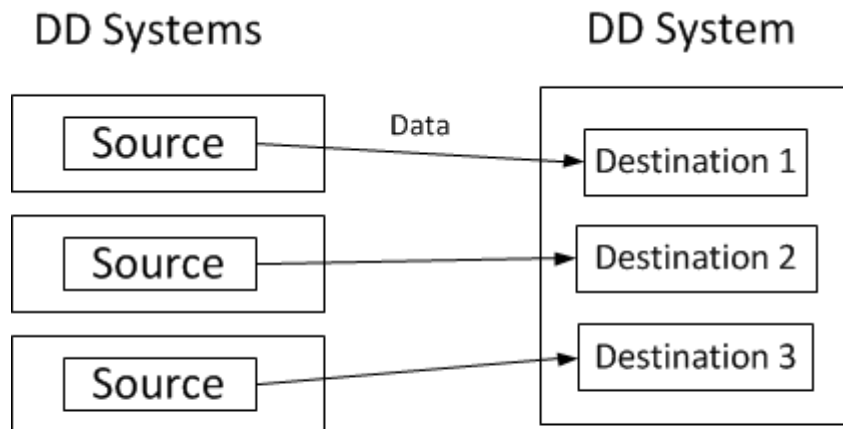


Data flows from a directory or MTree source system to many destination systems.

Many-to-one replication

In many-to-one replication, whether with MTree or directory, replication data flows from several source DD systems to a single destination DD system. This type of replication can be used to provide data recovery protection for several branch offices on a corporate headquarter's IT system.

Figure 13 Many-to-one replication



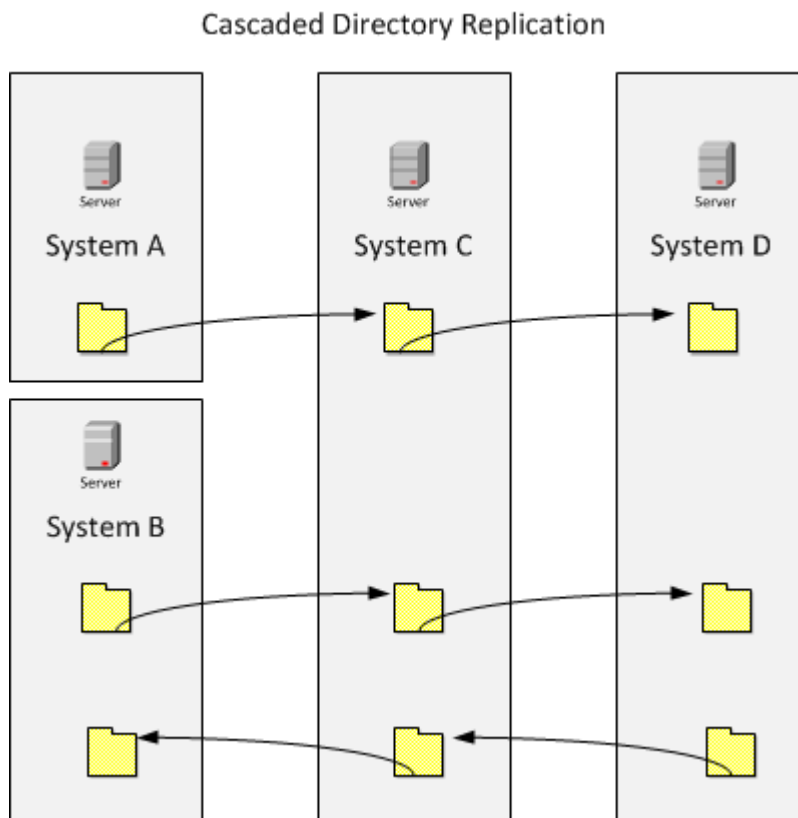
Data flows from many source systems to one destination system.

Cascaded replication

In a cascaded replication topology, a source directory or MTree is chained among three DD 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 14 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.

Managing replication

You can manage replication using the Data Domain System Manager (DD System Manager) or the Data Domain Operating System (DD OS) Command Line Interface (CLI).

To use a graphical user interface (GUI) to manage replication, log in to the DD System Manager.

Procedure

1. From the menu at the left of the DD System Manager, select **Replication**. If your license has not been added yet, select **Add License**.
2. Select **Automatic** or **On-Demand** (you must have a DD Boost license for on-demand).

CLI Equivalent

You can also log in at the CLI:

```
login as: sysadmin
Data Domain OS 6.0.x.x-12345
Using keyboard-interactive authentication.
Password:
```

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 a DD system, displaying aggregated information about the selected DD system – that is, summary information about the inbound and outbound replication pairs. The focus is the DD system, itself, and the inputs to it and outputs from it.

The Summary table can be filtered by entering a Source or Destination name, or by selecting a State (Error, Warning, or Normal).

Table 184 Replication Summary view

Item	Description
Source	System and path name of the 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.chaos.local/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.chaos.local/data/coll/MTree1</code> .
Type	Type of context: MTree, directory (Dir), or Pool.
State	Possible states of replication pair status include:

Table 184 Replication Summary view (continued)

Item	Description
	<ul style="list-style-type: none"> • Normal – If the replica is Initializing, Replicating, Recovering, Resyncing, or Migrating. • 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). • Warning – If there is an unusual delay for the first five states, or for the Uninitialized state. • Error – Any possible error states, such as Disconnected.
Synced As Of Time	Timestamp for last automatic replication sync operation performed by the 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 for a replication context

Selecting one replication context from the Summary view populates that context's information in Detailed Information, Performance Graph, Completion Stats, and Completion Predictor.

Table 185 Detailed Information

Item	Description
State Description	Message about state of replica.
Source	System and path name of source context, with format <code>system.path</code> . For example, for directory <code>dir1</code> on system <code>dd120-22</code> , you would see <code>dd120-22.chaos.local/data/coll/dir1</code> .
Destination	System and path name of destination context, with format <code>system.path</code> . For example, for MTree <code>MTree1</code> on system <code>dd120-44</code> , you would see <code>dd120-44.chaos.local/data/coll/MTree1</code> .
Connection Port	System name and listen port used for replication connection.

Table 186 Performance Graph

Item	Description
Pre-Comp Remaining	Pre-compressed data remaining to be replicated.

Table 186 Performance Graph (continued)

Item	Description
Pre-Comp Written	Pre-compressed data written on the source.
Post-Comp Replicated	Post-compressed data that has been replicated.

Table 187 Completion Stats

Item	Description
Synced As Of Time	Timestamp for last automatic replication sync operation performed by the 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.
Status	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"> • Replication • File System • Replication Lock • Encryption at Rest • Encryption over Wire • Available Space • Low Bandwidth Optimization • Compression Ratio • Low Bandwidth Optimization Ratio

Completion Predictor

The Completion Predictor is a widget for tracking a backup job's progress and for predicting when replication will complete, for a selected context.

Creating a replication pair

Before creating a replication pair, make sure the destination does not *exist*, or you will get an error.

Procedure

1. Select **Replication > Automatic > Summary tab > Create Pair** .

2. In the Create Pair dialog, add information to create an inbound or outbound MTree, directory, collection, or pool replication pair, as described in the next sections.

Adding a DD system for replication

You may need to add a DD system as either a host or a destination before you can create a replication pair.

Note

Make sure the system being added is running a compatible DD OS version.

Procedure

1. In the Create Pair dialog, select Add System.
2. For System, enter the hostname or IP address of the system to be added.
3. For User Name and Password, enter the sysadmin's user name and password.
4. Optionally, select **More Options** 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.

Note

IPv6 addresses are supported only when adding a DD OS 5.5 or later system to a management system using DD OS 5.5 or later.

5. Select **OK**.

Note

If the system is unreachable after adding it to DD System Manager, make sure that there is a route from the managing system to the system being added. If a hostname (either a fully qualified domain name (FQDN) or non-FQDN) is entered, make sure it is resolvable on the managed system. Configure a domain name for the managed system, ensure a DNS entry for the system exists, or ensure an IP address to hostname mapping is defined.

6. If the system certificate is not verified, the Verify Certificate dialog shows details about the certificate. Check the system credentials. Select **OK** if you trust the certificate, or select **Cancel**.

Creating a collection replication pair

See the *Collection replication* section for general information about this type of replication.

Before creating a collection replication pair, make sure:

- The storage capacity of the destination system is 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 destination has been destroyed, and subsequently re-created, but not enabled.
- Each destination and each source is in only one context at a time.
- The file system is disabled on the replica, while configuring and enabling encryption on the source.

- The file system is disabled on the source, while configuring and enabling encryption on the replica.

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. If you want to change any host connection settings, select the **Advanced** tab.
5. 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 an MTree, directory, or pool replication pair

See the *MTree replication* and *Directory replication* sections for general information about these types of replication.

When creating an MTree, directory, or pool replication pair:

- Make sure the replication is transiting/exiting the correct interface. When defining a replication context, the host names of the source and destination must resolve with forward and reverse lookups. To make the data transit alternate interfaces on the system, other than the default resolving interface, the replication context must be modified after creation. It may be necessary to set up host files to ensure that contexts are defined on non-resolving (cross-over) interfaces.
- You can “reverse” the context for an MTree replication, that is, you can switch the destination and the source.
- Subdirectories within an MTree cannot be replicated, because the MTree, in its entirety, is replicated.
- MTree replication is supported from DD Extended Retention-enabled systems to non-DD Extended Retention-enabled systems, if both are running DD OS 5.5 or later.
- The destination DD system must have available storage capacity of at least the post-compressed size of the expected maximum post-compressed size of the source directory or MTree.
- When replication is initialized, a destination directory is created automatically.
- A DD system can simultaneously be the source for one context and the destination for another context.

Procedure

1. In the Create Pair dialog, 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. If you want to change any host connection settings, select the **Advanced** tab.
7. Select **OK**.

The Replication from the source to the destination begins.

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.

CLI Equivalent

Here are examples of creating MTree or directory replication pairs at the CLI. The last example specifies the IP version used as a replication transport.

```
# replication add source mtree://ddsource.test.com/data/col1/
examplemtree destination mtree://ddtarget.test.com/data/col1/
examplemtree (Mtree example)
# replication add source dir://ddsource.test.com/data/col1/
directorytorep destination dir://ddtarget.test.com/backup/
directorytorep
# replication add source dir://ddsource.test.com/data/col1/
directorytorep destination dir://ddtarget.test.com/backup/
directorytorep ipversion ipv6
```

To start replication between a source and destination, use the `replication initialize` command on the source. This command checks that the configuration and connections are correct and returns error messages if any problems appear.

```
# replication initialize mtree://host3.test.com/data/col1/
mtree1/
```

Configuring bi-directional replication

To create a bi-directional replication pair, use the directory or MTree replication pair procedure (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.

Configuring one-to-many replication

To create a one-to-many replication pair, use the directory or MTree replication pair procedure (for example, using `mtree1`) on host A to: (1) `mtree1` on host B, (2) `mtree1` on host C, and (3) `mtree1` on host D. 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 a many-to-one replication pair, use the directory or MTree replication pair procedure [for example, (1) mtree1 from host A to mtree1 on host C and (2) mtree2 on host B to mtree2 on host C.]

Configuring cascaded replication

To create a cascaded replication pair, use the directory or MTree replication pair procedure: (1) mtree1 on host A to mtree1 on host B, and (2) 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.

Disabling and enabling 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 as an active connection to the source.

Procedure

1. Select one or more replication pairs in the Summary table, and select **Disable Pair**.
2. In the Display Pair dialog, 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**. Replication of data is resumed.

CLI Equivalent

```
# replication disable {destination | all}
# replication enable {destination | all}
```

Deleting a replication pair

When a directory or MTree replication pair is deleted, the destination directory or MTree, respectively, becomes writeable. When a collection replication pair is deleted, the destination DD system becomes a stand-alone read/write system, and the file system is disabled.

Procedure

1. Select one or more replication pairs in the Summary table, and select **Delete Pair**.
2. In the Delete Pair dialog, select **Next** and then **OK**. The replication pairs are deleted.

CLI Equivalent

Before running this command, always run the `filesys disable` command. Then, afterward, run the `filesys enable` command

```
# replication break {destination | all}
```

Changing host connection settings

To direct traffic out of a specific port, modify a current context by altering the connection host parameter using a host name previously defined in the local hosts file to address the alternate system. That host name will correspond to the destination.

The host entry will indicate an alternate destination address for that host. This may be required on both the source and destination systems.

Procedure

1. Select the replication pair in the Summary table, and select **Modify Settings**. You can also change these settings when you are performing Create Pair, Start Resync, or Start Recover by selecting the **Advanced** tab.
2. In the Modify Connection Settings dialog, modify any or all of these settings:

a. **Use Low Bandwidth Optimization** – For enterprises with small data sets and 6 Mb/s or less bandwidth networks, DD Replicator can further reduce the amount of data to be sent using *low bandwidth optimization*. This enables remote sites with limited bandwidth to use less bandwidth or to replicate and protect more of their data over existing networks. Low bandwidth optimization must be enabled on both the source and destination DD systems. If the source and destination have incompatible low bandwidth optimization settings, low bandwidth optimization will be inactive for that context. After enabling low bandwidth optimization on the source and destination, both systems must undergo a full cleaning cycle to prepare the existing data, so run `filesys clean start` on both systems. The duration of the cleaning cycle depends on the amount of data on the DD system, but takes longer than a normal cleaning. For more information on the `filesys` commands, see the *Data Domain Operating System Command Reference Guide*.

Important: Low bandwidth optimization is not supported if the DD Extended Retention software option is enabled on either DD system. It is also not supported for Collection Replication.

- b. **Enable Encryption Over Wire** – DD Replicator supports encryption of data-in-flight by using standard SSL (Secure Socket Layer) protocol version 1.0.1, which uses the ADH-AES256-GCM-SHA384 and DHE-RSA-AES256-GCM-SHA384 cipher suites to establish secure replication connections. Both sides of the connection must enable this feature for encryption to proceed.
- c. **Network Preference** – You may choose IPv4 or IPv6. An IPv6-enabled replication service can still accept connections from an IPv4 replication client if the service is reachable via IPv4. An IPv6-enabled replication client can still communicate with an IPv4 replication service if the service is reachable via IPv4.
- d. **Use Non-default Connection Host** – The source system transmits data to a destination system listen port. Since 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.

3. Select **Next** and then **Close**.

The replication pair settings are updated, and replication resumes.

CLI Equivalent

```
#replication modify <destination> connection-host <new-host-name> [port <port>]
```

Managing replication systems

You can add or delete Data Domain systems to be used for replication using the Manage Systems dialog.

Procedure

1. Select **Modify Settings**.
2. In the Manage Systems dialog, add and/or delete Data Domain systems, as required.
3. Select **Close**.

Recovering data from a replication pair

If source replication data becomes inaccessible, it can be *recovered* from the replication pair destination. The source must be empty before recovery can proceed. Recovery can be performed for all replication topologies, except for MTree replication.

Recovery of data from a directory pool, as well as from directory and collection replication pairs, is described in the next sections.

Recovering directory pool data

You can recover data from a directory-based pool, but not from an MTree-based pool.

Procedure

1. Select **More > Start Recover**.
2. In the Start Recover dialog, 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. If you want to change any host connection settings, select the **Advanced** tab.
7. Select **OK** to start the recovery.

Recovering collection replication pair data

To successfully recover collection replication pair data, the source file system must be in a pristine state, and the destination context must be fully initialized.

Procedure

1. Select **More > Start Recover** to display the Start Recover dialog.
2. Select **Collection** from the **Replication Type** menu.
3. Select the source system host name from the **System to recover to** menu.
4. Select the destination system host name 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. To change any host connection settings, select the **Advanced** tab.
7. Select **OK** to start the recovery.

Recovering directory replication pair data

To successfully recover directory replication pair data, the same directory used in the original context must be created (but left empty).

Procedure

1. Select **More > Start Recover** to display the Start Recover dialog.
2. Select **Directory** from the **Replication Type** menu.
3. Select the host name of the *system to which data needs to be restored* from the **System to recover to** menu.
4. Select the host name 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. To change any host connection settings, select the **Advanced** tab.
7. Select **OK** to start the recovery.

Aborting a replication pair recovery

If a replication pair recovery fails or must be terminated, you can 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

As soon as possible, you should restart recovery on the source.

Resyncing an MTree, directory, or pool replication pair

Resynchronization is the process of recovering (or bringing back into sync) the data between a source and a destination replication pair after a manual break. The replication pair are resynchronized so both endpoints contain the same data. Resynchronization is available for MTree, directory, and pool replication, but not for collection replication.

A replication resynchronization can also be used:

- To recreate a context that has been deleted.
- When a destination runs out of space, but the source still has data to replicate.
- To convert a directory replication pair to an MTree replication pair.

Procedure

1. Delete the context on both the replication source and replication destination systems.
2. From either the replication source or replication destination system, select **More > Start Resync** to display the Start Resync dialog.
3. Select the Replication Type to be resynced: **Directory**, **MTree**, or **Pool**.
4. Select the replication source system host name from the **Source System** menu.

5. Select the replication destination system host name from the **Destination System** menu.
6. Enter the replication source path in the **Source Path** text box.
7. Enter the replication destination path in the **Destination Path** text box.
8. To change any host connection settings, select the **Advanced** tab.
9. Select **OK**.

CLI Equivalent

```
# replication resync destination
```

Aborting a replication pair resynchronization

If a replication pair resynchronization fails or must be terminated, you can stop the resynchronization.

Procedure

1. From either the replication source or replication destination system, select **More > 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 resynchronization.
3. Select **OK**.

DD Boost view

The DD Boost view provides configuration and troubleshooting information to NetBackup administrators who have configured DD systems to use DD Boost AIR (Automatic Image Replication) or any DD Boost application that uses managed file replication.

See the *Data Domain Boost for OpenStorage Administration Guide* for DD Boost AIR configuration instructions.

The **File Replication** tab displays:

- 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** for more information about these remote system files.

The **Storage Unit Associations** tab 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.

An attempt is made to match both ends of a DD Boost path to form a pair and present this as one pair/record. If the match is impossible, for various reasons, the remote path will be listed as *Unresolved*.

Remote system files

The Show Details button provides information for the selected remote file replication system. File Replications shows starting and ending information, as well as size and data amount, for the selected remote file replication system. The Performance Graph shows performance over time for the selected remote file replication system.

Table 188 File Replications

Item	Description
Start	Starting point of time period.
End	Ending point of time period.
File Name	Name of specific replication file.
Status	Most recent status (Success, Failure).
Pre-Comp Size (MiB)	Amount of pre-compressed outbound and inbound data, as compared to network throughput or post-compressed data (in MiB).
Network Bytes (MiB)	Amount of network throughput data (in MiB).

Table 189 Performance Graph

Item	Description
Duration	Duration for replication (either 1d, 7d or 30d).
Interval	Interval for replication (either Daily or Weekly).
Pre-Comp Replicated	Amount of pre-compressed outbound and inbound data (in GiB).
Post-Comp Replicated	Amount of post-compressed data (in GiB).
Network Bytes	Amount of network throughput data (in GiB).
Files Succeeded	Number of files that were successfully replicated.
Files Failed	Number of files that failed to be replicated.
Show in new window	Brings up a separate window.
Print	Prints the graph.

Topology view

The Topology view shows how the selected replication pairs are configured in the network.

- The arrow – which is green (normal), yellow (warning), or red (error) – between DD systems represents one or more replication pairs.
- To view details, select a context , which opens the Context Summary dialog, with links to **Show Summary**, **Modify Options**, **Enable/Disable Pair**, **Graph Performance**, and **Delete Pair**.
- 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.

Performance view

The Performance view displays a graph that represents the fluctuation of data during replication. These are aggregated statistics of each replication pair for this DD system.

- **Duration** (x-axis) is 30 days by default.
- **Replication Performance** (y-axis) is in GibiBytes or MebiBytes (the binary equivalents of GigaBytes and MegaBytes).
- **Network In** is the total replication network bytes entering the system (all contexts).
- **Network Out** is the total replication network bytes leaving the system (all contexts).
- For a reading of a specific point in time, hover the cursor over a place on the graph.
- 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.

Advanced Settings view

Advanced Settings lets you manage throttle and network settings.

Throttle Settings

- **Throttle Override** – Displays throttle rate if configured, or 0 meaning all replication traffic is stopped.
- **Permanent Schedule** – Displays the time and days of the week on which scheduled throttling occurs.

Network Settings

- **Bandwidth** – Displays the configured data stream rate if bandwidth has been configured, or Unlimited (default) if not. The average data stream to the replication destination is at least 98,304 bits per second (12 KiB).

- **Delay** – Displays the configured network delay setting (in milliseconds) if it has been configured, or None (default) if not.
- **Listen Port** – Displays the configured listen port value if it has been configured, or 2051 (default) if not.

Adding throttle settings

To modify the amount of bandwidth used by a network for replication, you can set a *replication throttle* for replication traffic.

There are three types of replication throttle settings:

- **Scheduled throttle** – The throttle rate is set at a predetermined time or period.
- **Current throttle** – The throttle rate is set until the next scheduled change, or until a system reboot.
- **Override throttle** – The previous two types of throttle are overridden. This persists – even through reboot – until you select **Clear Throttle Override** or issue the `replication throttle reset override` command.

You can also set a default throttle or a throttle for specific destinations, as follows:

- **Default throttle** – When configured, all replication contexts are limited to this throttle, except for those destinations specified by destination throttles (see next item).
- **Destination throttle** – This throttle is used when only a few destinations need to be throttled, or when a destination requires a throttle setting different from the default throttle. When a default throttle already exists, this throttle takes precedence for the destination specified. For example, you can set the default replication throttle to *10 kbps*, but – using a destination throttle – you can set a single collection replication context to *unlimited*.

Note

Currently, you can set and modify destination throttle only by using the command-line interface (CLI); this functionality is not available in the DD System Manager. For documentation on this feature, see the `replication throttle` command in the *Data Domain Operating System Command Reference Guide*. If the DD System Manager detects that you have one or more destination throttles set, you will be given a warning, and you should use the CLI to continue.

Additional notes about replication throttling:

- Throttles are set only at the source. The only throttle that applies to a destination is the **0 Bps (Disabled)** option, which disables all replication traffic.
- The minimum value for a replication throttle is 98,304 bits per second.

Procedure

1. Select **Replication > Advanced Settings > 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 **Every Day** or by selecting checkbox(es) next to individual day(s).
3. Set the time that throttling is to start with the **Start Time** drop-down selectors for the hour:minute and AM/PM.
4. For **Throttle Rate**:
 - Select **Unlimited** to set no limits.

- Enter a number in the text box (for example, 20000), and select the rate from the menu (bps, Kbps, Bps, or KBps).
 - Select the **0 Bps (disabled)** option to disable all replication traffic.
5. Select **OK** to set the schedule. The new schedule is shown under **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

You can delete a single throttle setting or all throttle settings at once.

Procedure

1. Select **Replication > Advanced Settings > Delete Throttle Setting** to display the Delete Throttle Setting dialog.
2. Select the checkbox for the throttle setting to delete, or select the heading checkbox to delete all settings. This list can include settings for the “disabled” state.
3. Select **OK** to remove the setting.
4. In 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.

Procedure

1. Select **Replication > Advanced Settings > Set Throttle Override** to display the Throttle Override dialog.
2. Either set a new throttle override, or clear a previous override.
 - a. To set a new throttle override:
 - Select **Unlimited** to revert to the system-set throttle rate (no throttling performed), or
 - Set the throttling bit and rate in the text box (for example, 20000) and (bps, Kbps, Bps, or KBps), or
 - Select **0 Bps (Disabled)** to set the throttle rate to 0, effectively stopping all replication network traffic.
 - To enforce the change temporarily, select **Clear at next scheduled throttle event**.
 - b. To clear an override previously set, select **Clear Throttle Override**.
3. Select **OK**.

Changing network settings

Using the bandwidth and network-delay settings together, replication calculates the proper TCP (transmission control protocol) buffer size for replication usage. These network settings are global to the DD system and should be set only once per system.

Note the following:

- 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. This 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.

Procedure

1. Select **Replication > Advanced Settings > Change Network Settings** to display the Network Settings dialog.
2. In the Network Settings area, select **Custom Values**.
3. Enter **Delay** and **Bandwidth** values in the text boxes. The network delay setting is in milliseconds, and bandwidth is in bytes per second.
4. In the Listen Port area, enter a new value in the text box. 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 DD system.
5. Select **OK**. The new settings appear in the Network Settings table.

Monitoring replication

The DD System Manager provides many ways to track the status of replication – from checking replication pair status, to tracking backup jobs, to checking performance, to tracking a replication process.

Checking replication pair status

Several places in the Replication area of the DD System Manager provide status updates for replication pairs.

Procedure

1. Select **Replication > Topology**.
2. Check the colors of the arrows, which provide the status of the context.
3. Select the **Summary** tab.
4. 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.

Viewing estimated completion time for backup jobs

You can use the Completion Predictor to see the estimated time for when a backup replication job will be completed.

Procedure

1. Select **Replication > Summary**.
2. Select a Replication context for which to display Detailed Information.
3. 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 job will finish its replication to the destination. If the replication is finished, the area shows `Completed`.

Checking replication context performance

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.

Tracking status of a replication process

To display the progress of a replication initialization, resynchronization, or recovery operation, use the **Replication > Summary** view to check the current state.

CLI Equivalent

```
# replication show config all
CTX Source Destination
Connection Host and Port Enabled
-----
1 dir://host2/backup/dir2 dir://host3/backup/dir3
host3.company.com Yes
2 dir://host3/backup/dir3 dir://host2/backup/dir2
host3.company.com Yes
```

When specifying an IP version, use the following command to check its setting:

```
# replication show config rctx://2
CTX: 2
Source: mtree://ddbeta1.dallasrdc.com/data/coll/EDM1
Destination: mtree://ddbeta2.dallasrdc.com/data/coll/EDM_ipv6
Connection Host: ddbeta2-ipv6.dallasrdc.com
Connection Port: (default)
Ipsversion: ipv6
Low-bw-optim: disabled
Encryption: disabled
Enabled: yes
Propagate-retention-lock: enabled
```

Replication with HA

Floating IP addresses allow HA systems to specify a single IP address for replication configuration that will work regardless of which node of the HA pair is active.

Over IP networks, HA systems use a floating IP address to provide data access to the Data Domain HA pair, regardless of which physical node is the active node. The net config command provides the `[type {fixed | floating}]` option to configure a

floating IP address. The *Data Domain Operating System Command Reference Guide* provides more information.

If a domain name is needed to access the floating IP address, specify the HA system name as the domain name. Run the `ha status` command to locate the HA system name.

Note

Run the `net show hostname type ha-system` command to display the HA system name, and if required, run the `net set hostname ha-system` command to change the HA system name.

All file system access should be through the floating IP address. When configuring backup and replication operations on an HA pair, always specify the floating IP address as the IP address for the Data Domain system. Data Domain features such as DD Boost and replication will accept the floating IP address for the HA pair the same way as they accept the system IP address for a non-HA system.

Replication between HA and non-HA systems

If you want to set up a replication between a high-availability (HA) system and a system running DD OS 5.7.0.3 or earlier, you must create and manage that replication on the HA system if you want to use the DD System Manager graphical user interface (GUI).

However, you can perform replications from a non-HA system to an HA system using the CLI as well as from the HA system to the non-HA system.

Collection replication between HA and non-HA systems is not supported. Directory or MTree replication is required to replicate data between HA and non-HA systems.

Replicating a system with quotas to one without

Replicate a Data Domain system with a DD OS that supports quotas, to a system with a DD OS that does not have quotas.

- 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.

Note

Quotas were introduced as of DD OS 5.2.

Replication Scaling Context

The Replication Scaling Context feature gives you more flexibility when configuring replication contexts.

In environments with more than 299 replication contexts that include both directory and MTree replication contexts, this feature allows you to configure the contexts in any order. Previously, you had to configure the directory replication contexts first, followed by the MTree replication contexts.

The total number of replication contexts cannot exceed 540.

Note

This feature appears only on Data Domain systems running DD OS version 6.0.

Directory-to-MTree replication migration

The directory-to-MTree (D2M) replication optimization feature allows you to migrate existing directory replication contexts to new replication contexts based on MTrees, which are logical partitions of the file system. This feature also lets you monitor the process as it unfolds and verify that has successfully completed.

The D2M feature is compatible with Data Domain Operating System versions 6.0, 5.7, and 5.6.

The source Data Domain system must be running DD OS 6.0 to use this feature, but the destination system can be running 6.0, 5.7, or 5.6. However, the performance optimization benefits are seen only when both the source and destination systems are also running 6.0.

Note

Although you can use the graphical user interface (GUI) for this operation, it is recommended you use the Command Line Interface (CLI) for optimal performance.

Performing migration from directory replication to MTree replication

Do not shut down or reboot your system during directory-to-MTree (D2M) migration.

Procedure

1. Stop all ingest operations to the directory replication source directory.
2. Create an MTree on the source DD system: `mtree create /data/co11/
mtree-name`

Note

Do not create the MTree on the destination DD system.

3. (Optional) Enable DD Retention Lock on the MTree.

Note

If the source system contains retention-locked files, you might want to maintain DD Retention Lock on the new MTree.

See [Enabling DD Retention Lock Compliance on an MTree](#).

4. Create the MTree replication context on both the source and destination DD systems: `replication add source mtree://source-system-name/source
mtree replication add destination mtree://destination-system-
name/destination mtree`
5. Start the D2M migration: `replication dir-to-mtree start from rctx://1
to rctx://2`

In the previous example,
`rctx://1`

refers to the directory replication context, which replicates the directory backup backup/dir1 on the source system;

rctx://2

refers to the MTree replication context, which replicates the MTree /data/coll/mtree1 on the source system.

Note

This command might take longer than expected to complete. Do not press Ctrl-C during this process; if you do, you will cancel the D2M migration.

```
Phase 1 of 4 (precheck):
    Marking source directory /backup/dir1 as read-only...Done.

Phase 2 of 4 (sync):
    Syncing directory replication context...0 files flushed.
current=45 sync_target=47 head=47
current=45 sync_target=47 head=47
Done. (00:09)

Phase 3 of 4 (fastcopy):
    Starting fastcopy from /backup/dir1 to /data/coll/
mtree1...
    Waiting for fastcopy to complete...(00:00)
    Fastcopy status: fastcopy /backup/dir1 to /data/coll/
mtree1: copied 24
files, 1 directory in 0.13 seconds
    Creating snapshot 'REPL-D2M-
mtree1-2015-12-07-14-54-02'...Done

Phase 4 of 4 (initialize):
    Initializing MTree replication context...
(00:08) Waiting for initialize to start...
(00:11) Initialize started.

Use 'replication dir-to-mtree watch rctx://2' to monitor
progress.
```

Viewing directory-to-MTree migration progress

You can see which stage of the migration is currently in progress in the directory-to-MTree (D2M) replication.

Procedure

1. Enter `replication dir-to-mtree watch rctx://2` to see the progress.

```
rctx://2
```

specifies the replication context.

You should see the following output:

```
Use Control-C to stop monitoring.
Phase 4 of 4 (initialize).
(00:00) Replication initialize started...
(00:02) initializing:
(00:14)      100% complete, pre-comp: 0 KB/s, network: 0 KB/
s
(00:14) Replication initialize completed.
Migration for ctx 2 successfully completed.
```

Checking the status of directory-to-MTree replication migration

You can use the `replication dir-to-mtree status` command to check whether the directory-to-MTree migration (D2M) has successfully completed.

Procedure

1. Enter the following command; here, `rctx://2` represents the MTree replication context on the source system: `replication dir-to-mtree status rctx://2`

The output should be similar to the following:

```
Directory Replication CTX:      1
MTree Replication CTX:        2
Directory Replication Source:  dir://127.0.0.2/backup/dir1
MTree Replication Source:     mtree://127.0.0.2/data/
coll/mtreel
MTree Replication Destination: mtree://127.0.0.3/data/
coll/mtreel
Migration Status:              completed
```

If there is no migration in progress, you should see the following:

```
# replication dir-to-mtree status rctx://2
No migration status for context 2.
```

2. Begin ingesting data to the MTree on the source DD system when the migration process is complete.
3. (Optional) Break the directory replication context on the source and target systems.

See the *Data Domain Operating System Version 6.0 Command Reference Guide* for more information about the `replication break` command.

Aborting D2M replication

If necessary, you can abort the directory-to-MTree (D2M) migration procedure.

The `replication dir-to-mtree abort` command aborts the ongoing migration process and reverts the directory from a read-only to a read-write state.

Procedure

1. In the Command-Line Interface (CLI), enter the following command; here, `rctx://2` is the MTree replication context: `replication dir-to-mtree abort rctx://2`

You should see the following output:

```
Canceling directory to MTree migration for context dir-name.
Marking source directory dir-name as read-write...Done.
The migration is now aborted.
Remove the MTree replication context and MTree on both source
and destination
```

```
host by running 'replication break' and 'mtree delete'
commands.
```

2. Break the MTree replication context: `replication break rctx://2`
3. Delete the MTree on the source system: `mtree delete mtree-path`

Troubleshooting D2M

If you encounter a problem setting directory-to-MTree (D2M) replication, there is an operation you can perform to address several different issues.

The `dir-to-mtree abort` procedure can help cleanly abort the D2M process. You should run this procedure in the following cases:

- The status of the D2M migration is listed as aborted.
- The Data Domain system rebooted during D2M migration.
- An error occurred when running the `replication dir-to-mtree start` command.
- Ingest was not stopped before beginning migration.
- The MTree replication context was initialized before the `replication dir-to-mtree start` command was entered.

Note

Do not run `replication break` on the MTree replication context before the D2M process finishes.

Always run `replication dir-to-mtree abort` before running the `replication break` command on the `mrepl ctx`.

Running the `replication break` command prematurely will permanently render the `drepl` source directory as read-only.

If this occurs, please contact Support.

Procedure

1. Enter `replication dir-to-mtree abort` to abort the process.
2. Break the newly created MTree replication context on both the source and destination Data Domain systems.

In the following example, the MTree replication context is

```
rctx://2
```

```
.
```

```
replication break rctx://2
```

3. Delete the corresponding MTrees on both the source and destination systems.

```
mtree delete mtree-path
```

Note

MTrees marked for deletion remain in the file system until the `filesys clean` command is run.

See the *Data Domain Operating System Version 6.0 Command Reference Guide* for more information.

4. Run the `filesys clean start` command on both the source and destination systems.

For more information on the `filesys clean` commands, see the *Data Domain Operating System Version 6.0 Command Reference Guide*.

5. Restart the process.

See [Performing migration from directory replication to MTree replication](#).

Additional D2M troubleshooting

There are solutions available if you forgot to enable DD Retention Lock for the new MTree or an error occurs after directory-to-MTree migration has been initialized.

DD Retention Lock has not been enabled

If you forgot to enable DD Retention Lock for the new MTree and the source directory contains retention-locked files or directories, you have the following options:

- Let the D2M migration continue. However, you will not have DD Retention Lock information in the MTree after the migration.
- Abort the current D2M process as described in [Aborting D2M replication](#) on page 438 and restart the process with DD Retention Lock enabled on the source MTree.

An error occurs after initialization

If the `replication dir-to-mtree start` process finishes without error but you detect an error during the MTree replication initialization (phase 4 of the D2M migration process), you can perform the following steps:

1. Make sure there is no network issue.
2. Initialize the MTree replication context.

Using collection replication for disaster recovery with SMT

To use the destination system of a collection replication pair configured with SMT as a replacement system for disaster recovery, additional SMT configuration steps must be performed in addition to the other configuration steps required to bring a replacement system online.

Before you begin

Using the collection replication destination system in this manner requires autosupport reports to be configured and saved. The KB article *Collection replica with smt enabled*, available on <https://support.emc.com>, provides additional information.

The replacement system will not have the following SMT details:

- Alert notification lists for each tenant-unit
- All users assigned to the DD Boost protocol for use by SMT tenants, if DD Boost is configured on the system
- The default-tenant-unit associated with each DD Boost user, if any, if DD Boost is configured on the system

Complete the following steps to configure SMT on the replacement system.

Procedure

1. In the autosupport report, locate the output for the `smt tenant-unit show` detailed command.

```
Tenant-unit: "tul"
Summary:
Name      Self-Service  Number of Mtrees  Types          Pre-Comp (GiB)
-----
tul       Enabled       2                  DD Boost       2.0
-----

Management-User:
User      Role
-----
tul_ta   tenant-admin
tul_tu   tenant-user
tum_ta   tenant-admin
-----

Management-Group:
Group     Role
-----
qatest   tenant-admin
-----

DDBoost:
Name      Pre-Comp (GiB)  Status  User  Tenant-Unit
-----
sul       2.0             RW/Q    ddbul tul
-----

Q      : Quota Defined
RO     : Read Only
RW     : Read Write

Getting users with default-tenant-unit tul
DD Boost user  Default tenant-unit
-----
ddbul         tul
-----

Mtrees:
Name          Pre-Comp (GiB)  Status  Tenant-Unit
-----
/data/coll/m1      0.0  RW/Q    tul
/data/coll/sul    2.0  RW/Q    tul
-----

D      : Deleted
Q      : Quota Defined
RO     : Read Only
RW     : Read Write
RD     : Replication Destination
RLGE   : Retention-Lock Governance Enabled
RLGD   : Retention-Lock Governance Disabled
RLCE   : Retention-Lock Compliance Enabled

Quota:
Tenant-unit: tul
Mtree      Pre-Comp (MiB)  Soft-Limit (MiB)  Hard-Limit (MiB)
-----
/data/coll/m1      0                71680             81920
/data/coll/sul    2048              30720             51200
-----

Alerts:
Tenant-unit: "tul"
Notification list "tul_grp"
Members
-----
tom.tenant@abc.com
-----
```

```
No such active alerts.
```

2. On the replacement system, enable SMT if it is not already enabled.
3. On the replacement system, license and enable DD Boost if it is required and not already enabled.
4. If DD Boost is configured, assign each user listed in the `DD Boost` section of the "smt tenant-unit show detailed" output as a DD Boost User.


```
# ddbboost user assign ddbul
```
5. If DD Boost is configured, assign each user listed in the `DD Boost` section of the `smt tenant-unit show detailed` output to the default tenant-unit shown, if any, in the output.


```
# ddbboost user option set ddbul default-tenant-unit tul
```
6. Create a new alert notification group with the same name as the alert notification group in the `Alerts` section of the `smt tenant-unit show detailed` output.


```
# alert notify-list create tul_grp tenant-unit tul
```
7. Assign each email address in the alert notification group in the `Alerts` section of the `smt tenant-unit show detailed` output to the new alert notification group.


```
# alert notify-list add tul_grp emails tom.tenant@abc.com
```

CHAPTER 17

DD Secure Multitenancy

This chapter includes:

- [Data Domain Secure Multi-tenancy overview](#).....444
- [Provisioning a Tenant Unit](#)..... 447
- [Enabling Tenant Self-Service mode](#)..... 451
- [Data access by protocol](#)..... 451
- [Data management operations](#).....453

Data Domain Secure Multi-tenancy overview

Data Domain *Secure Multi-tenancy (SMT)* is the simultaneous hosting, by an internal IT department or an external provider, of an IT infrastructure for more than one consumer or workload (business unit, department, or Tenant).

SMT provides the ability to securely isolate many users and workloads in a shared infrastructure, so that the activities of one Tenant are not apparent or visible to the other Tenants.

A *Tenant* is a consumer (business unit, department, or customer) who maintains a persistent presence in a hosted environment.

Within an enterprise, a Tenant may consist of one or more business units or departments on a Data Domain system configured and managed by IT staff.

- For a business unit (BU) use case, the Finance and Human Resources departments of a corporation could share the same Data Domain system, but each department would be unaware of the presence of the other.
- For a service provider (SP) use case, the SP could deploy one or more Data Domain systems to accommodate different Protection Storage services for multiple end-customers.

Both use cases emphasize the segregation of different customer data on the same physical Data Domain system.

SMT architecture basics

Secure Multitenancy (SMT) provides a simple approach to setting up Tenants and Tenant Units, using MTrees. SMT setup is performed using DD Management Center and/or the DD OS command line interface. This administration guide provides the theory of SMT and some general command line instructions.

The basic architecture of SMT is as follows.

- A Tenant is created on the DD Management Center and/or DD system.
- A Tenant Unit is created on a DD system for the Tenant.
- One or more MTrees are created to meet the storage requirements for the Tenant's various types of backups.
- The newly created MTrees are added to the Tenant Unit.
- Backup applications are configured to send each backup to its configured Tenant Unit MTree.

Note

For more information about DD Management Center, see the *DD Management Center User Guide*. For more information about the DD OS command line interface, see the *DD OS Command Reference*.

Terminology used in Secure Multi-Tenancy (SMT)

Understanding the terminology that is used in SMT will help you better understand this unique environment.

MTrees

MTrees are logical partitions of the file system and offer the highest degree of management granularity, meaning users can perform operations on a specific MTree

without affecting the entire file system. MTrees are assigned to Tenant Units and contain that Tenant Unit's individualized settings for managing and monitoring SMT.

Multi-Tenancy

Multi-Tenancy refers to the hosting of an IT infrastructure by an internal IT department, or an external service provider, for more than one consumer/workload (business unit/department/Tenant) simultaneously. Data Domain SMT enables *Data Protection-as-a-Service*.

RBAC (role-based access control)

RBAC offers multiple roles with different privilege levels, which combine to provide the administrative isolation on a multi-tenant Data Domain system. (The next section will define these roles.)

Storage Unit

A *Storage Unit* is an MTree configured for the DD Boost protocol. Data isolation is achieved by creating a Storage Unit and assigning it to a DD Boost user. The DD Boost protocol permits access only to Storage Units assigned to DD Boost users connected to the Data Domain system.

Tenant

A *Tenant* is a consumer (business unit/department/customer) who maintains a persistent presence in a hosted environment.

Tenant Self-Service

Tenant Self-Service is a method of letting a Tenant log in to a Data Domain system to perform some basic services (add, edit, or delete local users, NIS groups, and/or AD groups). This reduces the bottleneck of always having to go through an administrator for these basic tasks. The Tenant can access only their assigned Tenant Units. Tenant Users and Tenant Admins will, of course, have different privileges.

Tenant Unit

A *Tenant Unit* is the partition of a Data Domain system that serves as the unit of administrative isolation between Tenants. Tenant units that are assigned to a tenant can be on the same or different Data Domain systems and are secured and logically isolated from each other, which ensures security and isolation of the control path when running multiple Tenants simultaneously on the shared infrastructure. Tenant Units can contain one or more *MTrees*, which hold all configuration elements that are needed in a multi-tenancy setup. Users, management-groups, notification-groups, and other configuration elements are part of a Tenant Unit.

Control path and network isolation

Control path isolation is achieved by providing the user roles of *tenant-admin* and *tenant-user* for a Tenant Unit. *Network isolation* for data and administrative access is achieved by associating a fixed set of *data access IP address(es)* and *management IP address(es)* with a Tenant Unit.

The *tenant-admin* and *tenant-user* roles are restricted in scope and capability to specific Tenant Units and to a restricted set of operations they can perform on those Tenant Units. To ensure a logically secure and isolated data path, a system administrator must configure one or more Tenant Unit MTrees for each protocol in an SMT environment. Supported protocols include DD Boost, NFS, CIFS, and DD VTL. Access is strictly regulated by the native access control mechanisms of each protocol.

Tenant-self-service sessions (through ssh) can be restricted to a fixed set of *management IP address(es)* on a DD system. Administrative access sessions (through ssh/http/https) can also be restricted to a fixed set of management IP address(es) on DD systems. By default, however, there are no management IP address(es) associated

with a Tenant Unit, so the only standard restriction is through the use of the *tenant-admin* and *tenant-user* roles. You must use `smt tenant-unit management-ip` to add and maintain management IP address(es) for Tenant Units.

Similarly, data access and data flow (into and out of Tenant Units) can be restricted to a fixed set of local or remote *data access IP address(es)*. The use of assigned data access IP address(es) enhances the security of the DD Boost and NFS protocols by adding SMT-related security checks. For example, the list of storage units returned over DD Boost RPC can be limited to those which belong to the Tenant Unit with the assigned local data access IP address. For NFS, access and visibility of exports can be filtered based on the local data access IP address(es) configured. For example, using `showmount -e` from the local data access IP address of a Tenant Unit will only display NFS exports belonging to that Tenant Unit.

The *sysadmin* must use `smt tenant-unit data-ip` to add and maintain data access IP address(es) for Tenant Units.

Note

If you attempt to mount an MTree in an SMT using a non-SMT IP address, the operation will fail.

If multiple Tenant Units are belong to the same tenant, they can share a default gateway. However, if multiple Tenant Units that belong to different tenants are oprevented from using the same default gateway.

Multiple Tenant Units belonging to the same tenant can share a default gateway. Tenant Units that belong to different tenants cannot use the same default gateway.

Understanding RBAC in SMT

In SMT (Secure Multi-Tenancy), permission to perform a task depends on the role that is assigned to a user. DD Management Center uses RBAC (role-based access control) to control these permissions.

All DD Management Center users can:

- View all Tenants
- Create, read, update, or delete Tenant Units belonging to any Tenant if the user is an administrator on the Data Domain system hosting the Tenant Unit
- Assign and unassign Tenant Units to and from a Tenant if the user is an administrator on the Data Domain system hosting the Tenant Unit
- View Tenant Units belonging to any Tenant if the user has any assigned role on the Data Domain system hosting the Tenant Unit

To perform more advanced tasks depends on the role of the user, as follows:

admin role

A user with an *admin* role can perform all administrative operations on a Data Domain system. An *admin* can also perform all SMT administrative operations on a Data Domain system, including setting up SMT, assigning SMT user roles, enabling Tenant self-service mode, creating a Tenant, and so on. In the context of SMT, the *admin* is typically referred to as the *landlord*. In DD OS, the role is known as the *sysadmin*.

To have permission to edit or delete a Tenant, you must be both a DD Management Center *admin* and a DD OS *sysadmin* on all Data Domain systems that are associated with the Tenant Units of that Tenant. If the Tenant does not have any Tenant Units, you need only to be a DD Management Center *admin* to edit or delete that Tenant.

limited-admin role

A user with a *limited-admin* role can perform all administrative operations on a Data Domain system as the *admin*. However, users with the *limited-admin* role cannot delete or destroy MTrees. In DD OS, there is an equivalent *limited-admin* role.

tenant-admin role

A user with a *tenant-admin* role can perform certain tasks only when *tenant self-service* mode is enabled for a specific Tenant Unit. Responsibilities include scheduling and running a backup application for the Tenant and monitoring resources and statistics within the assigned Tenant Unit. The *tenant-admin* can view audit logs, but RBAC ensures that only audit logs from the Tenant Unit(s) belonging to the *tenant-admin* are accessible. In addition, *tenant-admins* ensure administrative separation when Tenant self-service mode is enabled. In the context of SMT, the *tenant-admin* is usually referred to as the *backup admin*.

tenant-user role

A user with a *tenant-user* role can monitor the performance and usage of SMT components only on Tenant Unit(s) assigned to them and only when Tenant self-service is enabled, but a user with this role cannot view audit logs for their assigned Tenant Units. In addition, *tenant-users* may run the `show` and `list` commands.

none role

A user with a role of *none* is not allowed to perform any operations on a Data Domain system other than changing their password and accessing data using DD Boost. However, after SMT is enabled, the *admin* can select a user with a *none* role from the Data Domain system and assign them an SMT-specific role of *tenant-admin* or *tenant-user*. Then, that user can perform operations on SMT management objects.

management groups

BSPs (backup service providers) can use *management groups* defined in a single, external AD (active directory) or NIS (network information service) to simplify managing user roles on Tenant Units. Each BSP Tenant may be a separate, external company and may use a name-service such as AD or NIS.

With SMT management groups, the AD and NIS servers are set up and configured by the *admin* in the same way as SMT local users. The *admin* can ask their AD or NIS administrator to create and populate the group. The *admin* then assigns an SMT role to the entire group. Any user within the group who logs in to the Data Domain system is logged in with the role that is assigned to the group.

When users leave or join a Tenant company, they can be removed or added to the group by the AD or NIS administrator. It is not necessary to modify the RBAC configuration on a Data Domain system when users who are part of the group are added or removed.

Provisioning a Tenant Unit

Launching the configuration wizard begins the initial provisioning procedure for Secure Multitenancy (SMT). During the procedure, the wizard creates and provisions a new Tenant Unit based on Tenant configuration requirements. Information is entered by the administrator, as prompted. After completing the procedure, the administrator proceeds to the next set of tasks, beginning with enabling Tenant Self-Service mode. Following the initial setup, manual procedures and configuration modifications may be performed as required.

Procedure

1. Start SMT.

```
# smt enable
SMT enabled.
```

2. Verify that SMT is enabled.

```
# smt status
SMT is enabled.
```

3. Launch the SMT configuration wizard.

```
# smt tenant-unit setup
No tenant-units.
```

4. Follow the configuration prompts.

SMT TENANT-UNIT Configuration

```
Configure SMT TENANT-UNIT at this time (yes|no) [no]: yes
```

```
Do you want to create new tenant-unit (yes/no)? : yes
```

Tenant-unit Name

```
Enter tenant-unit name to be created
```

```
: SMT_5.7_tenant_unit
```

```
Invalid tenant-unit name.
```

```
Enter tenant-unit name to be created
```

```
: SMT_57_tenant_unit
```

Pending Tenant-unit Settings

```
Create Tenant-unit SMT_57_tenant_unit
```

```
Do you want to save these settings (Save|Cancel|Retry): save
```

```
SMT Tenant-unit Name Configurations saved.
```

SMT TENANT-UNIT MANAGEMENT-IP Configuration

```
Configure SMT TENANT-UNIT MANAGEMENT-IP at this time (yes|no) [no]: yes
```

```
Do you want to add a local management ip to this tenant-unit? (yes|no) [no]: yes
```

port	enabled	state	DHCP	IP address	netmask /prefix length	type	additional setting
ethMa	yes	running	no	192.168.10.57 fe80::260:16ff:fe49:f4b0**	255.255.255.0 /64	n/a	
eth3a	yes	running	ipv4	192.168.10.236* fe80::260:48ff:fe1c:60fc**	255.255.255.0* /64	n/a	
eth3b	yes	running	no	192.168.50.57 fe80::260:48ff:fe1c:60fd**	255.255.255.0 /64	n/a	
eth4b	yes	running	no	192.168.60.57 fe80::260:48ff:fe1f:5183**	255.255.255.0 /64	n/a	

```
* Value from DHCP
```

```
** auto_generated IPv6 address
```

```
Choose an ip from above table or enter a new ip address. New ip addresses will need to be created manually.
```

Ip Address

```
Enter the local management ip address to be added to this tenant-unit
```

```
: 192.168.10.57
```

```
Do you want to add a remote management ip to this tenant-unit? (yes|no) [no]:
```

Pending Management-ip Settings

```
Add Local Management-ip 192.168.10.57
```

```
Do you want to save these settings (Save|Cancel|Retry): yes  
unrecognized input, expecting one of Save|Cancel|Retry
```

```
Do you want to save these settings (Save|Cancel|Retry): save
```

```
Local management access ip "192.168.10.57" added to tenant-unit "SMT_57_tenant_unit".
```


SMT Tenant-unit Management-IP Configurations saved.

SMT TENANT-UNIT MANAGEMENT-IP Configuration

Do you want to add another local management ip to this tenant-unit? (yes|no) [no]:

Do you want to add another remote management ip to this tenant-unit? (yes|no) [no]:

SMT TENANT-UNIT DDBOOST Configuration

Configure SMT TENANT-UNIT DDBOOST at this time (yes|no) [no]:

SMT TENANT-UNIT MTREE Configuration

Configure SMT TENANT-UNIT MTREE at this time (yes|no) [no]: yes

Name	Pre-Comp (GiB)	Status	Tenant-Unit
/data/coll/laptop_backup	4846.2	RO/RD	-
/data/coll/random	23469.9	RO/RD	-
/data/coll/software2	2003.7	RO/RD	-
/data/coll/tsm6	763704.9	RO/RD	-

D : Deleted

Q : Quota Defined

RO : Read Only

RW : Read Write

RD : Replication Destination

RLGE : Retention-Lock Governance Enabled

RLGD : Retention-Lock Governance Disabled

RLCE : Retention-Lock Compliance Enabled

Do you want to assign an existing MTree to this tenant-unit? (yes|no) [no]:

Do you want to create a mtree for this tenant-unit now? (yes|no) [no]: yes

MTree Name

Enter MTree name

: SMT_57_tenant_unit

Invalid mtree path name.

Enter MTree name

:

SMT_57_tenant_unit

Invalid mtree path name.

Enter MTree name

: /data/coll/SMT_57_tenant_unit

MTree Soft-Quota

Enter the quota soft-limit to be set on this MTree (<n> {MiB|GiB|TiB|PiB}|none)

:

MTree Hard-Quota

Enter the quota hard-limit to be set on this MTree (<n> {MiB|GiB|TiB|PiB}|none)

:

Pending MTree Settings

Create MTree /data/coll/SMT_57_tenant_unit

MTree Soft Limit none

MTree Hard Limit none

Do you want to save these settings (Save|Cancel|Retry): save

MTree "/data/coll/SMT_57_tenant_unit" created successfully.

MTree "/data/coll/SMT_57_tenant_unit" assigned to tenant-unit "SMT_57_tenant_unit".

SMT Tenant-unit MTree Configurations saved.

SMT TENANT-UNIT MTREE Configuration

Name	Pre-Comp (GiB)	Status	Tenant-Unit
------	----------------	--------	-------------

```

/data/coll/laptop_backup          4846.2   RO/RD   -
/data/coll/random                 23469.9 RO/RD   -
/data/coll/software2              2003.7   RO/RD   -
/data/coll/tsm6                   763704.9 RO/RD   -
-----
D      : Deleted
Q      : Quota Defined
RO     : Read Only
RW     : Read Write
RD     : Replication Destination
RLGE   : Retention-Lock Governance Enabled
RLGD   : Retention-Lock Governance Disabled
RLCE   : Retention-Lock Compliance Enabled

Do you want to assign another MTree to this tenant-unit? (yes|no) [no]: yes

Do you want to assign an existing MTree to this tenant-unit? (yes|no) [no]:

Do you want to create another mtree for this tenant-unit? (yes|no) [no]:

SMT TENANT-UNIT SELF-SERVICE Configuration

Configure SMT TENANT-UNIT SELF-SERVICE at this time (yes|no) [no]: yes
Self-service of this tenant-unit is disabled

Do you want to enable self-service of this tenant-unit? (yes|no) [no]: yes

Do you want to configure a management user for this tenant-unit? (yes|no) [no]:

Do you want to configure a management group for this tenant-unit (yes|no) [no]: yes

Management-Group Name
Enter the group name to be assigned to this tenant-unit
: SMT_57_tenant_unit_group

What role do you want to assign to this group (tenant-user|tenant-admin) [tenant-user]:
tenant-admin

Management-Group Type
What type do you want to assign to this group (nis|active-directory)?
: nis

Pending Self-Service Settings
Enable Self-Service          SMT_57_tenant_unit
Assign Management-group      SMT_57_tenant_unit_group
Management-group role        tenant-admin
Management-group type        nis

Do you want to save these settings (Save|Cancel|Retry): save
Tenant self-service enabled for tenant-unit "SMT_57_tenant_unit"
Management group "SMT_57_tenant_unit_group" with type "nis" is assigned to tenant-unit
"SMT_57_tenant_unit" as "tenant-admin".

SMT Tenant-unit Self-Service Configurations saved.

SMT TENANT-UNIT SELF-SERVICE Configuration

Do you want to configure another management user for this tenant-unit? (yes|no) [no]:

Do you want to configure another management group for this tenant-unit? (yes|no)
[no]:

SMT TENANT-UNIT ALERT Configuration

Configure SMT TENANT-UNIT ALERT at this time (yes|no) [no]: yes
No notification lists.

Alert Configuration

Alert Group Name

```

```

Specify alert notify-list group name to be created
: SMT_57_tenant_unit_notify

Alert email addresses
Enter email address to receive alert for this tenant-unit
: dd_proserv@emc.com

Do you want to add more emails (yes/no)?
: no

Pending Alert Settings
Create Notify-list group   SMT_57_tenant_unit_notify
Add emails                 dd_proserv@emc.com

Do you want to save these settings (Save|Cancel|Retry): save
Created notification list "SMT_57_tenant_unit_notify" for tenant "SMT_57_tenant_unit".
Added emails to notification list "SMT_57_tenant_unit_notify".

SMT Tenant-unit Alert Configurations saved.

Configuration complete.

```

Enabling Tenant Self-Service mode

For administrative separation of duties and delegation of administrative/management tasks to implement Tenant Self-Service, which is required for control path isolation, the system administrator can enable this mode on a Tenant Unit and then assign users to manage the unit in the roles of tenant-admin or tenant-user. These roles allow users other than the administrator to perform specific tasks on the Tenant Unit to which they are assigned. In addition to administrative separation, Tenant Self-Service mode helps reduce the management burden on internal IT and service provider staff.

Procedure

1. View Tenant Self-Service mode status for one or all Tenant Units.

```
# smt tenant-unit option show { tenant-unit | all }
```

2. Enable Tenant Self-Service mode on the selected Tenant Unit.

```
# smt tenant-unit option set tenant-unit self-service { enabled
| disabled }
```

Data access by protocol

Secure data paths, with protocol-specific access controls, enable security and isolation for Tenant Units. In a Secure Multitenancy (SMT) environment, data access protocol management commands are also enhanced with a Tenant Unit parameter to enable consolidated reporting.

DD systems support multiple data access protocols simultaneously, including DD Boost, NFS, CIFS, and DD VTL. A DD system can present itself as an application-specific interface, such as a file server offering NFS or CIFS access over the Ethernet, a DD VTL device, or a DD Boost device.

The native access control mechanisms of each supported protocol ensure that the data paths for each Tenant remain separate and isolated. Such mechanisms include access control lists (ACLs) for CIFS, exports for NFS, DD Boost credentials, and Multi-User Boost credential-aware access control.

Multi-User DD Boost and Storage Units in SMT

When using Multi-User DD Boost with SMT (Secure Multi-Tenancy), user permissions are set by Storage Unit ownership.

Multi-User DD Boost refers to the use of multiple DD Boost user credentials for DD Boost Access Control, in which each user has a separate username and password.

A *Storage Unit* is an MTree configured for the DD Boost protocol. A user can be associated with, or “own,” one or more Storage Units. Storage Units that are owned by one user cannot be owned by another user. Therefore, only the user owning the Storage Unit can access the Storage Unit for any type of data access, such as backup/restore. The number of DD Boost user names cannot exceed the maximum number of MTrees. (See the “MTrees” chapter in this book for the current maximum number of MTrees for each DD model.) Storage Units that are associated with SMT must have the *none* role assigned to them.

Each backup application must authenticate using its DD Boost username and password. After authentication, DD Boost verifies the authenticated credentials to confirm ownership of the Storage Unit. The backup application is granted access to the Storage Unit only if the user credentials presented by the backup application match the user names associated with the Storage Unit. If user credentials and user names do not match, the job fails with a permission error.

Configuring access for CIFS

Common Internet File System (CIFS) is a file-sharing protocol for remote file access. In a Secure Multitenancy (SMT) configuration, backup and restores require client access to the CIFS shares residing in the MTree of the associated Tenant Unit. Data isolation is achieved using CIFS shares and CIFS ACLs.

Procedure

1. Create an MTree for CIFS and assign the MTree to the tenant unit.

```
# mtree create mtree-path tenant-unit tenant-unit
```

2. Set capacity soft and hard quotas for the MTree.

```
# mtree create mtree-path tenant-unit tenant-unit] [quota-soft-limit n{MiB|GiB|TiB|PiB} ] [quota-hard-limit n {MiB|GiB|TiB|PiB}
```

3. Create a CIFS share for *pathname* from the MTree.

```
# cifs share create share path pathname clients clients
```

Configuring NFS access

NFS is a UNIX-based, file-sharing protocol for remote file access. In a Secure Multitenancy (SMT) environment, backup and restores require client access to the NFS exports residing in the MTree of the associated Tenant Unit. Data isolation is achieved using NFS exports and network isolation. NFS determines if an MTree is associated with a network-isolated Tenant Unit. If so, NFS verifies the connection properties associated with the Tenant Unit. Connection properties include the destination IP address and interface or client hostname.

Procedure

1. Create an MTree for NFS and assign the MTree to the tenant unit.

```
# mtree create mtree-path tenant-unit tenant-unit
```

2. Set capacity soft and hard quotas for the MTree.

```
# mtree create mtree-path tenant-unit tenant-unit] [quota-soft-limit n{MiB|GiB|TiB|PiB} ] [quota-hard-limit n {MiB|GiB|TiB|PiB}
```

3. Create an NFS export by adding one or more clients to the MTree.

```
# nfs add path client-list
```

Configuring access for DD VTL

DD VTL Tenant data isolation is achieved using DD VTL access groups that create a virtual access path between a host system and the DD VTL. (The physical Fibre Channel connection between the host system and DD VTL must already exist.)

Placing tapes in the DD VTL allows them to be written to, and read by, the backup application on the host system. DD VTL tapes are created in a DD VTL pool, which is an MTree. Because DD VTL pools are MTrees, the pools can be assigned to Tenant Units. This association enables SMT monitoring and reporting.

For example, if a tenant-admin is assigned a Tenant Unit that contains a DD VTL pool, the tenant-admin can run MTree commands to display read-only information. Commands can run only on the DD VTL pool assigned to the Tenant Unit.

These commands include:

- `mtree list` to view a list of MTrees in the Tenant Unit
- `mtree show compression` to view statistics on MTree compression
- `mtree show performance` to view statistics on performance

Output from most `list` and `show` commands include statistics that enable service providers to measure space usage and calculate chargeback fees.

DD VTL operations are unaffected and continue to function normally.

Using DD VTL NDMP TapeServer

DD VTL Tenant data isolation is also achieved using NDMP. DD OS implements a NDMP (Network Data Management Protocol) tape server that allows NDMP-capable systems to send backup data to the DD system via a three-way NDMP backup.

The backup data is written to virtual tapes (which are in a pool) by a DD VTL assigned to the special DD VTL group *TapeServer*.

Because the backup data is written to tapes in a pool, information in the DD VTL topic regarding MTrees also applies to the Data Domain NDMP TapeServer.

Data management operations

Secure Multitenancy (SMT) management operations include monitoring Tenant Units and other objects, such as Storage Units and MTrees. For some SMT objects, additional configuration or modification may also be required.

Collecting performance statistics

Each MTree can be measured for performance or “usage” statistics and other real-time information. Historical consumption rates are available for DD Boost Storage Units. Command output lets the tenant-admin collect usage statistics and compression ratios for an MTree associated with a Tenant Unit, or for all MTrees and associated Tenant Units. Output may be filtered to display usage in intervals ranging from minutes to months. Results are passed to the administrator, who uses the

statistics as a chargeback metric. A similar method is used to gather usage statistics and compression ratios for Storage Units.

Procedure

1. Collect MTree real-time performance statistics.

```
# mtree show stats
```

2. Collect performance statistics for MTrees associated with a Tenant Unit.

```
# mtree show performance
```

3. Collect compression statistics for MTrees associated with a Tenant Unit.

```
# mtree show compression
```

Modifying quotas

To meet QoS criteria, a system administrator uses DD OS “knobs” to adjust the settings required by the Tenant configuration. For example, the administrator can set “soft” and “hard” quota limits on DD Boost Storage Units. Stream “soft” and “hard” quota limits can be allocated only to DD Boost Storage Units assigned to Tenant Units. After the administrator sets the quotas, the tenant-admin can monitor one or all Tenant Units to ensure no single object exceeds its allocated quotas and deprives others of system resources.

Quotas are set initially when prompted by the configuration wizard, but they can be adjusted or modified later. The example below shows how to modify quotas for DD Boost. (You can also use `quota capacity` and `quota streams` to deal with capacity and stream quotas and limits.)

Procedure

1. To modify soft and hard quota limits on DD Boost Storage Unit “su33”:

```
ddboost storage-unit modify su33 quota-soft-limit 10 Gib quota-hard-limit 20 Gib
```

2. To modify stream soft and hard limits on DD Boost Storage Unit “su33”:

```
ddboost storage-unit modify su33 write-stream-soft-limit 20 read-stream-soft-limit 6 repl -stream-soft-limit 20 combined-stream-soft-limit 20
```

3. To report physical size for DD Boost Storage Unit “su33”:

```
ddboost storage-unit modify su33 report-physical-size 8 GiB
```

SMT and replication

In case of disaster, user roles dictate how a user can assist in data recovery operations. Several replication types are available in an SMT configuration. (See the *DD Replicator* chapter for more detail on how to perform replication.)

Here are some points to consider regarding user roles:

- The admin can recover MTrees from a replicated copy.
- The tenant-admin can replicate MTrees from one system to another, using DD Boost managed file replication.
- The tenant-admin can recover MTrees from a replicated copy, also by using DD Boost managed file replication.

Collection replication

Collection replication replicates core Tenant Unit configuration information.

Secure replication over public internet

To protect against man-in-the-middle (MITM) attacks when replicating over a public internet connection, authentication includes validating SSL certificate-related information at the replication source and destination.

MTree replication (NFS/CIFS) using DD Boost managed file replication

MTree replication is supported on MTrees assigned to Tenant Units, using DD Boost managed file replication. During MTree replication, an MTree assigned to a Tenant Unit on one system can be replicated to an MTree assigned to a Tenant Unit on another system. MTree replication is not allowed between two different Tenants on the two DD systems. When security mode is set to *strict*, MTree replication is allowed only when the MTrees belong to same Tenants.

For backward compatibility, MTree replication from an MTree assigned to a Tenant Unit to an unassigned MTree is supported, but must be configured manually. Manual configuration ensures the destination MTree has the correct settings for the Tenant Unit. Conversely, MTree replication from an unassigned MTree to an MTree assigned to a Tenant Unit is also supported.

When setting up SMT-aware MTree replication, *security mode* defines how much checking is done on the Tenant. The *default* mode checks that the source and destination do not belong to different Tenants. The *strict* mode makes sure the source and destination belong to the same Tenant. Therefore, when you use strict mode, you must create a Tenant on the destination machine with the same UUID as the UUID of the Tenant on the source machine that is associated with the MTree being replicated.

DD Boost managed file replication (also with DD Boost AIR)

DD Boost managed file replication is supported between Storage Units, regardless of whether one Storage Unit, or both, are assigned to Tenant Units.

During DD Boost managed file replication, Storage Units are not replicated in total. Instead, certain files within a Storage Unit are selected by the backup application for replication. The files selected in a Storage Unit and assigned to a Tenant Unit on one system can be replicated to a Storage Unit assigned to a Tenant Unit on another system.

For backward compatibility, selected files in a Storage Unit assigned to a Tenant Unit can be replicated to an unassigned Storage Unit. Conversely, selected files in an unassigned Storage Unit can be replicated to a Storage Unit assigned to a Tenant Unit.

DD Boost managed file replication can also be used in DD Boost AIR deployments.

Replication control for QoS

An upper limit on replication throughput (`repl-in`) can be specified for an MTree. Since MTrees for each tenant are assigned to a Tenant Unit, each tenant's replication resource usage can be capped by applying these limits. The relation of this feature to SMT is that MTree Replication is subject to this throughput limit.

SMT Tenant alerts

A DD system generates *events* when it encounters potential problems with software or hardware. When an event is generated, an *alert* notification is sent immediately via email to members designated in the notification list and to the Data Domain administrator.

SMT alerts are specific to each Tenant Unit and differ from DD system alerts. When Tenant Self-Service mode is enabled, the tenant-admin can choose to receive alerts about the various system objects he or she is associated with and any critical events,

such as an unexpected system shutdown. A tenant-admin may only view or modify notification lists to which he or she is associated.

The example below shows a sample alert. Notice that the two event messages at the bottom of the notification are specific to a Multi-Tenant environment (indicated by the word “Tenant”). For the entire list of DD OS and SMT alerts, see the *Data Domain MIB Quick Reference Guide* or the SNMP MIB.

```
EVT-ENVIRONMENT-00021 - Description: The system has been shutdown by
abnormal method; for example, not by one of the following: 1) Via
IPMI chassis control command 2) Via power button 3) Via OS shutdown.
```

```
Action: This alert is expected after loss of AC (main power) event.
If this shutdown is not expected and persists, contact your
contracted support provider or visit us online at https://
my.datadomain.com.
```

```
Tenant description: The system has experienced an unexpected power
loss and has restarted.
```

```
Tenant action: This alert is generated when the system restarts after
a power loss. If this alert repeats, contact your System
Administrator.
```

Managing snapshots

A *snapshot* is a read-only copy of an MTree captured at a specific point in time. A snapshot can be used for many things, for example, as a restore point in case of a system malfunction. The required role for using `snapshot` is `admin` or `tenant-admin`.

To view snapshot information for an MTree or a Tenant Unit:

```
# snapshot list mtree mtree-path | tenant-unit tenant-unit
```

To view a snapshot schedule for an MTree or a Tenant Unit:

```
# snapshot schedule show [name | mtree mtree-list | tenant-
unit tenant-unit]
```

Performing a file system Fast Copy

A Fast Copy operation clones files and directory trees of a source directory to a target directory on a DD system. There are special circumstances regarding Fast Copy with Secure Multitenancy (SMT).

Here are some considerations when performing a file system Fast Copy with Tenant Self-Service mode enabled:

- A tenant-admin can Fast Copy files from one Tenant Unit to another when the tenant-admin is the tenant-admin for both Tenant Units, and the two Tenant Units belong to the same Tenant.
- A tenant-admin can Fast Copy files within the same Tenant Unit.
- A tenant-admin can Fast Copy files within the Tenant Units at source and destination.

To perform a file system Fast Copy:

```
# fileys fastcopy source <src> destination <dest>
```


CHAPTER 18

DD Cloud Tier

This chapter includes:

- [DD Cloud Tier overview](#) 458
- [Configuring Cloud Tier](#) 460
- [Configuring cloud units](#) 461
- [Data movement](#) 470
- [Using the Command Line Interface \(CLI\) to configure DD Cloud Tier](#) 474
- [Configuring encryption for DD cloud units](#) 478
- [Information needed in the event of system loss](#) 478
- [Using DD Replicator with Cloud Tier](#) 479
- [Using DD Virtual Tape Library \(VTL\) with Cloud Tier](#) 479
- [Displaying capacity consumption charts for DD Cloud Tier](#) 479
- [DD Cloud Tier logs](#) 480
- [Using the Command Line Interface \(CLI\) to remove DD Cloud Tier](#) 480

DD Cloud Tier overview

DD Cloud Tier is a native feature of DD OS 6.0 (or higher) for moving data from the active tier to low-cost, high-capacity object storage in the public, private, or hybrid cloud for long-term retention. DD Cloud Tier is best suited for long-term storage of infrequently accessed data that is being held for compliance, regulatory, and governance reasons. The ideal data for DD Cloud Tier is data that is past its normal recovery window.

DD Cloud Tier is managed using a single Data Domain namespace. There is no separate cloud gateway or virtual appliance required. Data movement is supported by the native Data Domain policy management framework. Conceptually, the cloud storage is treated as an additional storage tier (DD Cloud Tier) attached to the Data Domain system, and data is moved between tiers as needed. File system metadata associated with the data stored in the cloud is maintained in local storage, and also mirrored to the cloud. The metadata that resides in local storage facilitates operations such as deduplication, cleaning, Fast Copy, and replication. This local storage is divided into self-contained buckets, called cloud units, for ease of manageability.

Supported platforms

Cloud Tier is supported on physical platforms that have the necessary memory, CPU, and storage connectivity to accommodate another storage tier.

DD Cloud Tier is supported on these systems:

Table 190 DD Cloud Tier supported configurations

Model	Memory	Cloud capacity	Required number of SAS I/O modules	Supported disk shelf types for metadata storage	Number of ES30 shelves or DS60 disk packs required	Required capacity for metadata storage
DD990	256 GB	1140 TB	4	ES30	4	60 x 3 TB HDDs = 180 TB
DD3300 4 TB	16 GB	8 TB	N/A	N/A	N/A	1 x 1 TB virtual disk = 1 TB
DDD3300 16 TB	48 GB	32 TB	N/A	N/A	N/A	2 x 1 TB virtual disks = 2 TB
DD3300 32 TB	64 GB	64 TB	N/A	N/A	N/A	4 x 1 TB virtual disks = 4 TB
DD4200	128 GB	378 TB	3	DS60 or ES30	2	30 x 3 TB HDDs = 90 TB
DD4500	192 GB	570 TB	3	DS60 or ES30	2	30 x 4 TB HDDs = 120 TB
DD6800	192 GB	576 TB	2	DS60 or ES30	2	30 x 4 TB HDDs = 120 TB
DD7200	256 GB	856 TB	4	DS60 or ES30	4	60 x 4 TB HDDs = 240 TB

Table 190 DD Cloud Tier supported configurations (continued)

Model	Memory	Cloud capacity	Required number of SAS I/O modules	Supported disk shelf types for metadata storage	Number of ES30 shelves or DS60 disk packs required	Required capacity for metadata storage
DD9300	384 GB	1400 TB	2	DS60 or ES30	4	60 x 4 TB HDDs = 240 TB
DD9500	512 GB	1728 TB	4	DS60 or ES30	5	75 x 4 TB HDDs = 300 TB
DD9800	768 GB	2016 TB	4	DS60 or ES30	5	75 x 4 TB HDDs = 300 TB
DD VE 16 TB	32 GB	32 TB	N/A	N/A	N/A	1 x 500 GB virtual disk = 500 GB ^a
DD VE 64 TB	60 GB	128 TB	N/A	N/A	N/A	1 x 500 GB virtual disk = 500 GB ^a
DD VE 96 TB	80 GB	192 TB	N/A	N/A	N/A	1 x 500 GB virtual disk = 500 GB ^a

- a. The minimum metadata size is a hard limit. Data Domain recommends users start with 1 TB for metadata storage and expand in 1 TB increments. The *Data Domain Virtual Edition Installation and Administration Guide* provides more details about using DD Cloud Tier with DD VE.

Note

DD Cloud Tier is supported with Data Domain High Availability (HA). Both nodes must be running DD OS 6.0 (or higher), and they must be HA-enabled.

Note

DD Cloud Tier is not supported on any system not listed and is not supported on any system with the Extended Retention feature enabled.

Note

The Cloud Tier feature may consume all available bandwidth in a shared WAN link, especially in a low bandwidth configuration (1 Gbps), and this may impact other applications sharing the WAN link. If there are shared applications on the WAN, the use of QoS or other network limiting is recommended to avoid congestion and ensure consistent performance over time.

If bandwidth is constrained, the rate of data movement will be slow and you will not be able to move as much data to the cloud. It is best to use a dedicated link for data going to the Cloud Tier.

Note

Do not send traffic over onboard management network interface controllers (ethMx interfaces).

DD Cloud Tier performance

The Data Domain system uses internal optimizations to maximize DD Cloud Tier performance.

Large object size

DD Cloud Tier uses object sizes of 1 MB or 4 MB (depending on the cloud storage provider) to reduce the metadata overhead, and lower the number of objects to migrate to cloud storage.

Configuring Cloud Tier

To configure Cloud Tier, add the license and enclosures, set a system passphrase, and create a file system with support for data movement to the cloud.

- For Cloud Tier, the cloud capacity license is required.
- To license Cloud Tier, refer to the applicable *Data Domain Operating System Release Notes* for the most up-to-date information on product features, software updates, software compatibility guides, and information about Data Domain products, licensing, and service.
- To set a system passphrase, use the **Administration > Access > Administrator Access** tab.
If the system passphrase is not set, the **Set Passphrase** button appears in the Passphrase area. If a system passphrase is configured, the **Change Passphrase** button appears, and your only option is to change the passphrase.
- To configure storage, use the **Hardware > Storage** tab.
- To create a file system, use the File System Create Wizard.

Configuring storage for DD Cloud Tier

Cloud Tier storage on the DD system is required for the cloud units—it holds the metadata for the files, while the data resides in the cloud.

Procedure

1. Select **Hardware > Storage**.
2. In the Overview tab, expand **Cloud Tier**.
3. Click **Configure**.

The Configure Cloud Tier dialog box is displayed.

4. Select the checkbox for the shelf to be added from the Addable Storage section.



DD3300 systems require the use of 1 TB storage devices for DD Cloud Tier metadata storage.

5. Click the **Add to Tier** button.

6. Click **Save** to add the storage.
7. Select **Data Management > File System** and enable the Cloud Tier feature.
8. Click **Disable** (at the bottom of the screen) to disable the file system.
9. Click **OK**.
10. After the file system is disabled, select **Enable Cloud Tier**.
To enable the cloud tier, you must meet the storage requirement for the licensed capacity. Configure the cloud tier of the file system. Click **Next**.
A cloud file system requires a local store for a local copy of the cloud metadata.
11. Select **Enable file system**.
The cloud tier is enabled with the designated storage.
12. Click **OK**.
You must create cloud units separately, after the file system is created.

Configuring cloud units

The cloud tier consists of a maximum of two cloud units, and each cloud unit is mapped to a cloud provider, enabling multiple cloud providers per Data Domain system. The Data Domain system must be connected to the cloud and have an account with a supported cloud provider.

Configuring cloud units includes these steps:

- Configuring the network, including firewall and proxy settings
- Importing CA certificates
- Adding cloud units

Firewall and proxy settings

- Network firewall ports
 - Port 443 (HTTPS) and/or Port 80 (HTTP) must be open to the cloud provider networks for both the endpoint IP and the provider authentication IP for bi-directional traffic.
For example, for Amazon S3, both `s3-ap-southeast-1.amazonaws.com` and `s3.amazonaws.com` must have port 80 and/or port 443 unblocked and set to allow bi-directional IP traffic.

Note

Several public cloud providers use IP ranges for their endpoint and authentication addresses. In this situation, the IP ranges used by the provider need to be unblocked to accommodate potential IP changes.

- Remote cloud provider destination IP and access authentication IP address ranges must be allowed through the firewall.
- For ECS private cloud, local ECS authentication and web storage (S3) access IP ranges and ports 9020 (HTTP) and 9021 (HTTPS) must be allowed through local firewalls.

Note

ECS private cloud load balancer IP access and port rules must also be configured.

- Proxy settings
 - If there are any existing proxy settings that cause data above a certain size to be rejected, those settings must be changed to allow object sizes up to 4.5MB.
 - If customer traffic is being routed through a proxy, the self-signed/CA-signed proxy certificate must be imported. See "Importing CA certificates" for details.
 - OpenSSL cipher suites
 - Ciphers - ECDHE-RSA-AES256-SHA384, AES256-GCM-SHA384
-

Note

Default communication with all cloud providers is initiated with strong cipher.

- TLS Version: 1.2
 - Supported protocols
 - HTTP & HTTPS
-

Note

Default communication with all public cloud providers occurs on secure HTTP (HTTPS), but you can overwrite the default setting to use HTTP.

Importing CA certificates

Before you can add cloud units for Elastic Cloud Storage (ECS), Virtustream Storage Cloud, Amazon Web Services S3 (AWS), and Azure cloud, you must import CA certificates.

Before you begin

For AWS, Virtustream, and Azure public cloud providers, root CA certificates can be downloaded from <https://www.digicert.com/digicert-root-certificates.htm>.

- For an AWS cloud provider, download the Baltimore CyberTrust Root certificate.
- For a Virtustream cloud provider, download the DigiCert High Assurance EV Root CA certificate.
- For ECS, the root certificate authority will vary by customer. Implementing cloud storage on ECS requires a load balancer. If an HTTPS endpoint is used as an endpoint in the configuration, be sure to import the root CA certificate. Contact your load balancer provider for details.
- For an Azure cloud provider, download the Baltimore CyberTrust Root certificate.
- For an S3 Flexible provider, import the root CA certificate. Contact your S3 Flexible provider for details.

If your downloaded certificate has a .crt extension, it will likely need to be converted to a PEM-encoded certificate. If so, use OpenSSL to convert the file from .crt format to .pem (for example, `openssl x509 -inform der -in BaltimoreCyberTrustRoot.crt -out BaltimoreCyberTrustRoot.pem`).

Procedure

1. Select **Data Management > File System > Cloud Units**.
2. In the tool bar, click **Manage Certificates**.
The Manage Certificates for Cloud dialog is displayed.
3. Click **Add**.
4. Select one of these options:
 - **I want to upload the certificate as a .pem file.**
Browse to and select the certificate file.
 - **I want to copy and paste the certificate text.**
 - Copy the contents of the .pem file to your copy buffer.
 - Paste the buffer into the dialog.
5. Click **Add**.

Adding a cloud unit for Elastic Cloud Storage (ECS)

A Data Domain system or DD VE instance requires a close time synchronization with the ECS system to configure a Data Domain cloud unit. Configuring NTP on the Data Domain system or DD VE instance, and the ECS system addresses this issue.

Procedure

1. Select **Data Management > File System > Cloud Units**.
2. Click **Add**.
The Add Cloud Unit dialog is displayed.
3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.
The remaining fields in the Add Cloud Unit dialog pertain to the cloud provider account.
4. For **Cloud provider**, select **EMC Elastic Cloud Storage (ECS)** from the drop-down list.
5. Enter the provider **Access key** as password text.
6. Enter the provider **Secret key** as password text.
7. Enter the provider **Endpoint** in this format: `http://<ip/hostname>:<port>`. If you are using a secure endpoint, use `https` instead.

Note

Implementing cloud storage on ECS requires a load balancer.

By default, ECS runs the S3 protocol on port 9020 for HTTP and 9021 for HTTPS. With a load balancer, these ports are sometimes remapped to 80 for HTTP and 443 for HTTPS, respectively. Check with your network administrator for the proper ports.

8. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure for HTTP Proxy Server**.
Enter the proxy hostname, port, user, and password.
9. Click **Add**.

The File System main window now displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Adding a cloud unit for Virtustream

Virtustream offers a range of storage classes. The *Cloud Providers Compatibility Matrix*, available from the Data Protection community page at <https://inside.dell.com/community/active/data-protection> provides the most up to date information about the supported storage classes.

The following endpoints are used by the Virtustream cloud provider, depending on storage class and region. Be sure that DNS is able to resolve these hostnames before configuring cloud units.

- s-us.objectstorage.io
- s-eu.objectstorage.io
- s-eu-west-1.objectstorage.io
- s-eu-west-2.objectstorage.io
- s-us-central-1.objectstorage.io

Procedure

1. Select **Data Management > File System > Cloud Units**.
2. Click **Add**.

The Add Cloud Unit dialog is displayed.

3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.

The remaining fields in the Add Cloud Unit dialog pertain to the cloud provider account.

4. For **Cloud provider**, select **Virtustream Storage Cloud** from the drop-down list.
5. Select the storage class from the drop-down list.
6. Select the appropriate region corresponding to the type of account from the drop-down list.
7. Enter the provider **Access key** as password text.
8. Enter the provider **Secret key** as password text.
9. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.

Enter the proxy hostname, port, user, and password.

10. Click **Save**.

The File System main window now displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Adding a cloud unit for Amazon Web Services S3

AWS offers a range of storage classes. The *Cloud Providers Compatibility Matrix*, available from the Data Protection community page at <https://inside.dell.com/community/active/data-protection> provides the most up to date information about the supported storage classes.

For enhanced security, the Cloud Tier feature uses Signature Version 4 for all AWS requests. Signature Version 4 signing is enabled by default.

The following endpoints are used by the AWS cloud provider, depending on storage class and region. Be sure that DNS is able to resolve these hostnames before configuring cloud units.

- s3.amazonaws.com
- s3-us-west-1.amazonaws.com
- s3-us-west-2.amazonaws.com
- s3-eu-west-1.amazonaws.com
- s3-ap-northeast-1.amazonaws.com
- s3-ap-southeast-1.amazonaws.com
- s3-ap-southeast-2.amazonaws.com
- s3-sa-east-1.amazonaws.com
- ap-south-1
- ap-northeast-2
- eu-central-1

Note

The China region is not supported.

Note

The AWS user credentials must have permissions to create and delete buckets and to add, modify, and delete files within the buckets they create. S3FullAccess is preferred, but these are the minimum requirements:

- CreateBucket
- ListBucket
- DeleteBucket
- ListAllMyBuckets
- GetObject
- PutObject
- DeleteObject

Procedure

1. Select **Data Management > File System > Cloud Units**.
2. Click **Add**.
The Add Cloud Unit dialog is displayed.
3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.
The remaining fields in the Add Cloud Unit dialog pertain to the cloud provider account.
4. For **Cloud provider**, select **Amazon Web Services S3** from the drop-down list.
5. Select the storage class from the drop-down list.
6. Select the appropriate **Storage region** from the drop-down list.
7. Enter the provider **Access key** as password text.

8. Enter the provider **Secret key** as password text.
9. Ensure that port 443 (HTTPS) is not blocked in firewalls. Communication with the AWS cloud provider occurs on port 443.
10. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.
Enter the proxy hostname, port, user, and password.
11. Click **Add**.
The File System main window now displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Adding a cloud unit for Azure

Microsoft Azure offers a range of storage account types. The *Cloud Providers Compatibility Matrix*, available from the Data Protection community page at <https://inside.dell.com/community/active/data-protection> provides the most up to date information about the supported storage classes.

The following endpoints are used by the Azure cloud provider, depending on storage class and region. Be sure that DNS is able to resolve these hostnames before configuring cloud units.

- *Account name*.blob.core.windows.net

The account name is obtained from the Azure cloud provider console.

Procedure

1. Select **Data Management > File System > Cloud Units**.
2. Click **Add**.
The Add Cloud Unit dialog is displayed.
3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.
The remaining fields in the Add Cloud Unit dialog pertain to the cloud provider account.
4. For **Cloud provider**, select **Microsoft Azure Storage** from the drop-down list.
5. For **Account type**, select **Government** or **Public**.
6. Select the storage class from the drop-down list.
7. Enter the provider **Account name**.
8. Enter the provider **Primary key** as password text.
9. Enter the provider **Secondary key** as password text.
10. Ensure that port 443 (HTTPS) is not blocked in firewalls. Communication with the Azure cloud provider occurs on port 443.
11. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.
Enter the proxy hostname, port, user, and password.
12. Click **Add**.
The File System main window now displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Adding an S3 Flexible provider cloud unit

The Cloud Tier feature supports additional qualified S3 cloud providers under an S3 Flexible provider configuration option.

The S3 Flexible provider option supports the standard and standard-infrequent-access storage classes. The endpoints will vary depending on cloud provider, storage class and region. Be sure that DNS is able to resolve these hostnames before configuring cloud units.

Procedure

1. Select **Data Management > File System > Cloud Units**.
2. Click **Add**.
The Add Cloud Unit dialog is displayed.
3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.
The remaining fields in the Add Cloud Unit dialog pertain to the cloud provider account.
4. For **Cloud provider**, select **Flexible Cloud Tier Provider Framework for S3** from the drop-down list.
5. Enter the provider **Access key** as password text.
6. Enter the provider **Secret key** as password text.
7. Specify the appropriate **Storage region**.
8. Enter the provider **Endpoint** in this format: `http://<ip/hostname>:<port>`. If you are using a secure endpoint, use `https` instead.
9. For **Storage class**, select the appropriate storage class from the drop-down list.
10. Ensure that port 443 (HTTPS) is not blocked in firewalls. Communication with the S3 cloud provider occurs on port 443.
11. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.
Enter the proxy hostname, port, user, and password.
12. Click **Add**.

The File System main window now displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Modifying a cloud unit or cloud profile

Modify cloud unit credentials, an S3 Flexible provider name, or details of a cloud profile.

Modifying cloud unit credentials

Procedure

1. Select **Data Management > File System > Cloud Units**.
2. Click the pencil icon for the cloud unit whose credentials you want to modify.
The Modify Cloud Unit dialog is displayed.

3. For **Account name**, enter the new account name.
4. For **Access key**, enter the new provider access key as password text.

Note

Modifying the access key is not supported for ECS environments.

5. For **Secret key**, enter the new provider secret key as password text.
6. For **Primary key**, enter the new provider primary key as password text.

Note

Modifying the primary key is only supported for Azure environments.

7. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.
8. Click **OK**.

Modifying an S3 Flexible provider name

Procedure

1. Select **Data Management > File System > Cloud Units**.
2. Click the pencil icon for the S3 Flexible cloud unit whose name you want to modify.
The Modify Cloud Unit dialog is displayed.
3. For **S3 Provider Name**, enter the new provider name.
4. Click **OK**.

Using the CLI to modify a cloud profile

Procedure

1. Run the `cloud profile modify` command to modify the details of a cloud profile. The system prompts you to modify individual details of the cloud profile.

For Virtustream, AWS S3, or Azure profiles, run this command to add a storage class to an existing cloud profile.

The profile details that can be modified depend on the cloud provider:

- ECS supports modification of the secret key.
- Virtustream supports modification of the access key, and secret key.
- AWS S3 supports modification of the access key, and secret key.
- Azure supports modification of the access key, secret key, and primary key.
- S3 Flexible modification of the access key, secret key, and provider name.

Deleting a cloud unit

This operation results in the loss of all data in the cloud unit selected for deletion. Be sure to delete all files before deleting the cloud units.

Before you begin

- Check if data movement to the cloud is running (CLI command: `data-movement status`). If it is, stop data movement using the “`data-movement stop`” CLI command.
- Check if cloud cleaning is running for this cloud unit (CLI command: `cloud clean status`). If it is, stop cloud cleaning using the “`cloud clean`” CLI command.
- Check if a data movement policy is configured for this cloud unit (CLI command: `data-movement policy show`). If it is, remove this policy using the “`data-movement policy reset`” CLI command.

Procedure

1. Use the following CLI command to identify files in the cloud unit.

```
# fileysys report generate file-location
```

2. Delete the files that are in the cloud unit to be deleted.
3. Use the following CLI command to run cloud cleaning.

```
# cloud clean start unit-name
```

Wait for cleaning to complete. The cleaning may take time depending on how much data is present in the cloud unit.

4. Disable the file system.
5. Use the following CLI command to delete the cloud unit.

```
# cloud unit del unit-name
```

Internally, this marks the cloud unit as `DELETE_PENDING`.

6. Use the following CLI command to validate that the cloud unit is in the `DELETE_PENDING` state.

```
# cloud unit list
```

7. Enable the file system.

The file system initiates the procedure in the background to delete any remaining objects from the buckets in the cloud for this cloud unit and then delete the buckets. This process can take a long time, depending on how many objects were remaining in these buckets. Until the bucket cleanup completes, this cloud unit continues to consume a slot on the Data Domain system, which may prevent creation of a new cloud unit if both slots are occupied.

8. Periodically check the state using this CLI command:

```
# cloud unit list
```

The state remains `DELETE_PENDING` while the background cleanup is running.

9. Verify from the cloud provider S3 portal that all corresponding buckets have been deleted and the associated space has been freed up.

10. If needed, reconfigure data movement policies for affected MTrees and restart data movement.

Results

If you have difficulty completing this procedure, contact Support.

Data movement

Data is moved from the active tier to the cloud tier as specified by your individual data movement policy. The policy is set on a per-MTree basis. Data movement can be initiated manually or automatically using a schedule.

Adding data movement policies to MTrees

A file is moved from the Active to the Cloud Tier based on the date it was last modified. For data integrity, the entire file is moved at this time. The *Data Movement Policy* establishes the file age threshold, age range, and the destination.

Note

A data movement policy cannot be configured for the /backup MTree.

Procedure

1. Select **Data Management > MTree**.
2. In the top panel, select the MTree to which you want to add a data movement policy.
3. Click the **Summary** tab.
4. Under **Data Movement Policy** click **Add**.
5. For **File Age in Days**, set the file age threshold (**Older than**) and optionally, the age range (**Younger than**).

Note

The minimum number of days for **Older than** is 14. For nonintegrated backup applications, files moved to the cloud tier cannot be accessed directly and need to be recalled to the active tier before you can access them. So, choose the age threshold value as appropriate to minimize or avoid the need to access a file moved to the cloud tier.

6. For **Destination**, specify the destination cloud unit.
7. Click **Add**.

Moving data manually

You can start and stop data movement manually. Any MTree that has a valid data movement policy has its files moved.

Procedure

1. Select **Data Management > File System**.
2. At the bottom of the page, click **Show Status of File System Services**.

These status items are displayed:

- File System
 - Physical Capacity Measurement
 - Data Movement
 - Active Tier Cleaning
3. For **Data Movement**, click **Start**.

Moving data automatically

You can move data automatically, using a schedule and a throttle. Schedules can be daily, weekly, or monthly.

Procedure

1. Select **Data Management > File System > Settings**.
2. Click the **Data Movement** tab.
3. Set the throttle and schedule.

Note

The throttle is for adjusting resources for internal Data Domain processes; it does not affect network bandwidth.

Note

If a cloud unit is inaccessible when cloud tier data movement runs, the cloud unit is skipped in that run. Data movement on that cloud unit occurs in the next run if the cloud unit becomes available. The data movement schedule determines the duration between two runs. If the cloud unit becomes available and you cannot wait for the next scheduled run, you can start data movement manually.

Recalling a file from the Cloud Tier

For nonintegrated backup applications, you must recall the data to the active tier before you can restore the data. Backup administrators must trigger a recall or backup applications must perform a recall before cloud-based backups can be restored. Once a file is recalled, its aging is reset and will start again from 0, and the file will be eligible based on the age policy set. A file can be recalled on the source MTree only. Integrated applications can recall a file directly.

Note

If a file resides only in a snapshot, it cannot be recalled directly. To recall a file in a snapshot, use fastcopy to copy the file from the snapshot back to the active MTree, then recall the file from the cloud. A file can only be recalled from the cloud to an active MTree.

Procedure

1. Select **Data Management > File System > Summary**.
2. Do one of the following:
 - In the Cloud Tier section of the Space Usage panel, click **Recall**.

- Expand the File System status panel at the bottom of the screen and click **Recall**.

Note

The **Recall** link is available only if a cloud unit is created and has data.

3. In the Recall File from Cloud dialog, enter the exact file name (no wildcards) and full path of the file to be recalled, for example: `/data/coll/mt11/file1.txt`. Click **Recall**.
4. To check the status of the recall, do one of the following:
 - In the Cloud Tier section of the Space Usage panel, click **Details**.
 - Expand the File System status panel at the bottom of the screen and click **Details**.

The Cloud File Recall Details dialog is displayed, showing the file path, cloud provider, recall progress, and amount of data transferred. If there are unrecoverable errors during the recall, an error message is displayed. Hover your cursor over the error message to display a tool tip with more details and possible corrective actions.

Results

Once the file has been recalled to the active tier, you can restore the data.

Note

For nonintegrated applications, once a file has been recalled from the cloud tier to the active tier, a minimum of 14 days must elapse before the file is eligible for data movement. After 14 days, normal data movement processing will occur for the file. This restriction does not apply to integrated applications.

Note

For data-movement, nonintegrated applications configure an age-based data movement policy on the Data Domain system to specify which files get migrated to the cloud tier, and this policy applies uniformly to all files in an MTree. Integrated applications use an application-managed data movement policy, which lets you identify specific files to be migrated to the cloud tier.

Using the CLI to recall a file from the cloud tier

For nonintegrated backup applications, you must recall the data to the active tier before you can restore the data. Backup administrators must trigger a recall or backup applications must perform a recall before cloud-based backups can be restored. Once a file is recalled, its aging is reset and will start again from 0, and the file will be eligible based on the age policy set. A file can be recalled on the source MTree only. Integrated applications can recall a file directly.

Note

If a file resides only in a snapshot, it cannot be recalled directly. To recall a file in a snapshot, use `fastcopy` to copy the file from the snapshot back to the active MTree, then recall the file from the cloud. A file can only be recalled from the cloud to an active MTree.

Procedure

1. Check the location of the file using:

```
fileSYS report generate file-location [path {<path-name> |
all}] [output-file <filename>]
```

The pathname can be a file or directory; if it is a directory, all files in the directory are listed.

Filename	Location
/data/coll/mt11/file1.txt	Cloud Unit 1

2. Recall the file using:

```
data-movement recall path <path-name>
```

This command is asynchronous, and it starts the recall.

```
data-movement recall path /data/coll/mt11/file1.txt
Recall started for "/data/coll/mt11/file1.txt".
```

3. Monitor the status of the recall using

```
data-movement status [path {pathname | all | [queued]
[running] [completed] [failed]} | to-tier cloud | all]
```

```
data-movement status path /data/coll/mt11/file1.txt
Data-movement recall:
-----
Data-movement for "/data/coll/mt11/file1.txt": phase 2 of 3
(Verifying)
80% complete; time: phase XX:XX:XX total XX:XX:XX
Copied (post-comp): XX XX, (pre-comp) XX XX
```

If the status shows that the recall isn't running for a given path, the recall may have finished, or it may have failed.

4. Verify the location of the file using

```
fileSYS report generate file-location [path {<path-name> |
all}] [output-file <filename>]
```

Filename	Location
/data/coll/mt11/file1.txt	Active

Results

Once the file has been recalled to the active tier, you can restore the data.

Note

For nonintegrated applications, once a file has been recalled from the cloud tier to the active tier, a minimum of 14 days must elapse before the file is eligible for data movement. After 14 days, normal data movement processing will occur for the file. This restriction does not apply to integrated applications.

Note

For data-movement, nonintegrated applications configure an age-based data movement policy on the Data Domain system to specify which files get migrated to the cloud tier, and this policy applies uniformly to all files in an MTree. Integrated applications use an application-managed data movement policy, which lets you identify specific files to be migrated to the cloud tier.

Direct restore from the cloud tier

Direct restore lets nonintegrated applications read files directly from the cloud tier without going through the active tier.

Key considerations in choosing to use direct restore include:

- Direct restore does not require an integrated application and is transparent for nonintegrated applications.
- Reading from the cloud tier does not require copying first into the active tier.
- Histograms and statistics are available for tracking direct reads from the cloud tier.
- Direct restore is supported only for ECS cloud providers.
- Applications do experience cloud tier latency.
- Reading directly from the cloud tier is not bandwidth optimized.
- Direct restore supports a small number of jobs.

Direct restore is useful with nonintegrated applications that do not need to know about the cloud tier and won't need to restore cloud files frequently.

Using the Command Line Interface (CLI) to configure DD Cloud Tier

You can use the Data Domain Command Line Interface to configure DD Cloud Tier.

Procedure

1. Configure storage for both active and cloud tier. As a prerequisite, the appropriate capacity licenses for both the active and cloud tiers must be installed.
 - a. Ensure licenses for the features CLOUDTIER-CAPACITY and CAPACITY-ACTIVE are installed. To check the ELMS license:

```
# elicense show
```

If the license is not installed, use the `elicense update` command to install the license. Enter the command and paste the contents of the license file after this prompt. After pasting, ensure there is a carriage return, then press `Control-D` to save. You are prompted to replace licenses, and after answering yes, the licenses are applied and displayed.

```
# elicense update
```

```
Enter the content of license file and then press Control-D,  
or press Control-C to cancel.
```

- b. Display available storage:

```
# storage show all  
# disk show state
```

- c. Add storage to the active tier:

```
# storage add enclosures <enclosure no> tier active
```

- d. Add storage to the cloud tier:

```
# storage add enclosures <enclosure no> tier cloud
```

2. Install certificates.

Before you can create a cloud profile, you must install the associated certificates.

For AWS, Virtustream, and Azure public cloud providers, root CA certificates can be downloaded from <https://www.digicert.com/digicert-root-certificates.htm>.

- For an AWS cloud provider, download the Baltimore CyberTrust Root certificate.
- For a Virtustream cloud provider, download the DigiCert High Assurance EV Root CA certificate.
- For an Azure cloud provider, download the Baltimore CyberTrust Root certificate.
- For ECS, the root certificate authority will vary by customer. Contact your load balancer provider for details.

Downloaded certificate files have a .crt extension. Use openssl on any Linux or Unix system where it is installed to convert the file from .crt format to .pem.

```
$openssl x509 -inform der -in DigiCertHighAssuranceEVRootCA.crt -out DigiCertHighAssuranceEVRootCA.pem
```

```
$openssl x509 -inform der -in BaltimoreCyberTrustRoot.crt -out BaltimoreCyberTrustRoot.pem
```

```
# adminaccess certificate import ca application cloud
Enter the certificate and then press Control-D, or press
Control-C to cancel.
```

3. To configure the Data Domain system for data-movement to the cloud, you must first enable the “cloud” feature and set the system passphrase if it has not already been set.

```
# cloud enable
Cloud feature requires that passphrase be set on the system.
Enter new passphrase:
Re-enter new passphrase:
Passphrases matched.
The passphrase is set.
Encryption is recommended on the cloud tier.
  Do you want to enable encryption? (yes|no) [yes]:
Encryption feature is enabled on the cloud tier.
Cloud feature is enabled.
```

4. Configure the cloud profile using the cloud provider credentials. The prompts and variables vary by provider.

```
# cloud profile add <profilename>
```

Note

For security reasons, this command does not display the access/secret keys you enter.

Select the provider:

```
Enter provider name (aws|azure|ecs|s3_flexible|virtustream)
```

- AWS S3 requires access key, secret key, storage class, and region.
- Azure requires account name, whether or not the account is an Azure Government account, primary key, secondary key, and storage class.

- ECS requires entry of access key, secret key and endpoint.
- S3 Flexible providers require the provider name, access key, secret key, region, endpoint, and storage class.
- Virtustream requires access key, secret key, storage class, and region.

At the end of each profile addition you are asked if you want to set up a proxy. If you do, these values are required: *proxy hostname*, *proxy port*, *proxy username*, and *proxy password*.

5. Verify the cloud profile configuration:

```
# cloud profile show
```

6. Create the active tier file system if it is not already created:

```
# fileysys create
```

7. Enable the file system:

```
# fileysys enable
```

8. Configure the cloud unit:

```
# cloud unit add unitname profile profilename
```

Use the `cloud unit list` command to list the cloud units.

9. Optionally, configure encryption for the cloud unit.

a. Verify that the ENCRYPTION license is installed:

```
# elicence show
```

b. Enable encryption for the cloud unit:

```
# fileysys encryption enable cloud-unit unitname
```

c. Check encryption status:

```
# fileysys encryption status
```

10. Create one or more MTrees:

```
# mtree create /data/col1/mt11
```

11. Verify the DD Cloud Tier configuration:

```
# cloud provider verify
```

```
This operation will perform test data movement after creating a temporary profile and bucket.
```

```
Do you want to continue? (yes|no) [yes]:
```

```
Enter provider name (aws|azure|virtustream|ecs|s3_generic): aws
```

```
Enter the access key:
```

```
Enter the secret key:
```

```
Enter the region (us-east-1|us-west-1|us-west-2|eu-west-1|apnortheast-1|ap-southeast-1|ap-southeast-2|sa-east-1|ap-south-1|ap-northeast-2|eu-central-1):
```

```
Verifying cloud provider ...
```

```
This process may take a few minutes.
```

```
Cloud Enablement Check:
```

```
Checking Cloud feature enabled: PASSED
```

```
Checking Cloud volume: PASSED
```

```
Connectivity Check:
```

```
Checking firewall access: PASSED
```

```

Validating certificate PASSED

Account Validation:
Creating temporary profile: PASSED
Creating temporary bucket: PASSED

S3 API Validation:
Validating Put Bucket: PASSED
Validating List Bucket: PASSED
Validating Put Object: PASSED
Validating Get Object: PASSED
Validating List Object: PASSED
Validating Delete Object: PASSED
Validating Bulk Delete: PASSED

Cleaning Up:
Deleting temporary bucket: PASSED
Deleting temporary profile: PASSED

Provider verification passed.

```

12. Configure the file migration policy for this MTree. You can specify multiple MTrees in this command. The policy can be based on the age threshold or the range.
 - a. To configure the age-threshold (migrating files older than the specified age to cloud):

```
# data-movement policy set age-threshold age_in_days to-tier
cloud cloud-unit unitname mtrees mtreename
```

- b. To configure the age-range (migrating only those files that are in the specified age-range):

```
# data-movement policy set age-range min-age age_in_days max-age
age_in_days to-tier cloud cloud-unit unitname mtrees
mtreename
```

13. Export the file system, and from the client, mount the file system and ingest data into the active tier. Change the modification date on the ingested files such that they now qualify for data migration. (Set the date to older than the age-threshold value specified when configuring the data-movement policy.)
14. Initiate file migration of the aged files. Again, you can specify multiple MTrees with this command.

```
# data-movement start mtrees mtreename
```

To check the status of data-movement:

```
# data-movement status
```

You can also watch the progress of data-movement:

```
# data-movement watch
```

15. Verify that file migration worked and the files are now in the cloud tier:

```
# fileysys report generate file-location path all
```

16. Once you have migrated a file to the cloud tier, you cannot directly read from the file (attempting to do so results in an error). The file can only be recalled back to the active tier. To recall a file to the active tier:

```
# data-movement recall path pathname
```

Configuring encryption for DD cloud units

Encryption can be enabled at three levels: Data Domain system, Active Tier, and cloud unit. Encryption of the Active Tier is only applicable if encryption is enabled for the Data Domain system. Cloud units have separate controls for enabling encryption.

Procedure

1. Select **Data Management > File System > DD Encryption**.

Note

If no encryption license is present on the system, the Add Licenses page is displayed.

2. In the DD Encryption panel, do one of the following:
 - To enable encryption for **Cloud Unit x**, click **Enable**.
 - To disable encryption for **Cloud Unit x**, click **Disable**.

Note

You are prompted to enter security officer credentials to enable encryption.

3. Enter the security officer **Username** and **Password**. Optionally, check **Restart file system now**.
4. Click **Enable** or **Disable**, as appropriate.
5. In the File System Lock panel, lock or unlock the file system.
6. In the Key Management panel, click **Configure**.
7. In the Change Key Manager dialog, configure security officer credentials and the key manager.

Note

Cloud encryption is allowed only through the Data Domain Embedded Key Manager. External key managers are not supported.

8. Click **OK**.
9. Use the DD Encryption Keys panel to configure encryption keys.

Information needed in the event of system loss

Once Cloud Tier is configured on the Data Domain system, record the following information about the system and store it in a safe location apart from the Data Domain system. This information will be needed to recover the Cloud Tier data in case the Data Domain system is lost.

Note

This process is designed for emergency situations only and will involve significant time and effort from the Data Domain engineering staff.

- Serial number of the original Data Domain system
- System passphrase of the original Data Domain system
- DD OS version number of the original Data Domain system
- Cloud Tier profile and configuration information

Using DD Replicator with Cloud Tier

Collection replication is not supported on Cloud Tier enabled Data Domain systems.

Directory replication only works on the /backup MTree, and this MTree cannot be assigned to the Cloud Tier. So, directory replication is not affected by Cloud Tier.

Managed file replication and MTree replication are supported on Cloud Tier enabled Data Domain systems. One or both systems can have Cloud Tier enabled. If the source system is Cloud Tier enabled, data may need to be read from the cloud if the file was already migrated to the Cloud Tier. A replicated file is always placed first in the Active Tier on the destination system even when Cloud Tier is enabled. A file can be recalled from the Cloud Tier back to the Active Tier on the source MTree only. Recall of a file on the destination MTree is not allowed.

Note

If the source system is running DD OS 5.6 or 5.7 and replicating into a Cloud Tier enabled system using MTree replication, the source system must be upgraded to a release that can replicate to a Cloud Tier enabled system. Please see the *DD OS Release Notes* system requirements.

Note

Files in the Cloud Tier cannot be used as base files for virtual synthetic operations. The incremental forever or synthetic full backups need to ensure that the files remain in the Active Tier if they will be used in virtual synthesis of new backups.

Using DD Virtual Tape Library (VTL) with Cloud Tier

On systems configured with Cloud Tier and DD VTL, the cloud storage is supported for use as the VTL vault. To use DD VTL tape out to cloud, license and configure the cloud storage first, and then select it as the vault location for the VTL.

[DD VTL tape out to cloud](#) on page 353 provides additional information about using VTL with Cloud Tier.

Displaying capacity consumption charts for DD Cloud Tier

Three charts are available for displaying Cloud Tier consumption statistics—Space Usage, Consumption, and Daily Written.

Procedure

1. Select **Data Management > File System > Charts**.
2. For **Chart**, select one of the following:
 - Space Usage
 - Consumption

- Daily Written
3. For **Scope**, select **Cloud Tier**.
 - The Space Usage Tab displays space usage over time, in MiB. You can select a duration (one week, one month, three months, one year, or All). The data is presented (color-coded) as pre-compression used (blue), post-compression used (red), and the compression factor (green).
 - The Consumption Tab displays the amount of post-compression storage used and the compression ratio over time, which enables you to analyze consumption trends. You can select a duration (one week, one month, three months, one year, or All). The data is presented (color-coded) as capacity (blue), post-compression used (red), compression factor (green), cleaning (orange) and data movement (violet).
 - The Daily Written Tab displays the amount of data written per day. You can select a duration (one week, one month, three months, one year, or All). The data is presented (color-coded) as pre-compression written (blue), post-compression used (red), and the total compression factor (green).

DD Cloud Tier logs

If DD Cloud Tier suffers a failure of any kind, in configuration or operation, the system automatically creates a folder with a timestamp associated with the time of the failure.

The detailed logs for the DD Cloud Tier failure are created under `/ddvar/log/debug/verify_logs`. Mount the `/ddvar/log/debug` directory to access the logs.

Note

The output of the `log list view` command does not list all the detailed log files created for the DD Cloud Tier failure.

Using the Command Line Interface (CLI) to remove DD Cloud Tier

You can use the Data Domain Command Line Interface to remove the DD Cloud Tier configuration.

Before you begin

Delete all files in the cloud units before removing the DD Cloud Tier configuration from the system. Run the `filesys report generate file-location path all output-file file_loc` command to identify the files in the cloud units, and delete them from the NFS mount points of the MTrees.

Note

The command above creates the report `file_loc` in the `/ddr/var/` directory.

Procedure

1. Disable the file system.

```
# filesys disable
```

```
This action will disable the file system.
Applications may experience interruptions
```



```

while the file system is disabled.
  Are you sure? (yes|no) [no]: yes

ok, proceeding.

Please wait.....
The filesystem is now disabled.

```

2. List the cloud units on the system.

```

# cloud unit list
Name           Profile        Status
-----
cloud_unit-1   cloudProfile   Active
cloud_unit-2   cloudProfile2  Active
-----

```

3. Delete the cloud units individually.

```

# cloud unit del cloud_unit-1

This command irrevocably destroys all data
in the cloud unit "cloud_unit-1".
  Are you sure? (yes|no) [no]: yes

ok, proceeding.

Enter sysadmin password to confirm:

Destroying cloud unit "cloud_unit-1"
Cloud unit 'cloud_unit-1' deleted. The data in the cloud will be deleted asynchronously
on the filesystem startup.

# cloud unit del cloud_unit-2

This command irrevocably destroys all data
in the cloud unit "cloud_unit-2".
  Are you sure? (yes|no) [no]: yes

ok, proceeding.

Enter sysadmin password to confirm:

Destroying cloud unit "cloud_unit-2"
Cloud unit 'cloud_unit-2' deleted. The data in the cloud will be deleted asynchronously
on the filesystem startup.

```

4. Verify the delete operations are in progress.

```

# cloud unit list
Name           Profile        Status
-----
cloud_unit-1   cloudProfile   Delete-Pending
cloud_unit-2   cloudProfile2  Delete-Pending
-----

```

5. Restart the file system.

```

# fileysys enable
Please wait.....
The filesystem is now enabled.

```

6. Run the cloud unit list command to verify that neither cloud unit appears.

Contact Support if one or both cloud units still display with the status Delete-Pending.

7. Identify the disk enclosures that are assigned to DD Cloud Tier.

```
# storage show tier cloud

Cloud tier details:
Disk   Disks                Count  Disk   Additional
Group  -----                -----  Size   Information
-----
dgX    2.1-2.15, 3.1-3.15    30     3.6 TiB
-----
Current cloud tier size: 0.0 TiB
Cloud tier maximum capacity: 108.0 TiB
```

8. Remove the disk enclosures from DD Cloud Tier.

```
# storage remove enclosures 2, 3

Removing enclosure 2...Enclosure 2 successfully removed.
Updating system information...done
Successfully removed: 2 done

Removing enclosure 3...Enclosure 3 successfully removed.
Updating system information...done
Successfully removed: 3 done
```

CHAPTER 19

DD Extended Retention

This chapter includes:

- [DD Extended Retention overview](#) 484
- [Supported protocols in DD Extended Retention](#) 486
- [High Availability and Extended Retention](#) 486
- [Using DD Replicator with DD Extended Retention](#) 486
- [Hardware and licensing for DD Extended Retention](#) 488
- [Managing DD Extended Retention](#) 492
- [Upgrades and recovery with DD Extended Retention](#) 502

DD Extended Retention overview

Data Domain Extended Retention (DD Extended Retention) provides an internal tiering approach that enables cost-effective, long-term retention of backup data on a DD system. DD Extended Retention lets you leverage DD systems for long-term backup retention and minimize reliance on tape.

Note

DD Extended Retention was formerly known as *Data Domain Archiver*.

Two-Tiered File System

The internal two-tiered file system of a DD Extended Retention-enabled DD system consists of an *active tier* and a *retention tier*. The file system, however, appears as a single entity. Incoming data 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 your individual *Data Movement Policy*. For example, the active tier might retain weekly full and daily incremental backups for 90 days, while the retention tier might retain monthly fulls for seven years.

The retention tier is comprised of one or more retention units, each of which may draw storage from one or more shelves.

Note

As of DD OS 5.5.1, only one retention unit per retention tier is allowed. However, systems set up prior to DD OS 5.5.1 may continue to have more than one retention unit, but you will not be allowed to add any more retention units to them.

Transparency of Operation

DD Extended Retention-enabled DD systems support existing backup applications using simultaneous data access methods through NFS and CIFS file service protocols over Ethernet, through DD VTL for open systems and IBMi, or as a disk-based target using application-specific interfaces, such as DD Boost (for use with Avamar®, NetWorker®, GreenPlum, Symantec OpenStorage, and Oracle RMAN).

DD Extended Retention extends the DD architecture with automatic transparent data movement from the active tier to the retention tier. 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 is not affected by data movement. No partitioning of the file system is necessary to take advantage of the two-tiered file system.

Data Movement Policy

The *Data Movement Policy*, which you can customize, is the policy by which files are moved from the active to the retention tier. It is based on the time when the file was last modified. You can set a different policy for each different subset of data, because the policy can be set on a per-MTree basis. Files that may be updated need a policy different from those that never change.

Deduplication within Retention Unit

For fault isolation purposes, deduplication occurs entirely within the retention unit for DD Extended Retention-enabled DD systems. There is no cross-deduplication between active and retention tiers, or between different retention units (if applicable).

Storage Drawn from Each Tier

The concept of tiering extends to the storage level for a DD Extended Retention-enabled DD system. The active tier of the file system draws storage from the active tier of storage. The retention tier of the file system draws storage from the retention tier of storage.

Note

For both active and retention tiers, DD OS 5.2 and later releases support ES20 and ES30 shelves, and DD OS 5.7 and later supports DS60 shelves on certain models.. Different Data Domain shelf types cannot be mixed in the same shelf set, and the shelf sets must be balanced according to the configuration rules specified in the *ES30 Expansion Shelf Hardware Guide* or *DS60 Expansion Shelf Hardware Guide*. With DD Extended Retention, you can attach significantly more storage to the same controller. For example, you can attach up to a maximum of 56 ES30 shelves on a DD990 with DD Extended Retention. The active tier must include storage consisting of at least one shelf. For the minimum and maximum shelf configuration for the Data Domain controller models, refer to the expansion shelf hardware guides for ES30 and DS60.

Data Protection

On a DD Extended Retention-enabled DD system, data is protected with built-in fault isolation features, disaster recovery capability, and DIA (Data Invulnerability Architecture). 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 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 the data in the retention unit may be recovered even when other parts of the system are lost.

Note

Sanitization and some forms of Replication are not supported for DD Extended Retention-enabled DD systems.

Space Reclamation

To reclaim space that has been freed up by data moved to the retention tier, you can use *Space Reclamation* (as of DD OS 5.3), which runs in the background as a low-priority activity. It suspends itself when there are higher priority activities, such as data movement and cleaning.

Encryption of Data at Rest

As of DD OS 5.5.1, you can use the *Encryption of Data at Rest* feature on DD Extended Retention-enabled DD systems, if you have an encryption license. Encryption is not enabled by default.

This is an extension of the encryption capability already available, prior to DD OS 5.5.1, for systems not using DD Extended Retention.

Refer to the *Managing Encryption of Data at Rest* chapter in this guide for complete instructions on setting up and using the encryption feature.

Supported protocols in DD Extended Retention

DD Extended Retention-enabled DD systems support the protocols NFS, CIFS, and DD Boost. Support for DD VTL was added in DD OS 5.2, and support for NDMP was added in DD OS 5.3.

Note

For a list of applications supported with DD Boost, see the *DD Boost Compatibility List* on the Online Support site.

When you are using DD Extended Retention, data first lands in the active tier. Files are moved in their entirety into the retention unit in the retention tier, as specified by your Data Movement Policy. 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 all users, and all file system metadata is present in the active tier.

The trade-off in moving data from the active to the retention tier is larger capacity versus slightly slower access time if the unit to be accessed is not currently ready for access.

High Availability and Extended Retention

Data Domain systems with High Availability (HA) enabled do not support DD Extended Retention. DD OS cannot currently support Extended Retention with HA.

Using DD Replicator with DD Extended Retention

Some forms of replication are supported on DD Extended Retention-enabled DD systems.

Supported replication types depend on the data to be protected:

- To protect data on a system as a *source*, a DD Extended Retention-enabled DD system supports collection replication, MTree replication, and DD Boost managed file replication.
- To protect data from other systems as a *destination*, a DD Extended Retention-enabled DD system also supports directory replication, as well as collection replication, MTree replication, and DD Boost managed file replication.

Note

Delta (low bandwidth optimization) replication is not supported with DD Extended Retention. You must disable delta replication on all contexts before enabling DD Extended Retention on a DD system.

Collection replication with DD Extended Retention

Collection replication takes place between the corresponding active tier and retention unit of the two DD systems with DD Extended Retention enabled. If the active tier or 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 include:

- Both the source and destination systems must be configured as DD systems with DD Extended Retention enabled.
- The file system must not be enabled on the destination until the retention unit has been added to it, and replication has been configured.

Directory replication with DD Extended Retention

For directory replication, a DD Extended Retention-enabled DD system serves as a replication target and supports one-to-one and many-to-one topologies from any supported DD system. However, DD Extended Retention-enabled DD systems do not support bi-directional directory replication and cannot be a *source* of directory replication.

Note

To copy data using directory replication into a DD Extended Retention-enabled DD system, the source must be running DD OS 5.0 or later. Therefore, on systems running DD OS 5.0 or earlier, you must first import data into an intermediate system running DD OS 5.0 or later. For example, replication from a DD OS 4.9 Extended Retention-enabled system could be made into a DD OS 5.2 non-Extended Retention-enabled system. Then, replication could be made from the DD OS 5.2 system into the DD OS 4.9 system.

MTree replication with DD Extended Retention

You can set up MTree replication between two DD Extended Retention-enabled DD systems. Replicated data is first placed in the active tier on the destination system. The Data Movement Policy on the destination system then determines when the replicated data is moved to the retention tier.

Note that MTree replication restrictions and policies vary by DD OS release, as follows:

- As of DD OS 5.1, data can be replicated from a non-DD Extended Retention-enabled system to a DD Extended Retention-enabled system with MTree replication.
- As of DD OS 5.2, data can be protected within an active tier by replicating it to the active tier of a DD Extended Retention-enabled system.
- As of DD OS 5.5, MTree replication is supported from a DD Extended Retention-enabled system to a non-DD Extended Retention-enabled system if both are running DD OS 5.5 or later.
- For DD OS 5.3 and 5.4, if you plan to enable DD Extended Retention, do not set up replication for the /backup MTree on the source machine. (DD OS 5.5 and later do not have this restriction.)

Managed file replication with DD Extended Retention

For DD Extended Retention-enabled DD systems, the supported topologies for DD Boost managed file replication are one-to-one, many-to-one, bi-directional, one-to-many, and cascaded.

Note

For DD Boost 2.3 or later, you can specify how multiple copies are to be made and managed within the backup application.

Hardware and licensing for DD Extended Retention

Certain hardware configurations are required for DD Extended Retention-enabled DD systems. Licensing, specifically separate shelf capacity licenses, is also specific to this feature.

Hardware supported for DD Extended Retention

The hardware requirements for DD Extended Retention-enabled DD systems include memory requirements, shelves, NIC/FC cards, and so on. For details about the required hardware configurations for DD Extended Retention, see the installation and setup guide for your DD system, and the expansion shelf hardware guides for your expansion shelves.

The following DD systems support DD Extended Retention:

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 a DD860, the maximum usable storage capacity of an active tier is 142 TB. The retention tier can have a maximum usable capacity of 142 TB. The active and retention tiers have a total usable storage capacity of 284 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 a DD990, the maximum usable storage capacity of the active tier is 570 TB. The retention tier can have a maximum usable capacity of 570 TB. The active and retention tiers have a total usable storage capacity of 1140 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 a DD4200, the maximum usable storage capacity of the active tier is 192 TB. The retention tier can have a maximum usable capacity of 192 TB. The active and retention tiers have a total usable storage capacity of 384 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 a DD4500, the maximum usable storage capacity of the active tier is 285 TB. The retention tier can have a maximum usable capacity of 285 TB. The active and retention tiers have a total usable storage capacity of 570 TB. External connectivity is supported for DD Extended Retention configurations up to 24 shelves.

DD6800

- 192 GB of RAM
- 1 - NVRAM IO module (8 GB)
- 3 - Quad-port SAS IO modules
- 1 - 1 GbE port on the motherboard
- 0 to 4 - 1/10 GbE NIC cards for external connectivity
- 0 to 4 - Dual-Port FC HBA cards for external connectivity
- 0 to 4 - Combined NIC and FC cards
- Shelf combinations are documented in the installation and setup guide for your DD system, and the expansion shelf hardware guides for your expansion shelves.

If DD Extended Retention is enabled on a DD6800, the maximum usable storage capacity of the active tier is 288 TB. The retention tier can have a maximum usable capacity of 288 TB. The active and retention tiers have a total usable storage capacity of 0.6 PB. External connectivity is supported for DD Extended Retention configurations up to 28 shelves.

DD7200

- 256 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 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 a DD7200, the maximum usable storage capacity of the active tier is 432 TB. The retention tier can have a maximum usable capacity of 432 TB. The active and retention tiers have a total usable storage capacity of 864 TB. External connectivity is supported for DD Extended Retention configurations up to 32 shelves.

DD9300

- 384 GB of RAM
- 1 - NVRAM IO module (8 GB)
- 3 - Quad-port SAS IO modules
- 1 - 1 GbE port on the motherboard
- 0 to 4 - 1/10 GbE NIC cards for external connectivity
- 0 to 4 - Dual-Port FC HBA cards for external connectivity
- 0 to 4 - Combined NIC and FC cards
- Shelf combinations are documented in the installation and setup guide for your DD system, and the expansion shelf hardware guides for your expansion shelves.

If DD Extended Retention is enabled on a DD9300, the maximum usable storage capacity of the active tier is 720 TB. The retention tier can have a maximum usable capacity of 720 TB. The active and retention tiers have a total usable storage capacity of 1.4 PB. External connectivity is supported for DD Extended Retention configurations up to 28 shelves.

DD9500

- 512 GB of RAM
- 1 - NVRAM IO module (8 GB)
- 4 - Quad-port SAS IO modules
- 1 - Quad 1 GbE ports on the motherboard
- 0 to 4 - 10 GbE NIC cards for external connectivity
- 0 to 4 - Dual-Port 16 Gbe FC HBA cards for external connectivity

- Shelf combinations are documented in the installation and setup guide for your DD system, and the expansion shelf hardware guides for your expansion shelves.

If DD Extended Retention is enabled on a DD9500, the maximum usable storage capacity of the active tier is 864 TB. The retention tier can have a maximum usable capacity of 864 TB. The active and retention tiers have a total usable storage capacity of 1.7 PB. External connectivity is supported for DD Extended Retention configurations up to 56 shelves.

DD9800

- 768 GB of RAM
- 1 - NVRAM IO module (8 GB)
- 4 - Quad-port SAS IO modules
- 1 - Quad 1 GbE ports on the motherboard
- 0 to 4 - 10 GbE NIC cards for external connectivity
- 0 to 4 - Dual-Port 16 Gbe FC HBA cards for external connectivity
- Shelf combinations are documented in the installation and setup guide for your DD system, and the expansion shelf hardware guides for your expansion shelves.

If DD Extended Retention is enabled on a DD9800, the maximum usable storage capacity of the active tier is 1008 TB. The retention tier can have a maximum usable capacity of 1008 TB. The active and retention tiers have a total usable storage capacity of 2.0 PB. External connectivity is supported for DD Extended Retention configurations up to 56 shelves.

Licensing for DD Extended Retention

DD Extended Retention is a licensed software option installed on a supported DD system.

A separate shelf capacity license is needed for each storage shelf, for shelves installed in both the active tier and the retention tier. Shelf capacity licenses are 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 Data Domain model. You cannot use the additional storage without first applying the appropriate licenses.

Adding shelf capacity licenses for DD Extended Retention

Every shelf in a DD Extended Retention-enabled DD system must have a separate license.

Procedure

1. Select **Administration > Licenses**.
2. Click **Add Licenses**.
3. Enter one or more licenses, one per line, pressing the Enter key after each one. Click **Add** when you have finished. If there are any errors, a summary of the added licenses, and those not added because of the error, are listed. Select the erroneous License Key to fix it.

Results

The licenses for the DD system are displayed in two groups:

- Software option licenses, which are required for options such as DD Extended Retention and DD Boost.

- Shelf Capacity Licenses, which display shelf capacity (in TiB), the shelf model (such as ES30), and the shelf's storage tier (active or retention).

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 continue.

Configuring storage for DD Extended Retention

Additional storage for DD Extended Retention requires the appropriate license or licenses and enough installed memory on the DD system to support it. Error messages display if more licenses or memory are needed.

Procedure

1. Select **Hardware > Storage** tab.
2. In the Overview tab, select **Configure Storage**.
3. In the Configure Storage tab, select the storage to be added from the Addable Storage list.
4. Select the appropriate Tier Configuration (or **Active** or **Retention**) from the menu. The active tier is analogous to a standard DD 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.
5. Select the checkbox for the Shelf to be added.
6. Click the **Add to Tier** button.
7. Click **OK** to add the storage.
8. To remove an added shelf, select it in the Tier Configuration list, select **Remove from Tier**, and select **OK**.

Customer-provided infrastructure for DD Extended Retention

Before enabling DD Extended Retention, your environment and setup must meet certain requirements.

- **Specifications, site requirements, rack space, and interconnect cabling:** See the *Data Domain Installation and Setup Guide* for your DD system model.
- **Racking and cabling:** It is recommended that you rack your system with future expansion in mind. All shelves are attached to a single DD system.

Note

- See the *Data Domain Expansion Shelf Hardware Guide* for your shelf model (ES20, ES30, or DS60).

Managing DD Extended Retention

To set up and use DD Extended Retention on your DD system, you can use the DD System Manager and/or the DD CLI.

- The DD System Manager, formerly known as the Enterprise Manager, is a graphical user interface (GUI), which is described in this guide.
- The `archive` commands, entered at the DD Command Line Interface (CLI) are described in the *Data Domain Operating System Command Reference Guide*.

The only command not available when you use the DD System Manager is `archive report`.

Enabling DD systems for DD Extended Retention

Before using a DD system for DD Extended Retention, you must have the correct license and the correct file system setup.

Procedure

1. Ensure that the correct license is applied. Select **Administration > Licenses**, and check the Feature Licenses list for Extended Retention.
2. Select **Data Management > File System > More Tasks > Enable DD Extended Retention**.

This option is available only if your Data Domain system supports DD Extended Retention and the file system has not already been configured for DD Extended Retention. Be aware that after DD Extended Retention has been enabled, it cannot be disabled without destroying the file system.

- a. If the file system is already enabled (as a non-DD Extended Retention 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 a DD Extended Retention file system, the file system page is refreshed to include information about both tiers, and there is a new tab labeled **Retention Units**.

CLI Equivalent

You can also verify that the Extended Retention license has been installed at the CLI.

To use the legacy licensing method:

```
# license show
## License Key                               Feature
-----
1      AAAA-BBBB-CCCC-DDDD                     Replication
2      EEEE-FFFF-GGGG-HHHH                     VTL
-----
```

If the license is not present, each unit includes documentation – a quick install card – which shows the licenses that have been purchased. Enter the following command to populate the license key.

```
# license add license-code
```

Then, enable Extended Retention:

```
# archive enable
```

To use electronic licensing:

```
# elicense show
Feature licenses:
## Feature      Count Mode           Expiration Date
-----
1  REPLICATION  1     permanent (int) n/a
2  VTL          1     permanent (int) n/a
-----
```

If the license is not present, update the license file with the new feature license.

```
# elicense update mylicense.lic
New licenses: Storage Migration
```

```

Feature licenses:
## Feature          Count   Mode          Expiration Date
-----
1  REPLICATION      1      permanent (int)  n/a
2  VTL               1      permanent (int)  n/a
3  EXTENDED RETENTION 1      permanent (int)  n/a
-----
** This will replace all existing Data Domain licenses on the system with the above
EMC ELMS licenses.
Do you want to proceed? (yes|no) [yes]: yes
eLicense(s) updated.

```

Then, enable Extended Retention:

```
# archive enable
```

Creating a two-tiered file system for DD Extended Retention

DD Extended Retention has a two-tiered file system for the active and retention tiers. The DD system must have been enabled for DD Extended Retention before enabling this special file system.

Procedure

1. Select **Data Management > File System**.
2. If a file system exists, destroy it.
3. Select **More Tasks > Create file system**.
4. Select a retention-capable file system and click **Next**.
5. Click **Configure** in the File System Create dialog box.
Storage must be configured before the file system is created.
6. Use the Configure Storage dialog box to add and remove available storage from the Active and Retention Tiers, and click **OK** when you have finished.
The storage in the active tier is used to create the active file system tier, and the storage in the retention tier is used to create a retention unit.

Note

As of DD OS 5.5.1, only one retention unit per retention tier is allowed. However, systems set up prior to DD OS 5.5.1 may continue to have more than one retention unit, but you cannot add any more retention units to them.

7. Use the File System Create dialog box to:
 - a. Select the size of the retention unit from the drop-down list.
 - b. Select the **Enable file system after creation** option.
 - c. Click **Next**.

A Summary page shows the size of the active and retention tiers in the new file system.

8. Click **Finish** to create the file system.
The progress of each creation step is shown, and a progress bar monitors overall status.
9. Click **OK** after the file system execution has completed.

CLI Equivalent

To add additional shelves, use this command once for each enclosure:

```
# storage add tier archive enclosure 5
```

Create an archive unit, and add it to the file system. You are asked to specify the number of enclosures in the archive unit:

```
# filesys archive unit add
```

Verify that the archive unit is created and added to the file system:

```
# filesys archive unit list all
```

Check the file system, as seen by the system:

```
# filesys show space
```

File system panel for DD Extended Retention

After you have enabled a DD system for DD Extended Retention, the **Data Management > File System** panel will look slightly different (from a non-DD Extended Retention-enabled system).

- **State** shows that the file system is either enabled or disabled. You can change the state by using the Disable/Enable button immediately to the right.
- **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. There is also a **More Information** button that will display detailed information about starting and ending times, completion percentage, units reclaimed, space freed, etc.
- Selecting **More Tasks > Destroy** lets you delete all data in the file system, including virtual tapes. This can be done only by a system administrator.
- Selecting **More Tasks > Fast Copy** lets you clone files and MTrees of a source directory to a destination directory. Note that for DD Extended Retention-enabled systems, fast copy will not move data between the active and retention tiers.
- Selecting **More Tasks > Expand Capacity** lets you expand the active or retention tier.

Expanding the active or retention tier

When the file system is enabled, you can expand either the active or the retention tier.

To expand the **Active** tier:

Procedure

1. Select **Data Management > File System > More Tasks > Expand Capacity**.
2. In the Expand File System Capacity dialog, select **Active Tier**, then click **Next**.
3. 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 to the Expand File System Capacity dialog. Select **Finish** to complete the active tier expansion.

To expand the **retention** tier:

Procedure

1. Select **Data Management > File System > More Tasks > Expand Capacity**.
2. In the Expand File System Capacity dialog, select **Retention Tier**, then select **Next**.
3. If a retention unit is available, you will see the **Select Retention Unit** dialog. Select the retention unit you want to expand and then **Next**. If a retention unit is not available, you will see the **Create Retention Unit** dialog, and you must create a retention unit before proceeding.

Note

To ensure optimal performance of a DD system with DD Extended Retention enabled, you should always expand the retention tier in at least two-shelf increments. You should also not wait until the retention unit is nearly full before expanding it.

4. Select the size to expand the retention unit, then click **Configure**.
5. After configuration completes, you are returned to the Expand File System Capacity dialog. Click **Finish** to complete the retention tier expansion.

Reclaiming space in the retention tier

You can reclaim space from deleted data in the retention tier by running space reclamation (introduced in DD OS 5.3). Space reclamation also occurs during file system cleaning.

Procedure

1. Select **Data Management > File System**. Just above the tabs, **Space Reclamation Status** shows the amount of space that is reclaimed after deleting data in the retention tier.
2. 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.
3. Click **More Information** for details about the cycle name, start and end times, effective run time, percent completed (if in progress), units reclaimed, space freed on target unit, and total space freed.

Note

When you use the `archive space-reclamation` command, the system runs space-reclamation in the background until it is manually stopped unless you use the one-cycle option. You can also use the `archive space-reclamation schedule set` command to set the starting time for space-reclamation.

CLI Equivalent

To enable space reclamation:


```
# archive space-reclamation start
```

To disable space reclamation:

```
# archive space-reclamation stop
```

To show the status of space reclamation:

```
# archive space-reclamation status-detailed
```

```
Space-reclamation will start when 'archive data-movement'
completes.
```

```
Previous Cycle:
```

```
-----
```

```
Start time           : Feb 21 2014 14:17
End time             : Feb 21 2014 14:49
Effective run time   : 0   days, 00:32.
Percent completed    : 00 % (was stopped by user)
Units reclaimed      : None
Space freed on target unit : None
Total space freed    : None
```

File system tabs for DD Extended Retention

After you have enabled a DD system for DD Extended Retention, the **Data Management > File System** tabs will also look slightly different (from a non-DD Extended Retention-enabled system), and there will be one additional tab: **Retention Units**

Summary Tab

The Summary tab displays information about disk space usage and compression for both the active and retention tiers.

Space Usage: 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.

Active Tier and Retention Tier: Shows the pre-compression and post-compression values currently used and those written in the last 24 hours. Also shows the global, local, and total compression (reduction percentage) factors.

Retention Units Tab

The Retention Units tab displays the retention unit(s). As of DD OS 5.5.1.4, only one retention unit per retention tier is allowed. However, systems set up prior to DD OS 5.5.1.4 may continue to have more than one retention unit, but you will not be allowed to add any more retention units to them.

The following information is displayed: the unit's State (New, Empty, Sealed, Target, or Cleaning), its Status (Disabled, Ready, or Stand-by), its Start Date (when it was moved to the retention tier), and the Unit Size. The unit will be in the cleaning state if space reclamation is running. If the unit has been sealed, meaning no more data can be added, the Sealed Date is provided. Selecting the retention unit's checkbox displays additional information (Size, Used, Available, and Cleanable) in the Detailed Information panel.

There are two buttons: **Delete** (for deleting the unit) and **Expand** (for adding storage to a unit). The unit must be in a new or target state to be expanded.

Configuration Tab

The Configuration Tab lets you configure your system.

Selecting the Options **Edit** button displays the Modify Settings dialog, where you can change Local Compression Type [options are none, lz (the default), gz, and gzfast]

and Retention Tier Local Comp(ression) [options are none, lz, gz (the default), and gzfast], as well as enable Report Replica Writable.

Selecting the Clean Schedule **Edit** button displays the Modify Schedule dialog, where you can change the cleaning schedule, as well as the throttle percentage.

Selecting the Data Movement Policy **Edit** button displays the Data Movement Policy dialog, where you can set several parameters. File Age Threshold is a system-wide default that applies to all MTree for which you have not set a custom default. The minimum value is 14 days. Data Movement Schedule lets you establish how often data movement will be done; the recommended schedule is every two weeks. File System Cleaning lets you elect not to have a system cleaning after data movement; however, it is strongly recommended that you leave this option selected.

File Age Threshold per MTree Link

Selecting the **File Age Threshold per MTree** link will take you from the File System to the MTree area (also accessible by selecting **Data Management > MTree**), where you can set a customized File Age Threshold for each of your MTrees.

Select the MTree, and then select **Edit** next to Data Movement Policy. In the Modify Age Threshold dialog, enter a new value for File Age Threshold, and select **OK**. As of DD OS 5.5.1, the minimum value is 14 days.

Encryption Tab

The Encryption tab lets you enable or disable Encryption of Data at Rest, which is supported only for systems with a single retention unit. As of 5.5.1, DD Extended Retention supports only a single retention unit, so systems set up during, or after, 5.5.1 will have no problem complying with this restriction. However, systems set up prior to 5.5.1 may have more than one retention unit, but they will not work with Encryption of Data at Rest until all but one retention unit has been removed, or data has been moved or migrated to one retention unit.

Space Usage Tab

The Space Usage Tab lets you select one of three chart types [(entire) File System; Active (tier); Archive (tier)] to view space usage over time in MiB. You can also select a duration value (7, 30, 60, or 120 days) at the upper right. The data is presented (color-coded) as pre-compression written (blue), post-compression used (red), and the compression factor (black).

Consumption Tab

The Consumption Tab lets you select one of three chart types [(entire) File System; Active (tier); Archive (tier)] to view the amount of post-compression storage used and the compression ratio over time, which enables you to view consumption trends. You can also select a duration value (7, 30, 60, or 120 days) at the upper right. The Capacity checkbox lets you choose whether to display the post-compression storage against total system capacity.

Daily Written Tab

The Daily Written Tab lets you select a duration (7, 30, 60, or 120 days) to see the amount of data written per day. 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).

Expanding a retention unit

To ensure optimal performance, do not wait until a retention unit is nearly full before expanding it, and do not expand it in 1-shelf increments. Storage cannot be moved

from the active tier to the retention tier after the file system has been created. Only unused enclosures can be added to the retention tier.

Procedure

1. Select **Data Management > File System > Retention Units**.

2. Select the retention unit.

Note that if cleaning is running, a retention unit cannot be expanded.

3. Click **Expand**.

The system displays the current retention tier size, an estimated expansion size, and a total expanded capacity. If additional storage is available you can click the Configure link.

4. Click **Next**.

The system displays a warning telling you that you cannot revert the file system to its original size after this operation.

5. Click **Expand** to expand the file system.

Deleting 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, and then add it as a new retention unit.

Procedure

1. Select **Data Management > File System** and click **Disable** to disable the file system if it is running.
2. Select **Data Management > File System > Retention Units**.
3. Select the retention unit.
4. Click **Delete**.

Modifying retention tier local compression

You can modify the local compression algorithm for subsequent data movement to the retention tier.

Procedure

1. Select **Data Management > File System > Configuration**.
2. Click **Edit** to the right of **Options**.
3. Select one of the compression options from the Retention Tier Local Comp 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).

Understanding the Data Movement Policy

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. The *Data Movement Policy* establishes two things: a *File Age Threshold* and a *Data Movement Schedule*. If data has not changed during the period of days set by the File Age Threshold, it is

moved from the active to the retention tier on the date established by the Data Movement Schedule.

Note

As of DD OS 5.5.1, the File Age Threshold must be a minimum of 14 days.

You can specify different File Age Thresholds for each defined MTree. 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 in DD OS 5.3, it is recommended that you schedule data movement and file system cleaning on a bi-weekly (every 14 days) basis. By default, cleaning is always run after data movement completes. It is highly recommended that you do not change this default.

Avoid these common sizing errors:

- Setting a Data Movement Policy that is overly aggressive; data will be moved too soon.
- Setting a Data Movement Policy that is too conservative: after the active tier fills up, you will not be able to write data to the system.
- Having an undersized active tier and then setting an overly aggressive Data Movement Policy to compensate.

Be aware of the following caveats related to snapshots and file system cleaning:

- Files in snapshots are not cleaned, even after they have been moved to the retention tier. Space cannot be reclaimed until the snapshots have been deleted.
- It is recommended that you set the File Age Threshold for snapshots to the minimum of 14 days.

Here are two examples of how to set up a Data Movement Policy.

- You could segregate data with different degrees of change into two different MTrees and set the File Age Threshold to move data soon after the data stabilizes. Create MTree A for daily incremental backups and MTree B for weekly fulls. Set the File Age Threshold for MTree A so that its data is *never* moved, but set the File Age Threshold for MTree B to 14 days (the minimum threshold).
- For data that cannot be separated into different MTrees, you could do the following. Suppose the retention period of daily incremental backups is eight weeks, and the retention period of weekly fulls is three years. In this case, it would be best to set the File Age Threshold to nine weeks. If it were set lower, you would be moving daily incremental data that was actually soon to be deleted.

Modifying the Data Movement Policy

You can set different Data Movement Policies for each MTree.

Procedure

1. Select **Data Management > File System > Configuration**.
2. Click **Edit** to the right of **Data Movement Policy**.
3. In the Data Movement Policy dialog, specify the system-wide default File Age Threshold value in number of days. As of DD OS 5.5.1, this value must be greater than or equal to 14 days. This value 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 (see step 7). When data movement

starts, all files that have not been modified for the specified threshold number of days will be moved from the active to the retention tier.

4. Specify a Data Movement Schedule, that is, when data movement should take place; for example, daily, weekly, bi-weekly (every 14 days), monthly, or on the last day of the month. You can also pick a specific day or days, and a time in hours and minutes. It is highly recommended that you schedule data movement and file system cleaning on a bi-weekly (every 14 days) basis, to take advantage of the space reclamation feature (introduced in DD OS 5.3).
5. Specify a Data Movement Throttle, that is, the percentage of available resources the system uses for data movement. A value of 100% indicates that data movement will not be throttled.
6. By default, file system cleaning is always run after data movement completes. It is highly recommended that you leave **Start file system clean after Data Movement** selected.
7. Select OK.
8. Back in the Configuration tab, you can specify age threshold values for individual MTrees by using the **File Age Threshold per MTree** link at the lower right corner.

CLI Equivalent

To set the age threshold:

```
# archive data-movement policy set age-threshold {days|none}
mtrees mtree-list
```

If necessary, to set the *default* age threshold:

```
# archive data-movement policy set default-age-threshold days
```

To verify the age threshold setting:

```
# archive data-movement policy show [mtree mtree-list]
```

To specify the migration schedule:

```
# archive data-movement schedule set days days time time [no-clean]
```

Acceptable schedule values include:

- days sun time 00:00
- days mon,tue time 00:00
- days 2 time 10:00
- days 2,15 time 10:00
- days last time 10:00 - last day of the month

To verify the migration schedule:

```
# archive data-movement schedule show
```

To disable the file cleaning schedule:

Note

The reason for disabling the cleaning schedule is to eliminate a scheduling conflict between cleaning and data movement. At the conclusion of data movement, cleaning will automatically start. If you disable data movement, you should re-enable file system cleaning.

```
# filesystems clean set schedule never
```

Starting or stopping data movement on demand

Even when you have a regular Data Movement Policy, you can also start or stop data movement *on demand*.

Procedure

1. Select **Data Management > File System**.
2. Click **Start** to the right of **Data Movement Status**.
3. The Start Data Movement dialog warns that data is to be moved from the active to the retention tier, as defined by your Data Movement Policy, followed by a file system cleaning. Select **Start** to start the data movement.

If a file system cleaning happens to already be in progress, data movement will occur after that cleaning completes. However, another cleaning will be automatically started after this on-demand data movement completes, as well.

4. The **Start** button will be replaced by a **Stop** button.
5. At any time, if you want to stop data movement, click **Stop** and click **OK** in the Stop Data Movement dialog to confirm.

Using data movement packing

Data is compacted in the target partition after every file migration (as of DD OS 5.2). By default, this feature, which is called *data movement packing*, is enabled.

When this feature is enabled, the overall compression of the retention tier improves, but there is a slight increase in migration time.

To determine if this feature is enabled, select **Data Management > File System > Configuration**.

The current value for **Packing data during Retention Tier data movement** can be either Enabled or Disabled. Consult with a system engineer to change this setting.

Upgrades and recovery with DD Extended Retention

The following sections describe how to perform software and hardware upgrades, and how to recover data, for DD Extended Retention-enabled DD systems.

Upgrading to DD OS 5.7 with DD Extended Retention

The upgrade policy for a DD Extended Retention-enabled DD system is the same as for a standard DD system.

Upgrading from up to two major prior releases is supported. For instructions on how to upgrade the DD OS, refer to the upgrade instructions section of the *Release Notes* for the target DD OS version.

When upgrading a DD Extended Retention-enabled DD system to DD OS 5.7, be sure to update existing data movement schedules to bi-weekly (14 days) to take advantage of the space reclamation feature.

DD Extended Retention-enabled DD systems automatically run cleaning after data movement completes; therefore, do not schedule cleaning separately using the DD System Manager or CLI (command line interface).

If the active tier is available, the process upgrades the active tier and the retention unit, and puts the system into a 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 the retention tier has been upgraded. A subsequent upgrade is not permitted until this state is cleared.

If the active tier is not available, the upgrade process upgrades the system chassis and places it into a state where it is ready to create or accept a file system.

If the 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.

Upgrading hardware with DD Extended Retention

You can upgrade a DD Extended Retention-enabled DD system to a later or higher performance DD Extended Retention-enabled DD system. For example, you could replace a DD Extended Retention-enabled DD860 with a DD Extended Retention-enabled DD990

Note

Consult your contracted service provider, and refer to the instructions in the appropriate *System Controller Upgrade Guide*.

This type of upgrade affects DD Extended Retention as follows:

- If the new system has a more recent version of DD OS than the active and retention tiers, the active and retention tiers are upgraded to the new system's version. Otherwise, the new system is upgraded to the version of the active and retention tiers.
- The active and retention tiers that are connected to the new system become owned by the new system.
- If there is an active tier, the registry in the active tier is installed in the new system. Otherwise, the registry in the retention tier with the most recently updated registry is installed in the new system.

Recovering a DD Extended Retention-enabled system

If the active tier, and a subset of the retention units are lost, on a DD Extended Retention-enabled DD system, and there is no replica available, Support may be able to reconstitute any remaining sealed retention units into a new DD system.

A DD Extended Retention-enabled DD system 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, 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.

Using replication recovery

The replication recovery procedure for a DD Extended Retention-enabled DD system depends on the replication type

- Collection replication – The new source must be configured as a DD Extended Retention-enabled DD system with the same number (or more) of retention units

as the destination. The file system must not be enabled on the new source until the retention units have been added, and replication recovery has been initiated.

Note

If you need to recover only a portion of a system, such as one retention unit, from a collection replica, contact Support.

- MTree replication – See the *MTree Replication* section in the *Working with DD Replicator* chapter.
- DD Boost managed file replication – See the *Data Domain Boost for OpenStorage Administration Guide*.

Recovering from system failures

A DD Extended Retention-enabled DD system is equipped with tools to address failures in different parts of the system.

Procedure

1. Restore the connection between the system controller and the 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 DD Extended Retention through Support.

CHAPTER 20

DD Retention Lock

This chapter includes:

- [DD Retention Lock overview](#) 506
- [Supported data access protocols](#) 508
- [Enabling DD Retention Lock on an MTree](#) 509
- [Client-Side Retention Lock file control](#) 512
- [System behavior with DD Retention Lock](#) 517

DD Retention Lock overview

When data is locked on an MTree that is enabled with DD Retention Lock, DD Retention Lock helps ensure that data integrity is maintained. Any data that is locked cannot be overwritten, modified, or deleted for a user-defined retention period of up to 70 years.

There are two DD Retention Lock editions:

- *Data Domain Retention Lock Governance Edition* retains the functionality of Data Domain Retention Lock prior to DD OS 5.2. You can use Data Domain Retention Lock Governance 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.
- *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 Assessments - Summary and Conclusions – EMC Data Domain Retention Lock Compliance Edition* at:

<https://www.emc.com/collateral/analyst-reports/cohasset-dd-retention-lock-assoc-comp-assess-summ-ar.pdf>

(Login is required.)

Compliance with these standards ensures that files locked on a Data Domain system using Data Domain Retention Lock Compliance Edition software cannot be altered or destroyed before the retention period expires. Data Domain Retention Lock Compliance Edition 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 both the DD Retention Lock Governance and Compliance Editions. The differences in use stem from the system behavior for the DD Retention Lock Compliance Edition, since it places strict restrictions to meet compliance requirements. For an overview, see the *EMC Data Domain Retention Lock Software – A Detailed Review* (a white paper) available at:

<https://www.emc.com/collateral/hardware/white-papers/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 Data Domain 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 using DD Replicator 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.

DD Retention Lock Governance Edition is supported for on-premises, cloud-based, and DD3300 DD VE instances. DD Retention Lock Compliance Edition is not supported for on-premises, cloud-based, or DD3300 DD VE instances.

The topics that follow provide additional information on DD Retention Lock.

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.

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. Some archive applications and backup applications can issue this command when appropriately configured. 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*.

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.

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.
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.

Note

Windows clients may need to download utility programs for DD OS compatibility.

3. Optionally, extend file retention times using client-side commands.
4. Optionally, delete files with expired retention periods using client-side commands.

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, SourceOne, Cloud Tiering Appliance, DiskXtender, and so on. Customers using backup applications such as CommVault can also develop custom scripts to use Data Domain Retention Lock.

To check whether an archive application is tested and certified for DD Retention Lock, refer to the *Data Domain Archive Application Compatibility Guide*.

The protocol support of DD Retention Lock 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.
- DD VTL is supported with DD Retention Lock Governance, but not with DD Retention Lock Compliance.

Virtual tapes, here referred to as *tapes*, are represented as files on the 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 *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 DD 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.
- DD Boost is supported with both DD Retention Lock Governance and Compliance. If client-side scripts are used to retention-lock backup files or backup images, and if a backup application (Veritas NetBackup, for example) is also used on the system via DD Boost, be aware that the backup application may 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 Data Domain system.

Data Domain recommends that administrators change their retention period policy to align with the retention lock time. This applies to many of the backup

applications that are integrated with DD Boost, including Veritas NetBackup, Veritas Backup Exec, and NetWorker.

Setting retention lock during data ingest to a DD BOOST file in DSP mode is not allowed, and the client setting the RL receives an error. Retention lock should be set after the data ingest is complete.

Setting retention lock during data ingest to a DD BOOST file in OST mode, or to an NFS file is not allowed and the client writing the data receives error as soon as RL is set. The partial file written before RL is set is committed to disk as a worm file.

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

Add a DD Retention Lock Governance license to a system, and then enable DD Retention Lock Governance on one or more MTrees.

Procedure

1. Add the DD Retention Lock Governance license, if it is not listed under Feature Licenses.
 - a. Select **Administration > Licenses**
 - b. In the Licenses area click **Add Licenses**.
 - c. In the License Key text box, type the license key.

Note

License keys are case-insensitive. Include the hyphens when typing keys.

- d. Click **Add**.
2. Select an MTree for retention locking.
 - a. Select **Data Management > MTree**.
 - b. Select the MTree you want to use for retention locking. You can also create an empty MTree and add files to it later.
3. Click the MTree Summary tab to display information for the selected MTree.
4. Scroll down to Retention Lock area and click **Edit** to the right of Retention Lock.
5. Enable DD Retention Lock Governance on the MTree and change the default minimum and maximum retention lock periods for the MTree, if required.

Perform the following actions in the Modify Retention Lock dialog box:

- a. Select **Enable** to enable DD Retention Lock Governance on the MTree.

- b. To change the minimum or maximum retention period for the MTree, modify the minimum or maximum 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 the settings.

After you close the Modify Retention Lock dialog box, which is updated MTree information appears in the Retention Lock area.

6. 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 that is 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 that is generated for the MTree.

Note

To check retention lock configuration settings for any MTree, select the MTree in the Navigation Panel, then click the Summary tab.

After you finish

Retention-lock files in a retention-lock-enabled MTree.

Enabling DD Retention Lock Compliance on an MTree

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.

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 typing keys.

2. Set up one or more security officer users accounts according to Role-Base Access Control (RBAC) rules.

- a. In the system administrator 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 initialize and reload the system, which results 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 restart 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 or pool MTrees.

5. To change the default minimum and maximum retention lock periods for a compliance-enabled MTree, type 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

The retention *period* is specified in the format [number] [unit]. For example: 1 min, 1 hr, 1 day, 1 mo, or 1 year. 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

Retention lock files reside 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 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 *Data Domain Archive Application Compatibility Guide*.

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 DD Retention Lock on the Data Domain system.

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 Data Domain and should be used per customer environment.

- 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>

Note

The `touch` command for Windows may have a different format than the Linux examples in this chapter.

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.

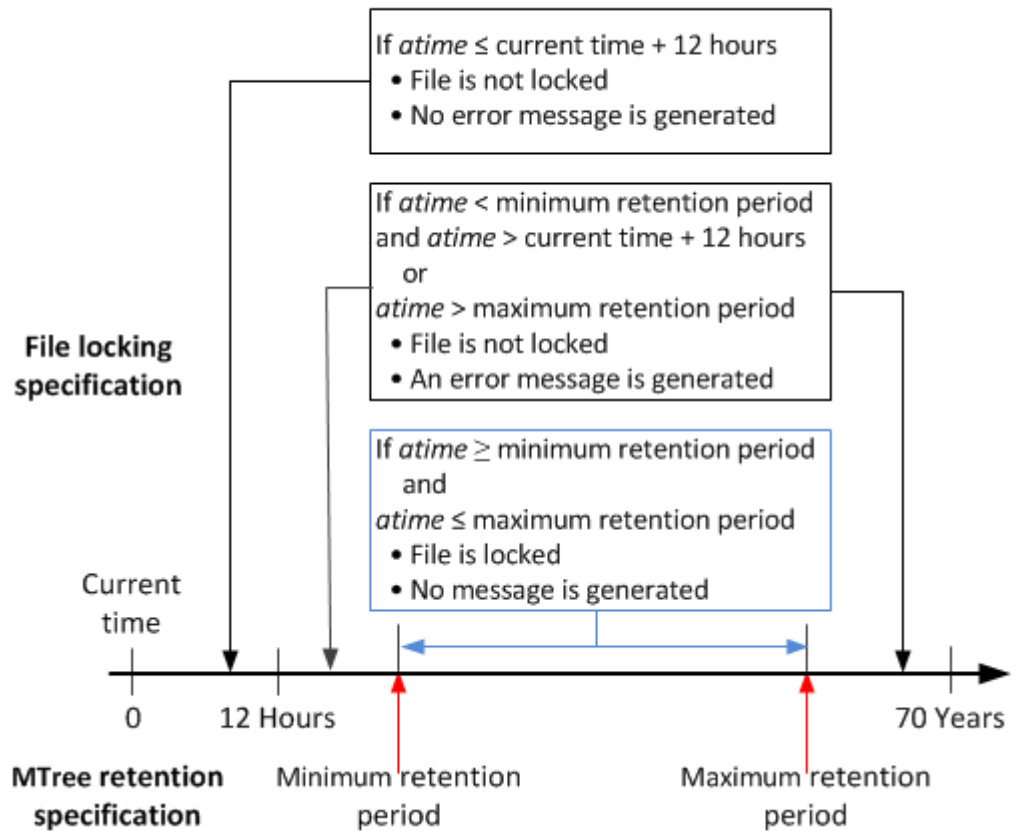
The topics that follow describe how to manage client-side retention lock file control.

Setting Retention Locking on a file

To perform retention locking on a file, change the last access time (*atime*) of the file to the desired retention time of the file, that is, the time when the file can be deleted.

This action is usually performed using the archive application, and all the archive applications that are qualified on Data Domain systems today (per the *Data Domain Archive Application Compatibility Guide*) 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 15 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:

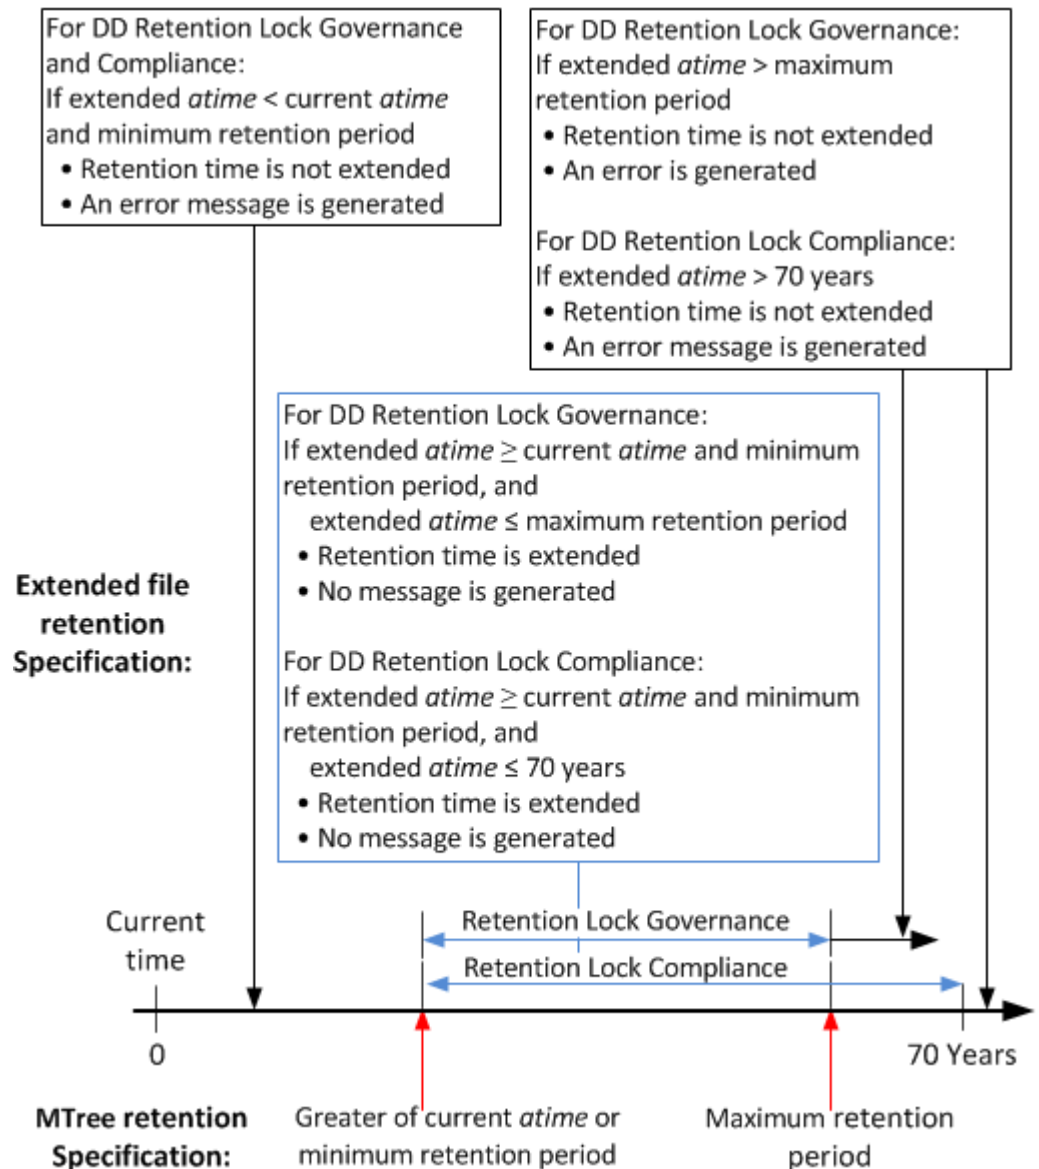
```
ClientOS# touch -a -t 201412312230 SavedData.dat
```

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 16 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

Use the command line to create a root directory containing the files for which access times will change.

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

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. You must have the appropriate access rights to delete the file, independent of DD Retention Lock.

Note

If the retention period of the retention-locked file has not expired, the delete operation results in a permission-denied error.

Privileged delete

For DD Retention Lock Governance (only), you can delete retention locked files using this two step process.

Procedure

1. Use the `mtree retention-lock revert path` command to revert the retention locked file.
2. Delete the file on the client system using the `rm filename` command.

Using ctime or mtime on Retention-Locked files

ctime is the last-metadata-change time of a file.

ctime

ctime 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.

Note

User access permissions for a retention-locked file are updated using the Linux command line tool `chmod`.

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

Collection replication, MTree replication, and directory replication replicate the locked or unlocked state of files.

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, the source system must have a DD Retention Lock Governance license installed—a license is not required on the destination system.

Replication is supported between systems that are:

- Running the same major DD OS version (for example, both systems are running DD OS 5.5.x.x).
- Running DD OS versions within the next two consecutive higher or lower major releases (for example, 5.3.x.x to 5.5.x.x or 5.5.x.x to 5.3.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 files are migrated to the cloud tier before the context is broken, the MTree replication resync overwrites all the data on the destination, so you will need to migrate the files to the cloud tier again.
- If the destination directory has DD Retention Lock enabled, but the source directory does not have DD Retention Lock enabled, then a resync of a directory replication will fail.
- With Mtree replication, resync will fail if the source MTree does not have retention lock enabled and the destination MTree has retention lock enabled.
- With Mtree replication, resync will fail if the source and destination MTrees are retention lock enabled but the propagate retention lock option is set to FALSE.

Fastcopy

When the `filesys fastcopy [retention-lock] source src destination dest` command is run on a system with a DD Retention Lock Governance enabled

MTree, the command preserves the retention lock attribute during the fastcopy operation.

Note

If the destination MTree is not retention lock enabled, the retention-lock file attribute is not preserved.

Filesys destroy

Effects of the `filesystems destroy` command when it is run on a system with a DD Retention Lock Governance enabled MTree.

- All data is destroyed, including retention-locked data.
 - All `filesystems` 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 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

An MTree enabled with DD 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 DD Retention Lock Compliance, or resync will fail.

Replication procedures

The topics in this section describe MTree and collection replication procedures supported for DD Retention Lock Compliance.

Note

For full descriptions of the commands referenced in the following topics, see the *Data Domain Operating System Command Reference Guide*.

Replicating an MTree: One-to-one topology

Replicate a DD Retention Lock Compliance enabled MTree from a source system to a destination system.

Before you begin

Enable DD Retention Lock on an MTree and configure client-side retention lock file control before replication.

Procedure

1. Until instructed otherwise, perform the following steps on the destination system only.
2. 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 typing keys.

3. Set up one or more security officer users accounts according to Role-Based Access Control (RBAC) rules.
 - a. In the system administrator role, add a security officer account.


```
user add user role security
```
 - b. Enable the security officer authorization.


```
authorization policy set security-officer enabled
```
4. 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 initialize and reload the system, which results 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 restart process is complete, enable DD Retention Lock Compliance on the system.

```
system retention-lock compliance enable
```

5. Create a replication context.

```
replication add source mtree://source-system-name/data/coll/mtree-name destination mtree://destination-system-name/data/coll/mtree-name
```

6. Perform the following steps on the source system only.

7. Create a replication context.

```
replication add source mtree://source-system-name/data/coll/mtree-name destination mtree://destination-system-name/data/coll/mtree-name
```

8. Initialize the replication context.

```
replication initialize mtree://destination-system-name/data/coll/mtree-name
```

9. 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

Replicate a DD Retention Lock Compliance enabled MTree from a source system to multiple destination systems.

Before you begin

Enable DD Retention Lock compliance on an MTree and configure client-side retention lock file control before replication.

Procedure

1. Until instructed otherwise, perform the following steps on the destination system only.
2. 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 typing keys.

3. Set up one or more security officer users accounts according to Role-Based Access Control (RBAC) rules.

- a. In the system administrator role, add a security officer account.

```
user add user role security
```

- b. Enable the security officer authorization.

```
authorization policy set security-officer enabled
```

4. 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 initialize and reload the system, which results 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 restart process is complete, enable DD Retention Lock Compliance on the system.

```
system retention-lock compliance enable
```

5. Create a replication context.

```
replication add source mtree://source-system-name/data/coll/mtree-name destination mtree://destination-system-name/data/coll/mtree-name
```

6. Perform the following steps on the source system only.

7. 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
```

8. Initialize the replication context for each destination system MTree.

```
replication initialize mtree://destination-system-name/data/coll/mtree-name
```

9. 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 DD Retention Lock Compliance protection to an existing MTree replication pair

Add DD Retention Lock Compliance protection to an existing MTree replication pair that is not enabled for retention locking.

Procedure

1. Until instructed otherwise, perform the following steps on both the source and destination systems.

2. Log in to the DD System Manager.

The DD System Manager window appears with **DD Network** in the Navigation panel.

3. Select a Data Domain system.

In the Navigation panel, expand **DD Network** and select a system

4. Add the DD Retention Lock Governance license, if it is not listed under Feature Licenses.

- a. Select **Administration > Licenses**

- b. In the Licenses area click **Add Licenses**.

- c. In the License Key text box, type the license key.

Note

License keys are case-insensitive. Include the hyphens when typing keys.

- d. Click **Add**.

5. Break the current MTree context on the replication pair.

```
replication break mtree://destination-system-name/data/
coll/mtree-name
```

6. 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
```

7. Perform the following steps on the source system only.

8. 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.)

9. Click the MTree Summary tab to display information for the selected MTree.

10. Lock files in the compliance-enabled MTree.

11. Ensure that both source and destination (replica) MTrees are the same.

```
replication resync mtree://destination-system-name/data/
coll/mtree-name
```

12. Check the progress of resync.

```
replication watch mtree://destination-system-name/data/
coll/mtree-name
```

13. 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

A procedure 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 DD 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 DD 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 DD 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.

```
filesystems disable
```

b. Break the collection replication context.

```
replication break col://destination-system-name
```

c. Enable the file system. (Security officer authorization may be required.)

```
filesystems enable
```

- d. Add a replication context for each DD 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/
coll/mtree-name detailed
```

This command reports 0 pre-compressed bytes remaining when replication is finished.

Performing collection replication

Replicate /data/coll from a compliance-enabled source system to a compliance-enabled destination system.

Note

For collection replication the same security officer account must be used on both the source and destination systems.

Procedure

1. Until instructed to do differently, perform the following steps on the source system only.
 2. Log in to the DD System Manager.
The DD System Manager window appears with **DD Network** in the Navigation Panel.
 3. Select a Data Domain system.
In the Navigation Panel, expand **DD Network** and select a system.
 4. Add the DD Retention Lock Governance license, if it is not listed under Feature Licenses.
 - a. Select **Administration > Licenses**
 - b. In the Licenses area click **Add Licenses**.
 - c. In the License Key text box, type the license key.
-

Note

License keys are case-insensitive. Include the hyphens when typing keys.

d. Click **Add**.

5. Create the replication context.

```
replication add source col://source-system-name
destination col://destination-system-name
```

6. Until instructed to do differently, perform the following steps on the destination system only.
7. Destroy the file system.

```
filesystem destroy
```

8. Log in to the DD System Manager.

The DD System Manager window appears with **DD Network** in the Navigation Panel.

9. Select a Data Domain system.

In the Navigation Panel, expand **DD Network** and select a system.

10. Create a file system, but do not enable it.

```
filesystem create
```

11. Create the replication context.

```
replication add source col://source-system-name
destination col://destination-system-name
```

12. 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.)

13. Perform the following steps on the source system only.

14. Initialize the replication context.

```
replication initialize source col://source-system-name
destination col://destination-system-name
```

15. 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 DD Retention Lock Compliance protection to an existing collection replication pair

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. Until instructed otherwise, perform the following steps on both the source and destination systems.
2. Disable the replication.

```
replication disable col://destination-system-name
```

3. Log in to the DD System Manager.

The DD System Manager window appears with **DD Network** in the Navigation Panel.

4. Select a Data Domain system.

In the Navigation Panel, expand **DD Network** and select a system.

5. Until instructed otherwise, perform the following steps on the source system.
6. 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.)

7. Enable the replication context.

```
replication enable col://destination-system-name
```

8. Until instructed otherwise, perform the following steps on the destination system.

9. 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.)

10. Enable the replication context.

```
replication enable col://destination-system-name
```

Fastcopy

When the `filesystems fastcopy [retention-lock] source src destination dest` command is run on a system with a DD Retention Lock Compliance enabled MTree, the command preserves the retention lock attribute during the fastcopy operation.

Note

If the destination MTree is not retention lock enabled, the retention-lock file attribute is not preserved.

CLI usage

Considerations for a Data Domain system with DD Retention Lock Compliance.

- Commands that break compliance cannot be run. The following commands are disallowed:
 - `filesystems archive unit del archive-unit`
 - `filesystems 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 file system 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.

```
filesystems enable
```

2. At the prompt, confirm that you want to quit the `filesystems 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.

```
filesystems 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 21

DD Encryption

This chapter includes:

- [DD encryption overview](#)..... 532
- [Configuring encryption](#).....533
- [About key management](#).....533
- [Key manager setup](#)..... 545
- [Changing key managers after setup](#).....550
- [Checking settings for encryption of data at rest](#).....551
- [Enabling and disabling encryption of data at rest](#)..... 552
- [Locking and unlocking the file system](#)..... 553

DD encryption overview

Data encryption protects user data if the Data Domain system is stolen or if the physical storage media is lost during transit, and it eliminates accidental exposure of a failed drive if it is replaced.

When data enters the Data Domain system using any of the supported protocols (NFS, CIFS, DD VTL, DD Boost, and NDMP Tape Server), the stream is segmented, fingerprinted, and de-duplicated (global compression). It is then grouped into multi-segment compression regions, locally compressed, and encrypted before being 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 DD Encryption feature is enabled does not automatically get encrypted. To protect all of the data on the system, be sure to enable the option to encrypt existing data when you configure encryption.

Additional Notes:

As of DD OS 5.5.1.0, Encryption of Data at Rest is supported for DD Extended Retention-enabled systems with a single retention unit. As of 5.5.1.0, DD Extended Retention supports only a single retention unit, so systems set up during, or after, 5.5.1.0 will have no problem complying with this restriction. However, systems set up prior to 5.5.1.0 may have more than one retention unit, but they will not work with Encryption of Data at Rest until all but one retention unit has been removed, or data has been moved or migrated to one retention unit.

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 *Data Domain Operating System Command Reference Guide*.

Encryption of Data at Rest supports all of the currently supported backup applications described in the Backup Compatibility Guides available through Online Support at <http://support.emc.com>.

Data Domain Replicator can be used with encryption, enabling encrypted data to be replicated using collection, directory, MTree, or application-specific managed file replication with the various topologies. Each replication form works uniquely with encryption and offers the same level of security. For more information, see the section on using encryption of data at rest with replication.

Files locked using Data Domain Retention Lock 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

This procedure includes configuring a key manager.

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.

Note

The system passphrase must be set in order to enable encryption.

Provide the following information:

- 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 three)
 - 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, type a rotation interval between 1-12 months.
 - RSA DPM Key Manager
 - SafeNet KeySecure Key Manager

Note

See the section about key management for an explanation about how the Embedded Key Manager, the RSA DPM Key Manager, and SafeNet KeySecure Key Manager work.

The Summary shows the selected configuration values. Review them for correctness. To change a value, click **Back** to browse 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 generated by the user 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) or SafeNet KeySecure Key Manager. Support for Key Management Interoperability Protocol (KMIP) is introduced with DD OS 6.1.

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 or SafeNet KeySecure 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 the Embedded and DPM key managers provide multiple keys, although the system uses only one key at a time to encrypt data coming into a Data Domain system. If the external Key Manager is configured and enabled, the Data Domain systems use keys that are provided by the RSA DPM Key Manager Server. If the same DPM Key Manager manages multiple Data Domain systems, all systems have the same active key (if they are using the same key class) when the systems are synced and the 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 normal basis, depending on the key class. The Embedded Key Manager key rotation is managed on the Data Domain system. The Key Manager key rotation is managed on the external Key Manager server.

KeySecure

KeySecure 8.5 is supported, which is a KMIP compliant key manager product from Safenet Inc/Gemalto Keysecure. To be able to use KMIP key manager, users have to configure both the key manager and the Data Domain system/DD VE, to trust each other. Users have to pre-create keys on the key manager. A Data Domain system will retrieve these keys and their states from KeySecure after establishing a secure TLS connection. See the *Data Domain Operating System and Gemalto KeySecure Integration Guide* for more information on how to create keys and use them on a Data Domain system.

Rectifying lost or corrupted keys

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 a Data Domain system. After exporting, it is recommended that you save the key file in a secure file server 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

All Key Managers support all DD OS file system protocols.

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

If the RSA DPM Key Manager is configured and enabled, the Data Domain systems use keys provided by the RSA DPM Key Manager Server. If the same DPM Key Manager manages multiple Data Domain systems, all systems will have the same active key (if they are using the same key class) when the systems are synced and the file system has been restarted. Key rotation is managed on the RSA DPM Key Manager server.

If the RSA DPM Key Manager is configured and enabled, the Data Domain systems use keys provided by the RSA DPM Key Manager Server. If the same DPM Key Manager manages multiple Data Domain systems, all systems will have the same active key (if they are using the same key class) when the systems are synced and the file system has been restarted. Key rotation is managed on the RSA DPM Key Manager server.

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 cleaning 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 DD System Manager and go to the **Data Management > File System > Encryption** tab. Keys are listed by ID number in the **Encryption Keys** section of the Encryption tab. 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 (see DPM Encryption Key States Supported by Data Domain), and its post-compression size. If the system is licensed for Extended Retention, the following fields are also displayed:

Active Size (post comp)

The amount of physical space on the active tier encrypted with the key.

Retention Size (post comp)

The amount of physical space on the retention tier encrypted with the key.

Click on a Key MUID and the system displays the following information for the key in the Key Details dialog: Tier/Unit (example: Active, Retention-unit-2), creation date,

valid until date, state (see DPM Encryption Key States Supported by Data Domain), and post compression size. Click Close to close the dialog.

Table 191 DPM encryption key states supported by Data Domain

State	Definition
Pending-Activated	The key has just been created. After a file system restart, 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 cleaning 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	<p>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.</p> <hr/> <p>Note</p> <p>A key is not destroyed in the Data Domain system until a cleaning operation is run and completed.</p> <hr/>

Keeping keys 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 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 DD System Manager, 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.

2. Select **Data Management > File System > Encryption**.
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. This procedure requires security officer credentials.

Note

For information about the security officer, see the sections regarding creating local users and enabling security authorization.

To change an RSA DPM key to a state in which it can be destroyed:

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, select **Data Management > File System > Encryption**.
4. In the Encryption Keys section, select the key in the list to be destroyed.
5. Click **Destroy....**

The system displays the Destroy dialog that includes the tier and state for the key.

6. Enter your security officer user name and password.
 7. Confirm that you want to destroy the key by clicking **Destroy**.
-

Note

After a file system clean has run, the key state will change to Destroyed.

Deleting a key

You can delete Key Manager keys that are in the Destroyed or Compromised-Destroyed states. However, you only need to delete a key when the number of keys

has reached the maximum 254 limit. This procedure requires security officer credentials.

Note

To reach the Destroyed state, the Destroying a Key procedure (for either the Embedded Key Manager or the RSA DPM Key Manager) must be performed on the key and a system cleaning must be run.

Procedure

1. Select **Data Management > File System > Encryption**.
2. In the Encryption Keys section, select the key or keys in the list to be deleted.
3. Click **Delete....**
The system displays the key to be deleted, and the tier and state for the key.
4. Type your security officer user name and password.
5. Confirm that you want to delete the key or keys by clicking **Delete**.

Working with the Embedded Key Manager

When the Embedded Key Manager is selected, the Data Domain System creates its own keys.

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 must perform a file system restart to activate the new key and deactivate the old key. You can disable the key rotation policy by clicking the disable button associated with the Embedded Key Manager Key's rotation status.

Creating a key (Embedded Key Manager)

Create an encryption key for the Embedded Key Manager.

Procedure

1. Select **Data Management > File System > DD Encryption**.
2. In the Encryption Keys section, click **Create... .**
3. Type your security officer user name and password.
4. Click **Restart the filesystem now** if you want to restart the file system.
A new Data Domain key will be created. After the file system is restarted, the previous key will become deactivated and the new key will become activated.
5. Click **Create**.

Destroying a key (Embedded Key Manager)

Destroy an encryption key for the Embedded Key Manager.

Procedure

1. Select **Data Management > File System > Encryption**.
2. In the Encryption Keys section, select the key in the list to be destroyed.
3. Click **Destroy....**

The system displays the Destroy dialog that includes the tier and state for the key.

4. Type your security officer user name and password.
 5. Confirm that you want to destroy the key by clicking **Destroy**.
-

Note

After a file system clean has run, the key state changes to Destroyed.

Working with KeySecure Key Manager

KeySecure Key Manager supports external key managers by using Key Management Interoperability Protocol (KMIP) and centrally manages encryption keys in a single, centralized platform.

- Keys will be pre-created on the Key Manager.
- KMIP Key Manager cannot be enabled on systems that have encryption enabled on one or more cloud units.

Using DD System Manager to set up and manage the KeySecure Key Manager

This section describes how to use Data Domain System Manager (DD SM) to manage the KeySecure Key Manager.

Creating a key for the KeySecure Key Manager

Create an encryption key for the KeySecure Key Manager (KMIP).

Procedure

1. Scroll down to the **Key Manager Encryption Keys** table.
2. Click **Add** to create a new Key Manager encryption key.
 - a. Enter the Security Officer username and password.
 - b. Click **Restart the file system now**.
 - c. Click **Create**.
3. Click **Restart the file system now** to make the changes take effect.

A new KIMP key is created. After the file system is restarted, the previous key is deactivated and the new key is activated.

Modifying the state of an existing key in KeySecure Key Manager

Use DD SM to modify the state of an existing KIMP encryption key.

Before you begin

Review the conditions for changing a key state:

- When a key already exists (is active) and a new key is created, the new key will change to the `Pending-Activated` state until the user restarts the file system.
- Users can deactivate a key in an `Activated-RW` state only if there is a `Pending-Activated` key to take its place.
- A key in a `Pending-Activated` state is deactivated only if there is another `Pending-Activated` key to take its place.
- A key in an `Activated-RO` key requires no conditions. Deactivate at any time.

Procedure

1. Select **Data Management > File System > DD Encryption**.
2. Scroll down to view the **Key Manager Encryption Keys** table.
3. Select the appropriate key from the **Key Manager Encryption Keys** table.
4. To deactivate a key:
 - a. Click on any key that shows an `Activated` state.
 - b. Enter the security officer username and password.
 - c. Click **DEACTIVATE**.

Figure 17 Change KMIP key to a Deactivated state

5. Click **Restart the filesystem now**.

Results

The state of an existing key is changed.

Configuring the KeySecure Key Manager

Use DD SM to set the key rotation policy from the Data Domain system.

Before you begin

Confirm the desired Key rotation period (weeks or months), the Key rotation start date, and the Next key rotation date.

Procedure

1. Select **Data Management > File System > DD Encryption**.
2. In the **Key Management** section, click **Configure**. The **Change Key Manager** dialog box opens.
3. Enter your security officer user name and password.
4. Select **KeySecure Key Manager** from the **Key Manager Type** drop down menu. The Change Key Manager information appears.
5. Set the key rotation policy:

Note

The rotation policy is specified in weeks and months. The minimum key rotation policy increment is one week, and the maximum key rotation policy increment is 52 weeks (or 12 months).

- a. Enable the Key Rotation policy. Set the **Enable Key rotation policy** button to enable.
- b. Enter the appropriate dates in the Key rotation schedule field.
- c. Select the appropriate number of weeks or months from the **Weeks or Months** drop down menu.
- d. Click **OK**.
- e. Click **Restart the filesystem now** if you want to restart the file system to make the changes take effect immediately. per Fig 3

Results

The key rotation policy is set or changed.

Using the Data Domain CLI to manage the KeySecure Key Manager

This section describes how to use the CLI to manage the KeySecure Key Manager.

Create a new active key on the KeySecure Key Manager

Use the Data Domain CLI to create a new active key.

Before you begin

Ensure that you have the appropriate user credentials. The security role is required to run these commands.

Procedure

1. Log into the Data Domain system using the security role:

Username: `sec`

Password: `<security officer password>`

2. Create a new active key:

```
# filesys encryption key-manager keys create
```

3. Output that is similar to the following appears:

```
New encryption key was successfully created.  
The filesystem must be restarted to activate the new key.
```

Results

A new active key is created.

Modify the state of an existing key in the KeySecure Key Manager

Use the Data Domain CLI to modify the state of an existing key to a deactivated state.

Before you begin

Ensure that you have the appropriate user credentials. The security role is required to run these commands.

Procedure

1. Log into the Data Domain system using the security role:

Username: `sec`

Password: `<security officer password>`

2. Modify the state of an existing key:

```
# fileSYS encryption key-manager keys modify{<key-id> | muid
<key-muid>}state deactivated
```

For example:

```
# fileSYS encryption key-manager keys modify muid
740D711374A8C964A62817B4AD193C8DC44374A6ED534C85642782014F2E9D
41 state deactivated
```

3. Output that is similar to the following appears:

```
Key state modified.
```

Results

The state of an existing key is modified.

Set or reset a key rotation policy in the KeySecure Key Manager

Use the Data Domain CLI to set the key rotation policy on the Data Domain system to periodically rotate keys. Note that the rotation policy is specified in weeks and months. The minimum key rotation policy increment is one week, and the maximum key rotation policy increment is 52 weeks (or 12 months).

Before you begin

Ensure that you have the appropriate user credentials. The security role is required to run these commands.

Procedure

1. Log into the Data Domain system using the security role:

Username: `sec`

Password: `<security officer password>`

2. Set a key rotation policy for the first time. In our example, we will set the rotation policy to **three weeks**:

```
# fileysys encryption key-manager set key-rotation-policy
{every <n> {weeks | months} | none}
```

For example:

```
# fileysys encryption key-manager set key-rotation-policy
every 3 weeks
```

Output that is similar to the following appears:

```
Key-rotation-policy is set. Encryption key will be rotated
every 3 weeks.
```

3. Subsequently, run this command if you choose to change the existing key rotation policy. In our example, we will change the rotation policy from **three weeks** to **four months**:

Note

Log into the Data Domain system using the security role (where Username is `sec`, and the password is the `<security officer password>`).

```
# fileysys encryption key-manager reset [key-rotation-policy]
```

For example:

```
fileysys encryption key-manager set key-rotation-policy every
4 months
```

Output that is similar to the following appears:

```
Key-rotation-policy is set. Encryption key will be rotated
every 4 months.
```

4. Display the current key rotation policy, or verify that the policy is set correctly:

```
# fileysys encryption key-manager show
```

Output that is similar to the following appears:

```
The current key-manager configuration is:
Key Manager: Enabled
```



```

Server Type:                               KeySecure
Server:                                     <IP address of
KMIP server>
Port:                                       5696
Fips-mode:                                 enabled
Status:                                    Online
Key-class:                                 <key-class>
KMIP-user:                                 <KMIP username>
Key rotation period:                       2 months
Last key rotation date:                   03:14:17 03/19
2018
Next key rotation date:                   01:01:00 05/17
2018

```

Results

The key rotation policy is set or changed.

How the cleaning operation works

Encryption affects the performance of cleaning operations when 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.

Key manager setup

Follow the instructions for the type of key manager you are using.

For more information about SafeNet KeySecure Key Manager setup, see the *Data Domain Operating System and Gemalto KeySecure Integration Guide*.

RSA DPM Key Manager encryption setup

RSA 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).

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 identity 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 the section about RSA DPM key classes.

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 the section about importing certificates for more information.

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

After obtaining certificates, import them to the Data Domain system.

Before you begin

- The Host certificate should be in PKCS12 format.
- The CA certificate should be in PEM format.
- 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.
- 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 the section regarding changing the system passphrase in the Managing the Data Domain Systems chapter.

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.

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.

The following information describes how to respond to a couple of alerts that might appear during certificate management.

- 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:

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 *Data Domain Operating System 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

Configure encryption on the Data Domain System Manager using the DPM Key Manager.

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.

If you are working in a dual stack environment, and the system displays the following error message: "RKM is not configured correctly," then use the `net hosts add ipaddr host-list` command to add the IPv4 address of the Data Domain System to the `/etc/hosts` file.

Note

A DPM server cannot be enabled in an environment using only IPv6 addresses.

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 the section regarding configuring encryption and select the **DPM Key Manager**. If encryption has already been set up, follow the instructions in the section regarding changing key managers after setup.

Setting up KMIP key manager

With KMIP support, a Data Domain appliance can retrieve symmetric key objects that are used for data at rest encryption from KMIP key managers.

Procedure

1. Set up a KeySecure instance with IP address <IP1>.
2. Create and install an SSL server certificate on the KeySecure.
3. Enable KMIP by navigating to **Device > Key Server**.

Ensure <IP1> is the address that is used and Port is <Port1> and the server certificate from Step 2 is used.

4. Create a certificate signing request (CSR) for the system on the Data Domain system/DD VE or Linux computer.
 - a. Log in to the Data Domain.
 - b. Issue the command `adminaccess certificate cert-signing-request generate`.

If the command is successful, it generates the file `CertificateSigningRequest.csr`, which is located in `/ddvar/certificates/`.

By default, NFS exports do not have permissions to access the certificates folder, even to a root user.

```
# mount 16tbddve:/ddvar /mnt/DDVE
# cd /mnt/DDVE/certificates/
bash: cd: /mnt/DDVE/certificates/: Permission denied
# ls -al /mnt/DDVE/
total 800292
drwxr-xr-x 25 root staff      4096 Apr 10 08:32 .
drwxr-xr-x 26 root root       4096 Oct 24 12:11 ..
-rwxr-xr-x  1 root staff       180 Apr 10 08:36 .bashrc
drwxrwsr-x  2 root staff     4096 Aug 18  2016 benchmark
drwxr-sr-x  3 root staff     4096 Apr  4 15:49 cacerts
drwxrwsr-x  2 root staff     4096 Apr  4 12:50 cdes
```

```
drwxrws--- 2 root staff 4096 Apr 11 2017 certificates
drwxrwsr-x 3 root staff 4096 Jul 1 2016 core
```

5. Take this CSR and have it issued/signed by the CA on the KeySecure.

If the command is successful, it generates the file

`CertificateSigningRequest.csr`, which is located in `/ddvar/certificates/`.

6. Download that signed certificate (x.509 pem file] on to the Data Domain system and use the private key of the CSR to create a pkcs#12 file.

Rename `csr` to `pem` in the file name.

7. Download the root CA certificate from the CA of the KeySecure (**Security > Local CAs**).
8. On the Data Domain system/DD VE, use `adminaccess` CLI to install the pkcs#12 client certificate and the CA certificate. Use application type as **keysecure**.
9. On the KeySecure, create a symmetric key with AES-256 as the algorithm and key length.
 - a. Set the owner to the user that will use as KMIP on the Data Domain system/DD VE.
 - b. Select the `Exportable` option.
 - c. Under **Security > Keys > Attributes** for the key, ensure to set **Application Namespace** to `DD_DARE_KEYS`. Ensure to set **Application Data** to key-class that you are planning to use on the Data Domain system/DD VE.
10. Use `filesys encryption key-manager set` command to configure ALL the parameters to access the keysecure key manager.
11. Enable the external key-manager by using the command `filesys encryption key-manager enable`.
12. Enable encryption by using the commands `filesys encryption enable` and `filesys restart`.
This action restarts the file system.
13. Keys should be automatically retrieved from the keysecure key-manager should be seen in the local key table.

Sample output of local key table for `filesys encryption keys` show:

Active Tier:

Key Id	Key MJID	State	Size post-comp
0.1	e56	Deactivated	0
0.2	953C694E2128F977FC8B18D7F8A51E44F8847A8D171D0BBDC8C01576FF5DE1D5	Activated-RW	0

* Post-comp size is based on last cleaning of Tue Feb 14 10:02:02 2017.

The current active key is used to encrypt any data being ingested.

14. Sync the key states.
 - a. On the keysecure web interface, create a new active key as previously described.
 - b. On the keysecure web interface, deactivate the old key by clicking the key and going under the **Life Cycle** tab. Click **Edit State**. Set the **Cryptographic State** to **Deactivated**. Click **Save**.

- On the Data Domain system, sync the local key table by running the `filesys encryption keys sync` command.

Sample output of local key table for `filesys encryption keys` show:

Active Tier:

Key Id	Key MUID	State	Size post-comp
0.1	e56	Deactivated	0
0.2	953C694E2128F977FC8B18D7F8A51E44F8847A8D171DOBBD8C01576FF5DE1D5	Deactivated	0
0.3	851631E574D6F02886CAEF2795896D4C401EBC57A0997EFE04A146E584E9A99A	Activated-RW	0

* Post-comp size is based on last cleaning of Tue Feb 14 10:12:05 2017.

Note

Keys can be marked as versioned keys. When 2nd and 3rd versions of a specific key are generated, KMIP queries currently don't pick up these keys and may be an issue if that key is being used by a Data Domain system or DD VE.

Changing key managers after setup

Select from the Embedded Key Manager or the RSA DPM Key Manager.

Before you begin

To manage certificates for a system, you must start DD System Manager on that system.

Procedure

- Select **Data Management > File System > Encryption**.
- Under Key Management, click **Configure**.
- Type your security officer username and password.
- 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**.
- Click **Manage Certificates** to add certificates.

Managing certificates for RSA Key Manager

You must use both host and CA certificates with RSA Key Manager.

Note

Certificates are only necessary for RSA Key Manager. Embedded Key Manager does not use certificates.

Adding CA certificates for RSA Key Manager

Upload or copy and paste CA certificates.

Procedure

1. Select one of the following:
 - Select the option to upload a CA certificate as a .pem file and click **Browse** to find the file.
 - Select the option to copy and paste the CA certificate and paste the certificate contents into the field provided.
2. Click **Add** to add the certificate.

Adding a host certificate for RSA Key Manager

Upload the certificate as a .p12 file, or upload a public key as a .pem file and use a generated private key.

To start, select the first or second of the following steps:

Procedure

1. Select the option to upload the certificate as a .p12 file.
 - a. Enter a password.
 - b. Click **Browse** to find the .p12 file.
2. Select the option to upload the public key as a .pem file and use a generated private key.
 - a. Click **Browse** to find the .pem file.
3. Click **Add**.

Deleting certificates

Select a certificate with the correct fingerprint.

Procedure

1. Select a certificate to delete.
2. Click **Delete**.

The system displays a Delete Certificate dialog with the fingerprint of the certificate to be deleted.

3. Click **OK**.

Checking settings for encryption of data at rest

Check the settings for the DD Encryption feature.

Click the **Data Management > File System > Encryption** tabs. The currently used Key Manager is shown as Enabled. For a description of the DD Encryption settings, see the section about the encryption view.

Enabling and disabling encryption of data at rest

After configuring DD Encryption, the status is enabled and the Disabled button is active. When DD Encryption is disabled, the Enabled button is active.

Enabling encryption of data at rest

Use the DD System Manager to enable the DD Encryption feature.

Procedure

1. Using the DD System Manager, select the Data Domain system you are working with in the Navigation panel.
2. In the Encryption view, click the **Enable** button.
3. Both of the following options are available:
 - Select **Apply to existing data** and click **OK**. Encryption of existing data will occur during the first cleaning cycle after the file system is restarted.
 - Select **Restart the file system now** and click **OK**. DD Encryption will be enabled after the file system is restarted.

After you finish

Note

Applications may experience an interruption while the file system is restarted.

Disabling encryption of data at rest

Use the DD System Manager to disable the DD Encryption feature.

Procedure

1. Using the DD System Manager, select the Data Domain system you are working with in the Navigation panel.
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 one of the following:
 - Select **Apply to existing data** and click **OK**. Decryption of existing data will occur during the first cleaning cycle after the file system is restarted.
 - Select **Restart the file system now** and click **OK**. DD Encryption will be disabled after the file system is restarted.

After you finish

Note

Applications may experience an interruption while the file system is restarted.

Locking and unlocking the file system

Use this procedure when an DD 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. Select **Data Management > File System > Encryption** .

In the File System Lock area, the Status shows whether the file system is Locked or Unlocked.

2. Disable the file system by clicking **Disabled** in the File System status area.
3. Use the procedure to lock or unlock the file system.

Locking the file system

To lock the file system, DD Encryption must be enabled and the file system must be disabled.

Procedure

1. Select **Data Management > File System > Encryption** and click **Lock File System**.
2. In the text fields of the Lock File System dialog box, provide:
 - The username 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.
3. Click **OK**.

This procedure re-encrypts the encryption keys with the new passphrase. This process destroys the cached copy of the current passphrase (both in-memory and on-disk).

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.

4. Shut down the system:

CAUTION

Do not use the chassis power switch to power off the system. Type the following command at the command prompt instead.

```
# system poweroff The 'system poweroff' command shuts down
the system and turns off the power. Continue? (yes|no|?)
[no]:
```

5. Transport the system or remove the disk being replaced.
6. Power on the system and use the procedure to unlock the file system.

Unlocking the file system

This procedure prepares an encrypted file system for use after it has arrived at its destination.

Procedure

1. Select **Data Management > File System > Encryption** and click **Unlock File System**.
2. In the text fields, type the passphrase that was used to lock the file system.
3. Click **OK**.
4. Click **Close** to exit.

If the passphrase is incorrect, the file system does not start and the system reports the error. Type the correct passphrase, as directed in the previous step.

Changing the encryption algorithm

Reset the encryption algorithm if necessary, or select options to encrypt new and existing data or just new data.

Procedure

1. Select **Data Management > File System > Encryption**
2. To change the Encryption Algorithm used to encrypt the Data Domain system, click **Change Algorithm**.

The Change Algorithm dialog box is displayed. Supported encryption algorithms are:

- AES-128 CBC
- AES-256 CBC
- AES-128 GCM
- AES-256 GCM

3. 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**.

4. 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**.
5. The status is displayed. Click **Close** when the process is complete.
-

Note

Applications may experience an interruption while the file system is restarted.
