

Dell EMC Data Domain[®] Operating System

Version 6.1

Command Reference Guide

302-003-762

REV. 05

Copyright © 2010-2019 Dell Inc. or 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

Figures		15
Preface		17
	Revision history	21
Chapter 1	adminaccess	23
	adminaccess change history.....	24
	adminaccess guidelines and restrictions.....	24
	adminaccess add.....	25
	adminaccess authentication.....	25
	adminaccess certificate.....	26
	adminaccess del.....	34
	adminaccess disable.....	34
	adminaccess enable.....	35
	adminaccess ftp.....	35
	adminaccess ftps.....	35
	adminaccess http.....	36
	adminaccess option.....	36
	adminaccess reset.....	38
	adminaccess show.....	38
	adminaccess ssh.....	39
	adminaccess telnet.....	41
	adminaccess trust.....	41
	adminaccess web.....	42
Chapter 2	alerts	43
	alerts change history.....	44
	alerts clear.....	44
	alerts notify-list.....	44
	alerts show.....	47
	alerts test events.....	49
Chapter 3	alias	51
	alias change history.....	52
	alias add.....	52
	alias del.....	53
	alias reset.....	53
	alias show.....	53
Chapter 4	archive	55
	archive change history.....	56
	archive data-movement.....	56
	archive disable.....	57

	archive enable.....	58
	archive option.....	58
	archive report.....	58
	archive show.....	59
	archive space-reclamation.....	59
Chapter 5	authentication	61
	authentication change history.....	62
	authentication kerberos.....	63
	authentication ldap.....	64
	authentication nis.....	67
Chapter 6	authorization	69
	authorization change history.....	70
	authorization guidelines and restrictions.....	70
	authorization policy.....	70
	authorization show.....	70
Chapter 7	autosupport	71
	autosupport change history.....	72
	autosupport guidelines and restrictions.....	72
	autosupport add.....	72
	autosupport del.....	72
	autosupport reset.....	73
	autosupport send.....	73
	autosupport set.....	74
	autosupport show.....	75
	autosupport test.....	75
Chapter 8	boostfs	77
	boostfs change history.....	78
	boostfs mount.....	78
	boostfs lockbox.....	79
	boostfs kerberos.....	79
Chapter 9	cifs	81
	cifs change history.....	82
	cifs disable.....	82
	cifs enable.....	82
	cifs local-group.....	82
	cifs option.....	83
	cifs reset.....	83
	cifs restart.....	84
	cifs set.....	84
	cifs share.....	85
	cifs show.....	86
	cifs status.....	86
	cifs troubleshooting.....	86
Chapter 10	client-group	89
	client-group change history.....	90

	client-group guidelines and restrictions.....	90
	client-group add.....	91
	client-group compression show.....	91
	client-group create.....	92
	client-group data-access-permit-list.....	92
	client-group del.....	93
	client-group destroy.....	93
	client-group rename.....	93
	client-group show.....	93
	client-group stats options.....	102
	client-group stream-limit.....	104
Chapter 11	cloud	105
	cloud change history.....	106
	cloud clean.....	106
	cloud enable.....	106
	cloud profile.....	106
	cloud provider.....	107
	cloud status.....	108
	cloud unit.....	108
Chapter 12	compression	109
	compression change history.....	110
	compression physical-capacity-measurement.....	110
Chapter 13	config	117
	config change history.....	118
	config reset.....	118
	config set.....	118
	config setup.....	119
	config show.....	119
Chapter 14	data-movement	121
	data-movement change history.....	122
	data-movement policy.....	122
	data-movement recall.....	122
	data-movement resume.....	122
	data-movement schedule.....	123
	data-movement start.....	123
	data-movement status.....	124
	data-movement stop.....	124
	data-movement suspend.....	124
	data-movement throttle.....	124
	data-movement watch.....	125
Chapter 15	ddboost	127
	ddboost change history.....	129
	ddboost guidelines and restrictions.....	129
	ddboost association.....	129
	ddboost clients.....	131
	ddboost destroy.....	133
	ddboost disable.....	133

	ddboost enable.....	133
	ddboost event.....	133
	ddboost fc.....	135
	ddboost file-replication.....	136
	ddboost ifgroup.....	140
	ddboost option.....	143
	ddboost reset.....	144
	ddboost set.....	144
	ddboost show.....	144
	ddboost status.....	147
	ddboost storage-unit.....	148
	ddboost streams.....	151
	ddboost user.....	152
Chapter 16	disk	155
	disk change history.....	156
	disk beacon.....	156
	disk fail.....	156
	disk multipath.....	157
	disk port.....	157
	disk release.....	158
	disk rescan.....	158
	disk reset.....	158
	disk set.....	159
	disk show.....	159
	disk status.....	164
	disk unfail.....	165
Chapter 17	elicense	167
	elicense change history.....	168
	elicense reset.....	168
	elicense show.....	168
	elicense update.....	169
Chapter 18	enclosure	171
	enclosure change history.....	172
	enclosure guidelines and restrictions.....	172
	enclosure beacon.....	172
	enclosure release.....	172
	enclosure show.....	172
	enclosure test.....	177
Chapter 19	filesystems	179
	filesystems change history.....	180
	filesystems archive.....	180
	filesystems clean.....	181
	filesystems create.....	183
	filesystems destroy.....	183
	filesystems disable.....	183
	filesystems enable.....	183
	filesystems encryption.....	184
	filesystems expand.....	193
	filesystems fastcopy.....	193

	fileSYS option.....	194
	fileSYS report.....	197
	fileSYS restart.....	198
	fileSYS show.....	198
	fileSYS status.....	202
	fileSYS sync.....	202
Chapter 20	ha	203
	ha change history.....	204
	ha guidelines and restrictions.....	204
	ha create.....	205
	ha destroy.....	206
	ha failover.....	206
	ha offline.....	206
	ha online.....	206
	ha status.....	207
Chapter 21	help	209
Chapter 22	ifgroup	211
	ifgroup change history.....	212
	ifgroup add.....	212
	ifgroup create.....	213
	ifgroup del.....	213
	ifgroup destroy.....	213
	ifgroup disable.....	213
	ifgroup enable.....	213
	ifgroup option.....	214
	ifgroup rename.....	214
	ifgroup replication assign.....	215
	ifgroup replication unassign.....	215
	ifgroup reset.....	215
	ifgroup show config.....	216
	ifgroup show connections.....	216
Chapter 23	ipmi	219
	ipmi change history.....	220
	ipmi guidelines and restrictions.....	220
	ipmi config.....	220
	ipmi disable.....	220
	ipmi enable.....	220
	ipmi remote.....	220
	ipmi reset.....	221
	ipmi show.....	221
	ipmi user.....	221
Chapter 24	license	223
	license change history.....	224
	license guidelines and restrictions.....	224
	license add.....	224
	license delete.....	224
	license reset.....	224

	license show.....	225
Chapter 25	log	227
	log change history.....	228
	log host.....	228
	log list.....	228
	log view.....	228
	log watch.....	230
Chapter 26	migration	231
	migration change history.....	232
	migration abort.....	232
	migration commit.....	232
	migration receive.....	233
	migration send.....	234
	migration show stats.....	236
	migration status.....	236
	migration watch.....	236
Chapter 27	mtree	239
	mtree change history.....	240
	mtree create.....	240
	mtree delete.....	241
	mtree list.....	241
	mtree modify.....	242
	mtree option.....	243
	mtree rename.....	243
	mtree retention-lock.....	244
	mtree show.....	246
	mtree undelete.....	248
Chapter 28	ndmpd	251
	ndmpd change history.....	252
	ndmpd disable.....	252
	ndmpd enable.....	252
	ndmpd option.....	252
	ndmpd show.....	252
	ndmpd status.....	253
	ndmpd stop.....	253
	ndmpd user.....	253
Chapter 29	net	255
	net change history.....	256
	net guidelines and restrictions.....	256
	net aggregate.....	256
	net config.....	259
	net congestion-check.....	263
	net create.....	267
	net ddns.....	267
	net destroy.....	268
	net disable.....	269
	net enable.....	269

	net failover.....	269
	net filter.....	271
	net hosts.....	274
	net iperf	275
	net lookup.....	278
	net modify.....	278
	net option.....	278
	net ping.....	278
	net reset.....	279
	net route.....	279
	net set.....	285
	net show.....	286
	net tcpdump.....	292
	net troubleshooting.....	292
Chapter 30	nfs	293
	nfs change history.....	294
	nfs add.....	294
	nfs del.....	297
	nfs disable.....	297
	nfs enable.....	297
	nfs export add.....	298
	nfs export create.....	300
	nfs export del	301
	nfs export destroy.....	302
	nfs export modify.....	302
	nfs export rename.....	303
	nfs export show.....	303
	nfs option.....	305
	nfs reset.....	306
	nfs show	307
	nfs status.....	308
Chapter 31	ntp	309
	ntp change history.....	310
	ntp guidelines and restrictions.....	310
	ntp add.....	310
	ntp del.....	310
	ntp disable.....	310
	ntp enable.....	311
	ntp reset.....	311
	ntp show.....	311
	ntp status.....	311
Chapter 32	qos	313
	qos change history.....	314
	qos randomio.....	314
Chapter 33	quota	315
	quota change history.....	316
	quota capacity.....	316
	quota disable.....	317
	quota enable.....	317

	quota reset.....	317
	quota set.....	318
	quota show.....	318
	quota status.....	318
	quota streams.....	318
Chapter 34	replication	321
	replication change history.....	322
	replication abort.....	322
	replication add.....	322
	replication break.....	324
	replication dir-to-mtree.....	324
	replication disable.....	325
	replication enable.....	325
	replication initialize.....	325
	replication modify.....	326
	replication option.....	327
	replication reauth.....	328
	replication recover.....	328
	replication resync	328
	replication show.....	329
	replication status.....	333
	replication sync.....	334
	replication throttle.....	334
	replication watch.....	336
Chapter 35	route	339
	route change history.....	340
	route guidelines and restrictions.....	340
	route add.....	340
	route del.....	342
	route reset.....	342
	route set.....	343
	route show.....	343
	route trace.....	343
Chapter 36	scsitarget	345
	scsitarget change history.....	346
	scsitarget device.....	346
	scsitarget disable.....	346
	scsitarget enable.....	346
	scsitarget endpoint.....	346
	scsitarget group.....	350
	scsitarget initiator.....	352
	scsitarget option.....	353
	scsitarget persistent-reservation.....	355
	scsitarget port.....	355
	scsitarget reset.....	359
	scsitarget service.....	359
	scsitarget show.....	359
	scsitarget status.....	359
	scsitarget trace.....	360
	scsitarget transport.....	361

Chapter 37	smt	365
	smt change history.....	366
	smt disable.....	366
	smt enable.....	366
	smt status.....	366
	smt tenant.....	366
	smt tenant-unit.....	368
Chapter 38	snapshot	375
	snapshot change history.....	376
	snapshot create.....	376
	snapshot expire.....	376
	snapshot list.....	377
	snapshot rename.....	377
	snapshot schedule.....	378
Chapter 39	snmp	381
	snmp change history.....	382
	snmp guidelines and restrictions.....	382
	snmp add.....	382
	snmp_debug.....	383
	snmp del.....	383
	snmp disable.....	384
	snmp enable.....	384
	snmp reset.....	384
	snmp set.....	385
	snmp show.....	385
	snmp status.....	386
	snmp user.....	386
Chapter 40	storage	387
	storage change history.....	388
	storage guidelines and restrictions.....	388
	storage add.....	388
	storage migration.....	389
	storage remove.....	393
	storage sanitize.....	394
	storage show.....	394
Chapter 41	support	401
	support change history.....	402
	support bundle.....	403
	support connectemc.....	403
	support coredump.....	404
	support notification.....	406
Chapter 42	system	407
	system change history.....	408
	system availability.....	408
	system bash.....	408
	system headswap.....	408
	system option.....	408

	system package.....	409
	system passphrase.....	410
	system poweroff.....	412
	system reboot.....	413
	system retention-lock.....	413
	system sanitize.....	413
	system set.....	414
	system show.....	414
	system status.....	423
	system upgrade.....	424
Chapter 43	user	427
	user change history.....	428
	user add.....	428
	user change.....	429
	user del.....	429
	user disable.....	430
	user enable.....	430
	user password.....	430
	user reset.....	435
	user show.....	435
Chapter 44	vdisk	439
	vdisk change history.....	440
	vdisk guidelines and restrictions.....	440
	vdisk config	441
	vdisk device.....	442
	vdisk device-group.....	443
	vdisk disable.....	444
	vdisk enable.....	444
	vdisk group.....	445
	vdisk pool	446
	vdisk property.....	447
	vdisk reset.....	448
	vdisk show.....	448
	vdisk static-image.....	450
	vdisk status.....	452
	vdisk trace.....	452
	vdisk user.....	452
Chapter 45	vtl	455
	vtl change history.....	457
	vtl add.....	457
	vtl cap.....	457
	vtl config.....	458
	vtl debug.....	460
	vtl del.....	462
	vtl disable.....	462
	vtl drive.....	462
	vtl enable.....	464
	vtl export.....	464
	vtl group.....	464
	vtl import.....	466
	vtl option.....	468

	vtl pool.....	469
	vtl readahead.....	471
	vtl rename.....	471
	vtl reset.....	471
	vtl show.....	472
	vtl slot.....	473
	vtl status.....	473
	vtl tape.....	473
Appendix A	Time Zones	477
	Time zones overview.....	478
	Africa.....	478
	America.....	479
	Antarctica.....	480
	Asia.....	480
	Atlantic.....	481
	Australia.....	481
	Brazil.....	481
	Canada.....	482
	Chile.....	482
	Etc.....	482
	Europe.....	482
	GMT.....	483
	Indian (Indian Ocean).....	483
	Mexico.....	483
	Miscellaneous.....	483
	Pacific.....	484
	US (United States).....	484
	Aliases.....	484

CONTENTS

FIGURES

1	Output: net show hardware.....	288
2	Output: net show settings.....	289
3	Output: net show settings for HA (active).....	290
4	Output: net show settings for HA (standby).....	290
5	Output: storage add enclosure 2 tier cache.....	389
6	Output: storage migration finalize.....	390
7	Output: storage migration show history.....	391
8	Output: storage show all.....	396
9	Output: disk show hardware.....	397
10	Output: user show list.....	436

FIGURES

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 describes the Data Domain operating system (DD OS) commands and provides an overview of how they are used. For more specific, task-based instructions, see the *Data Domain Operating System Administration Guide*.

Related documentation

The following Data Domain system documents provide additional information.

- Installation and setup guide for your system, for example, *Data Domain DD 2500 Storage System, Installation and Setup 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 Product Security Guide*
- *Data Domain Operating System Administration Guide*
- *Data Domain Operating System MIB Quick Reference*
- *Data Domain Operating System Offline Diagnostics Suite User's Guide*
- Hardware overview guide for your system, for example, *Data Domain DD4200, DD4500, and DD7200 Systems, Hardware Overview*
- Field replacement guides for your system components, for example, *Field Replacement Guide, Data Domain DD4200, DD4500, and DD7200 Systems, I/O Module and Management Module Replacement or Upgrade*
- *Data Domain System Controller Upgrade Guide*
- *Data Domain Expansion Shelf Hardware Guide* (for shelf model ES20 or ES30)
- *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*

Special notice conventions used in this document

Data Domain uses the following conventions for special notices.

NOTICE

A notice identifies content that warns of 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

The following table describes the type style conventions used in this document.

Table 1 Typographical Conventions

Bold	Use for names of interface elements, 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>	Use for full titles of publications referenced in text
Monospace	Use for: <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>	Use for variables
Monospace bold	Use 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.

Revision history

Table 2 Document Revision History

Revision	Date	Description
05 (6.1.2)	February 2019	This revision includes updated information about these commands: <ul style="list-style-type: none">• <code>cifs option reset server-signing</code>• <code>cifs option set server-signing</code>
04 (6.1.2)	July 2018	This revision includes information about these command changes supporting Data Domain 6.1.2: <ul style="list-style-type: none">• <code>adminaccess</code> commands: Four new, four with modified arguments.• <code>authentication</code> commands: Sixteen new.• <code>cifs</code> commands: Three with modified arguments and modified outputs.• <code>client-group</code> commands: Four new, one with modified arguments.• <code>config</code> commands: One with modified arguments, one with modified output.• <code>elicense</code> commands: Added new options in support of the subscription license, also known as the Flexible Consumption Model (FCM) in DD VE.• <code>filesystems encryption key-manager</code> commands for KMIP support: Four new.• <code>support</code> commands: One new, five with modified arguments, six removed.• <code>mtree</code> commands: One with modified behavior.• <code>nfs</code> commands: One with modified arguments and modified behavior.• <code>vtl</code> commands: Two removed.
03 (6.1.1)	February 2018	This revision includes updated information about these commands: <ul style="list-style-type: none">• <code>ha create</code>• <code>ha destroy</code>
02 (6.1.1)	January 2018	This revision includes information about these command changes supporting Data Domain 6.1.1: <ul style="list-style-type: none">• <code>cloud</code> commands: One new, one with modified arguments.

Table 2 Document Revision History (continued)

Revision	Date	Description
		<ul style="list-style-type: none"> • data-movement commands: One with modified arguments and modified behavior, one with modified output. • snmp commands: Two new. • system commands: One new.
01 (6.1)	June 2017	<p>This revision includes information about these command changes supporting Data Domain 6.1:</p> <ul style="list-style-type: none"> • adminaccess commands: Three new. • client-group commands: Three new. • cloud commands: Three new. • data-movement commands: One with modified arguments, one with modified behavior. • ddbost commands: Six with modified arguments, two with modified output. • filesystems commands: One with modified arguments. • nfs commands: Fifteen new. • vtl commands: Four new, one with modified arguments, one with modified output.

CHAPTER 1

adminaccess

The `adminaccess` command manages access control and enables users to import host and CA certificates. Command options also enable remote hosts to use the FTP, FTPS, Telnet, HTTP, HTTPS, SSH, and SCP administrative protocols on the Data Domain system. SSH is open to the default user `sysadmin` and to users added by the administrator.

A Certificate Signing Request (CSR) can now be generated for a host certificate. Also, host certificates can now be imported in PKCS12 or PEM formats. Data Domain uses the SHA1 RSA encryption algorithm for a CSR and PBE-SHA1-3DES encryption algorithm for the PKCS12 key and certificate.

This chapter contains the following topics:

• adminaccess change history	24
• adminaccess guidelines and restrictions	24
• adminaccess add	25
• adminaccess authentication	25
• adminaccess certificate	26
• adminaccess del	34
• adminaccess disable	34
• adminaccess enable	35
• adminaccess ftp	35
• adminaccess ftps	35
• adminaccess http	36
• adminaccess option	36
• adminaccess reset	38
• adminaccess show	38
• adminaccess ssh	39
• adminaccess telnet	41
• adminaccess trust	41
• adminaccess web	42

adminaccess change history

New commands in DD OS 6.1.2

`adminaccess option reset password-hash`

Resets the password hash to its default value of MD5. Role required: admin.

`adminaccess option set password-hash {md5 | sha512}`

Specifies the hashing algorithm used to store local user passwords. After changing the hashing algorithm, update the passwords for existing local users to be compatible with the new algorithm. Role required: admin.

`adminaccess ssh option set ciphers cipher-list`

Sets the ciphers to be used by SSH daemon.

`adminaccess ssh option set macs MAC-list`

Sets the MACs to be used by SSH daemon.

Modified arguments in DD OS 6.1.2

`adminaccess ssh option reset [ciphers | macs | server-port | session-timeout]`

The `ciphers` and `macs` parameters have been added.

`adminaccess certificate delete { subject subject-name | fingerprint fingerprint} [application {all | cloud | ddbboost | ldap | login-auth | https | keysecure | rkm | support | application-list}]`

The `ldap` parameter has been added.

`adminaccess certificate import {host application {all | ddbboost | https | keysecure | rkm | application-list} | ca application {all | cloud | ddbboost | ldap | login-auth | keysecure | rkm | application-list}} [file file-name]`

The `ldap` parameter has been added for importing the x.509 host certificate.

`adminaccess certificate show [detailed] [imported-host [application {all | ddbboost | https | keysecure | rkm | application-list}] | imported-ca [application {all | cloud | ddbboost | ldap | login-auth | keysecure | rkm | support | application-list}] | host | ca | subject subject-name | fingerprint fingerprint`

The `ldap` parameter has been added.

adminaccess guidelines and restrictions

- FTP and FTPS are mutually exclusive. Only one or the other can be enabled, not both.
- SCP can be enabled only when SSH is enabled.
- The following characters are invalid for SCP:

- dollar sign (\$)
 - parenthesis [(and)]
 - plus sign (+)
 - square brackets ([and])
 - semi-colon (;)
 - comma (,)
 - apostrophe (unslanted single quotation mark)
 - single slanted quotation mark (‘)
 - ampersand (&)
 - number sign (#)
 - less-than sign (<)
 - greater-than sign (>)
 - vertical bar (|)
 - exclamation mark (!)
- FTP and Telnet are disabled by default.

adminaccess add

```
adminaccess add ssh-keys [user username]
```

Add an SSH public key to the SSH authorized keys file on the Data Domain system. Admin role users can add and delete ssh-keys for other users. User role users can add or delete ssh-keys for their username only. Specify a username to associate the user with the key. When prompted, enter the key, press **Enter**, and press **Ctrl-D**. Role required: admin, limited-admin, security, user, or backup-operator.

Note

For Data Domain high availability (HA) systems, SSH keys created on the active node take 30 seconds to one minute to propagate to the standby node.

adminaccess authentication

```
adminaccess authentication add {cifs}
```

Allow Windows domain users with no local account on the Data Domain system to access the system through SSH, Telnet, and FTP using Windows domain group credentials. For administrative access, the user must be in the standard Windows Domain Admins group or in a group that you create named Data Domain. Users from both group names are always accepted as administrative users. The command also gives user-level access (no administrative operations allowed) to all other users from the domain. Users must be from the domain that includes the Data Domain system or a related, trusted domain.

The SSH, Telnet, or FTP command that accesses the Data Domain system must include the domain name, a backslash, and the user name in double quotation marks.

Note

CIFS must be enabled and the Data Domain system must be part of a Windows domain.

Role required: admin, limited-admin.

```
adminaccess authentication del {cifs}
```

Prevent authentication of a Windows domain. Allow admin role only for users with local user accounts on the Data Domain system. Role required: admin, limited-admin.

```
adminaccess authentication reset {cifs}
```

Reset the Windows user access to the default of requiring a local account for administrative access to the Data Domain system. Role required: admin, limited-admin.

```
adminaccess authentication show
```

Display whether CIFS authentication is enabled or disabled. Role required: admin, limited-admin, security, user, or backup-operator.

adminaccess certificate

```
adminaccess certificate cert-signing-request delete
```

Delete the certificate signing request. To see if there is a certificate signing request on the system, enter `adminaccess certificate cert-signing-request show`.

Role required: admin, limited-admin.

```
adminaccess certificate cert-signing-request generate [key-
strength {1024bit | 2048bit | 3072bit | 4096bit}] [country
country-code] [state state] [city city] [org-name organization-
name] [org-unit organization-unit] [common-name common-name]
```

Generate a Certificate Signing Request (CSR) file 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 Certificate Authority (CA). After you receive a signed CSR file from a CA, you can import the certificate using `adminaccess certificate import`.

If a previously generated CSR exists, the system prompts you to approve or reject certificate regeneration, which replaces the existing CSR.

If the user does not specify values, default values are used.

Note

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

Role required: admin, limited-admin.

Argument Definitions

If the value you set for a variable includes any space characters, enclose the variable string in quotes.

key strength

Enumeration values allowed are 1024 bit, 2048 bit, 3072 bit, or 4096 bit. Default is 2048 bit.

country

Default is US. Abbreviation for country cannot exceed two characters. No special characters are allowed.

state

Default is California. Maximum entry is 128 characters.

city

Default is Santa Clara. Maximum entry is 128 characters.

org-name

Default is My Company Ltd. Maximum entry is 64 characters.

org-unit

Default value is empty string. Maximum entry is 64 characters.

common name

Default value is the system host name. Maximum entry is 64 characters.

```
adminaccess certificate cert-signing-request show
```

Show the certificate signing request stored on the system. Role required: admin, security, user, backup-operator, or none.

```
adminaccess certificate delete { subject subject-name |
fingerprint fingerprint} [application {all | cloud | ddbboost |
ldap | login-auth | https | keysecure | rkm | support |
application-list}]
```

Delete a certificate for the specified application.

Note

Log out from the browser session before deleting an HTTPS host certificate. Otherwise HTTPS browser sessions (using imported host certificates) are closed. After deleting the host certificate, refresh or restart the browser to proceed.

Role required: admin, limited-admin.

Argument Definitions

If the value you set for a variable includes any space characters, enclose the variable string in quotes.

all

Deletes the certificates for all applications.

application-list

To delete the certificates for multiple applications, replace *application-list* with the application names, separated by commas or spaces (for example, *ddbboost, rkm*).

cloud

Deletes the certificate for the cloud application.

ddbboost

Deletes the certificate for the DD Boost application.

https

Deletes the certificate for the HTTPS application.

keysecure

Deletes the certificate for the KeySecure application.

ldap

Deletes the certificate for the LDAP application.

imported-ca application

Indicates that the certificate to be deleted is a CA certificate for the specified applications.

imported-host application

Indicates that the certificate to be deleted is a host certificate for the specified applications.

login-auth

Deletes login authorization.

rkm

Deletes the certificate for the RSA Key Manager (RKM) application.

support

Deletes the certificate for the support application.

```
adminaccess certificate delete { subject subject-name |
fingerprint fingerprint} [{all | cloud | ddbboost | ldap |
login-auth | https | keysecure | rkm | support | application-
list}]
```

Delete a certificate for the specified subject, fingerprint, or application.

Note

Log out from the browser session before deleting an HTTPS host certificate. Otherwise HTTPS browser sessions (using imported host certificates) are closed. After deleting the host certificate, refresh or restart the browser to proceed.

Role required: admin, limited-admin.

Argument Definitions

If the value you set for a variable includes any space characters, enclose the variable string in quotes.

all

Deletes the certificates for all applications.

application-list

To delete the certificates for multiple applications, replace *application-list* with the application names, separated by commas or spaces (for example, *ddbboost, rkm*).

cloud

Deletes the certificate for the cloud application.

ddbboost

Deletes the certificate for the DD Boost application.

fingerprint

Specifies the fingerprint of a certificate to be deleted. To display the available certificates and their footprints, enter `adminaccess certificate show`.

keysecure

Deletes the certificate for the KeySecure application.

ldap

Deletes the certificate for the LDAP application.

login-auth

Deletes login authorization.

https

Deletes the certificate for the HTTPS application.

rkm

Deletes the certificate for the RSA Key Manager (RKM) application.

subject

Specifies the subject name of a certificate to be deleted. To display the available certificates and their subject names, enter `adminaccess certificate show`.

support

Deletes the certificate for the support application.

```
adminaccess certificate export imported-ca {subject subject-name | fingerprint fingerprint} [file file-name]
```

Export a CA certificate for the specified subject name or fingerprint. The certificate appears on screen after the CLI command. Role required: admin, limited-admin.

Argument Definitions

If the value you set for a variable includes any space characters, enclose the variable string in quotes.

file

Saves a copy of the certificate in the `/ddvar/certificates` directory using the specified filename.

fingerprint

Specifies the fingerprint of a certificate to be exported. To display the available certificates and their footprints, enter `adminaccess certificate show`.

subject

Specifies the subject name of a certificate to be exported. To display the available certificates and their subject names, enter `adminaccess certificate show`.

```
adminaccess certificate export imported-ca-for-host application {ddboost | rkm | application-list} [file file-name]
```

Export a CA certificate for the specified host applications. The certificate appears on screen after the CLI command. Role required: admin, limited-admin.

Argument Definitions

If the value you set for a variable includes any space characters, enclose the variable string in quotes.

application-list

To export the certificates for multiple applications, replace *application-list* with the application names, separated by commas or spaces (for example, `ddboost, rkm`).

ddboost

Exports the certificate for the DD Boost application.

rkm

Exports the certificate for the RSA Key Manager (RKM) application.

file

Saves a copy of the certificate in the `/ddvar/certificates` directory using the specified filename.

```
adminaccess certificate generate self-signed-cert [regenerate-ca]
```

Generate a self-signed CA certificate and host certificate. The `regenerate-ca` option invalidates existing trust with external systems. Secure communication to trusted hosts is interrupted until mutual trust is reestablished. Role required: admin, limited-admin.

```
adminaccess certificate import {host application {all | ddboost | https | keysecure | rkm | application-list} | ca application {all | cloud | ddboost | ldap | login-auth | keysecure | rkm | application-list}} [file file-name]
```

Imports a certificate for one or more applications. You can import only one certificate per application, but you can use the same certificate for multiple applications.

To prepare for importing a certificate, use SCP, FTP, or FTPS to copy the host or CA certificate to the directory: `/ddvar/certificates`. Optionally, you can copy and paste the entire PEM file of the host certificate, and then run the import command without specifying the certificate filename. An error is generated if users mistakenly import a mismatched certificate; for example, importing a host certificate as a CA certificate, or vice versa.

Directly importing an RKM auto-registered certificate and bulk importing multiple certificates is not supported. After a public host certificate is imported, any related CSR is deleted from the system.

- Users must provide the PKCS12 file and password to decrypt the PKCS12 file.
- CA certificates must be imported in PEM format.

When importing or deleting certificates on an encrypted Data Domain system on which the system passphrase is set, the imported host PKCS12 certificate is reencrypted with the system passphrase. If the system passphrase is not set, an error is generated during the import.

When the system passphrase is changed, the imported host PKCS12 certificate, if present on Data Domain system, is reencrypted using the new system passphrase.

The correct server or client extensions must also be set. See the sections “Basic Constraints,” “Key Usage,” and “Extended Key Usage” in RFC 5280 for details (<http://www.ietf.org/rfc/rfc5280.txt>). Extensions are provided for host certificates during the certificate signing process.

Note

When a certificate is imported for HTTPS (which is used by DD System Manager), running this command closes any current browser sessions. It is a good practice to log out of the DD System Manager sessions prior to running this command.

Role required: admin, limited-admin.

Argument Definitions

If the value you set for a variable includes any space characters, enclose the variable string in quotes.

all

Imports the same certificate for all applications.

application-list

To import a certificate for multiple applications, replace *application-list* with the application names, separated by commas or spaces (for example, *ddbboost, rkm*).

ca application

Indicates that the certificate to be imported is a CA certificate.

cloud

Imports the certificate for the cloud application.

ddbboost

Imports the certificate for the DD Boost application.

host application

Indicates that the certificate to be imported is a host certificate.

ldap

Imports the certificate for the LDAP application.

keysecure

Imports the certificate for the KeySecure application.

login-auth

Imports login authorization.

rkm

Imports the certificate for the RSA Key Manager (RKM) application.

file

Specifies a file from which to import the certificate. A copy of the certificate file must be in the `/ddvar/certificates` directory.

Example 1

On the local system, enter:

```
# scp host.p12 <administrator_role>@<DD>:/ddvar/certificates
```

On the Data Domain system, enter:

```
# adminaccess certificate import host application https file host.p12
```

```
adminaccess certificate show [detailed] [imported-host
[application {all | ddbboost | https | keysecure | rkm |
application-list}] | imported-ca [application {all | cloud |
ddbboost | ldap | login-auth | keysecure | rkm | support |
application-list}] | host | ca | subject subject-name |
fingerprint fingerprint
```

Display certificates for the imported host, CA, imported CA, or support bundle server trusted CA. All users may run this command option. Role required: admin, limited-admin, security, user, backup-operator, or none.

Argument Definitions

If the value you set for a variable includes any space characters, enclose the variable string in quotes.

all

Displays the certificates in use for all applications.

application-list

Specifies a list of applications for which certificates are displayed.

ca application

Displays CA certificates for the specified applications.

cloud

Displays the certificate for the cloud application.

ddboost

Displays the certificate for the DD Boost application.

fingerprint

Displays the certificate with the specified fingerprint. To display the available certificates and their footprints, enter `adminaccess certificate show`.

host application

Displays host certificates for the specified applications.

https

Displays the certificate for the HTTPS application.

imported-ca

Specifies the CA certificates are to be displayed for the specified applications.

imported-host

Specifies the host certificates are to be displayed for the specified applications.

ldap

Displays the certificate for the LDAP application.

keysecure

Displays the certificate for the KeySecure application.

login-auth

Displays login authorizations.

rkm

Displays the certificate for the RSA Key Manager (RKM) application.

subject

Displays all certificate that use the specified subject. To display the available certificates and their subject names, enter `adminaccess certificate show`.

support

Displays the certificate for the support application.

```
adminaccess certificate cert-revoke-list delete {issuer issuer-name | fingerprint fingerprint} [application {all | cloud | ddboost | login-auth | application-list}]
```


Deletes the certificate revoke list for the specified issuer, fingerprint, or application.

Role required: admin, limited-admin.

Argument Definitions

If the value you set for a variable includes any space characters, enclose the variable string in quotes.

all

Remove all CRL related to any application on system.

application-list

To remove the certificate revoke lists for multiple applications, replace *application-list* with the application names, separated by commas or spaces (for example, `ddbboost, rkm`).

ddbboost

Remove CRL which are imported for DDBoost application.

fingerprint

Specifies the fingerprint of a certificate revoked list to be removed. To display the available certificates and their footprints, enter `adminaccess certificate show`.

issuer

Specifies the issuer of a certificate on the revoked list to be removed. To display the available certificates and their issuer, enter `adminaccess certificate show`.

login-auth

Remove CRL which are imported for making decision for certificate based login

```
adminaccess certificate cert-revoke-list import application
{all | cloud | ddbboost | login-auth | application-list} [file
file-name ]
```

Import certificate revocation list file for applications.

Can transfer the file first on DD under `/ddr/var/certificates` directory via `scp` and then import the certificate.

Argument Definitions

If the value you set for a variable includes any space characters, enclose the variable string in quotes.

all

Import all CRL related to any application on system.

application-list

To import the certificate revoke lists for multiple applications, replace *application-list* with the application names, separated by commas or spaces (for example, `ddbboost, rkm`).

ddbboost

Import CRL which are imported for DDBoost application.

login-auth

Import CRL which are imported for making decision for certificate based login

```
adminaccess certificate cert-revoke-list show [detailed]
{application {all | cloud | ddbboost | login-auth | application-
list} | issuer issuer-name | fingerprint fingerprint}
```

Show the certificate revocation list file present on the system using application, CRL issuer or fingerprint options.

Argument Definitions

If the value you set for a variable includes any space characters, enclose the variable string in quotes.

all

Show all CRL related to any application on system.

application-list

To show the certificate revoke lists for multiple applications, replace *application-list* with the application names, separated by commas or spaces (for example, *ddbboost, rkm*).

cloud

Show CRL which are imported for cloud application.

ddbboost

Show CRL which are imported for DDBoost application.

fingerprint

Specifies the fingerprint of a certificate revoked list to be shown. To display the available certificates and their footprints, enter `adminaccess certificate show`.

issuer

Specifies the issuer of a certificate on the revoked list to be shown. To display the available certificates and their issuer, enter `adminaccess certificate show`.

login-auth

Show CRL which are imported for making decision for certificate based login

adminaccess del

```
adminaccess del ssh-keys lineno [user username]
```

Delete an SSH key from the key file. Users may delete their own keys, and users in admin role may delete user keys. Run the command option `adminaccess show ssh-keys` to view line number values. Role required: admin, limited-admin, security, user, backup-operator, or none.

adminaccess disable

```
adminaccess disable {http | https | ftp | ftps | telnet | ssh |
scp | web-service | all}
```

Disable system access using the specified protocol. Disabling FTP or Telnet does not affect entries in the access lists. If all access is disabled, the Data Domain system is available only through a serial console or keyboard and monitor. Role required: admin, limited-admin.

```
#adminaccess disable web-service
Web Service:    disabled
```

adminaccess enable

```
adminaccess enable {http | https | ftp | ftps | telnet | ssh |
scp | web-service | all}
```

Enable a protocol on the Data Domain system. By default, SSH, HTTP, HTTPS, and web-service are enabled and FTP and Telnet are disabled. HTTP and HTTPS allow users to log in from System Manager. The web-service allows the use of REST APIs. To use FTP and Telnet, users with admin role permissions must add host machines to the access lists. Role required: admin, limited-admin.

```
#adminaccess enable web-service
Web Service:    enabled
```

adminaccess ftp

```
adminaccess ftp add host-list
```

Add one or more hosts to the FTP list. You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address. Host entries cannot include a space. Multiple entries may be separated by commas, spaces, or both. Role required: admin, limited-admin.

Note

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

```
adminaccess ftp del host-list
```

Remove one or more hosts (IP addresses, hostnames, or asterisks) from the FTP list. Multiple entries may be separated by commas, spaces, or both. Role required: admin, limited-admin.

```
adminaccess ftp option reset [session-timeout]
```

Reset the FTP options to default values. The default timeout setting never times out a session. Role required: admin, limited-admin.

```
adminaccess ftp option set session-timeout timeout-in-secs
```

Set the FTP client session timeout. The timeout range is 60 to 31,536,000 seconds. The default setting never times out a session. Role required: admin, limited-admin.

```
adminaccess ftp option show
```

Show the current FTP options. Role required: admin, limited-admin, security, user, backup-operator, or none.

adminaccess ftps

```
adminaccess ftps add host-list
```

Add one or more hosts to the FTPS list. You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address. Host entries cannot include a space. Multiple entries may be separated by commas, spaces, or both. Role required: admin, limited-admin.

Note

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

```
adminaccess ftps del host-list
```

Remove one or more hosts (IP addresses, hostnames, or asterisk) from the FTPS list. Host entries may be separated by commas, spaces, or both. Role required: admin, limited-admin.

```
adminaccess ftps option reset [session-timeout]
```

Resets the FTPS options to default values. The default timeout setting never times out a session. Role required: admin, limited-admin.

```
adminaccess ftps option set session-timeout timeout-in-secs
```

Sets the FTPS client session timeout. The timeout range is 60 to 31,536,000 seconds. The default setting never times out a session. Role required: admin, limited-admin.

```
adminaccess ftps option show
```

Shows the current FTPS options. Role required: admin, limited-admin, security, user, backup-operator, or none.

adminaccess http

```
adminaccess http add host-list
```

Add one or more hosts to the HTTP/HTTPS list. You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address. Host entries cannot include a space. Multiple entries may be separated by commas, spaces, or both. Role required: admin, limited-admin.

```
adminaccess http del host-list
```

Remove one or more hosts (IP addresses, hostnames, or asterisk) from the HTTP/HTTPS list. Host entries may be separated by commas, spaces, or both. Role required: admin, limited-admin.

adminaccess option

```
adminaccess option reset [login-max-attempts | login-unlock-timeout | login-max-active]
```

Resets the specified option to the default value, which is 4 for maximum attempts, 120 seconds for the unlock timeout, and 10 for active sessions. Note that the limit set for active sessions does not apply to the sysadmin user. Role required: admin, limited-admin.

Example 2

```
# adminaccess option reset login-max-active
Adminaccess option "login-max-active" is reset to default (10).
```

```
adminaccess option reset cipher-list
```

Resets the cipher list to be used for server communication. Role required: admin, limited-admin.

```
adminaccess option reset password-auth
```

Resets password-based authentication method to its default value of enabled. Role required: admin.

```
adminaccess option reset password-hash
```

Resets the password hash to its default value of MD5. Role required: admin.

```
adminaccess option set cipher-list
```

Sets the cipher list to be used for server communication. Role required: admin, limited-admin.

```
adminaccess option set login-max-active {unlimited | count}
```

Specifies the maximum number of active sessions for each user. The minimum value is 1, and the default value is 10. Note that the limit set for active sessions does not apply to the sysadmin user. Role required: admin, limited-admin.

Example 3

```
# adminaccess option set login-max-active 5
Adminaccess option "login-max-active" set to "5".
```

```
adminaccess option set login-max-attempts count
```

Specifies the maximum number of login attempts before a mandatory lock is applied to an account. The user cannot log in while the account is locked. The range is 4 to 20, and the default value is 4. Role required: admin, limited-admin.

Example 4

```
# adminaccess option set login-max-attempts 5
Adminaccess option "login-max-attempts" set to "5".
```

```
adminaccess option set login-unlock-timeout timeout-in-secs
```

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 3600 seconds, and the default period is 120 seconds. Role required: admin, limited-admin.

Example 5

```
# adminaccess option set login-unlock-timeout 120
Adminaccess option "login-unlock-timeout" set to "120".
```

```
adminaccess option set password-auth {enabled | disabled}
```

Enable or disable password-based authentication. Password-based authentication cannot be disabled unless an SSH key exists for the administrative user. The command will provide a warning if no CA certificates for login authentication are imported, and prompt you to continue with disabling password-based authentication. The CA certificates are required to allow users to login after certificate-based authentication is configured. If a security policy is configured, the command will prompt for security officer credentials before making the configuration change. Role required: admin.

```
adminaccess option set password-hash {md5 | sha512}
```

Specifies the hashing algorithm used to store local user passwords. After changing the hashing algorithm, update the passwords for existing local users to be compatible with the new algorithm. Role required: admin.

Example 6

```
# adminaccess option set password-hash sha512
Adminaccess option "password-hash" set to "sha512".
Please update existing local users' passwords, to hash with "sha512".
```

Example 6 (continued)

```
adminaccess option show
```

Shows the current configuration for the `adminaccess option` command. Role required: admin, limited-admin.

Example 7

```
# adminaccess option sh
Option                               Value
-----                               -
login-unlock-timeout                 120
login-max-attempts                    4
login-max-active                      10
cipher-list                           default*
password-auth                         enabled
password-hash                         md5
-----                               -
(*) Run 'adminaccess option show cipher-list' for detail.
```

```
adminaccess option show cipher-list
```

Shows the current configuration for the `adminaccess option cipher-list` command. Role required: admin, limited-admin.

adminaccess reset

```
adminaccess reset {http | https | ftp | ftps | telnet | ssh |
scp | all}
```

Reset one or more protocols to their default states and clear the access lists of host entries. Output shows the running state of each protocol. Role required: admin, limited-admin.

Note

Because SCP works together with SSH, output appears the same for both. However, due to the registry configuration, output could be misleading. For example, if SSH is disabled, SCP also shows as disabled; however, SCP is enabled at the registry level. This is expected behavior and does not affect functionality. When a user resets SCP, the SCP registry entry changes to enabled and output for SSH shows as enabled.

```
adminaccess reset ssh-keys [user username]
```

Remove the authorized SSH keys file for the specified user from the Data Domain system. After removing the file, every SSH connection requires password authentication. Role required: varies as listed below.

- Users may reset their own keys only.
- *Admin* role users may reset the keys of any user.
- *Security* role users and *none* role users may not reset keys.

adminaccess show

```
adminaccess show
```

Lists the access services available on a Data Domain system and displays option values for the access services that are enabled. Role required: admin, limited-admin, security, user, backup-operator, or none.

- N/A means the service does not use an access list.
- A hyphen means the service can use an access list, but the access list does not contain host names.
- An asterisk means the service allows all hosts.

```
# adminaccess show
Service      Enabled    Allowed Hosts
-----
ssh          yes       -
scp          yes       (same as ssh)
telnet      no        -
ftp         yes       *
ftps        no        -
http        yes       -
https       yes       -
web-service yes       N/A
-----

Ssh/Scp options:
Option      Value
-----
session-timeout default (infinite)
server-port default (22)
-----

Telnet options:
Option      Value
-----
session-timeout default (infinite)
-----

Ftp options:
Option      Value
-----
session-timeout default (infinite)
-----

Ftps options:
Option      Value
-----
session-timeout default (infinite)
-----

Web options:
Option      Value
-----
http-port   80
https-port  443
session-timeout 10800
-----
```

```
adminaccess show ssh-keys [user username]
```

Displays the authorized SSH key file with a line number for each entry. *Admin* role users can view the SSH key files of any user. Users in other roles can view only their own SSH key file.

adminaccess ssh

```
adminaccess ssh add host-list
```

Add one or more hosts to the SSH list. You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address. Host entries cannot include a space. Multiple entries may be separated by commas, spaces, or both. Role required: admin, limited-admin.

```
adminaccess ssh del host-list
```

Remove one or more hosts (IP addresses, hostnames, or asterisks) from the SSH list. Host entries may be separated by commas, spaces, or both. Role required: admin, limited-admin.

```
adminaccess ssh option reset [ciphers | macs | server-port | session-timeout]
```

Resets the ssh options to their default values. Role required: admin, limited-admin.

Table 3 Default option values

Option	Value
session-timeout	default (infinite)
server-port	default (22)
ciphers	aes128-cbc, chacha20-poly1305@openssh.com, aes256-gcm@openssh.com, aes128-gcm@openssh.com, aes256-ctr, aes192-ctr, aes128-ctr
macs	hmac-sha1, hmac-sha2-512-etm@openssh.com, hmac-sha2-256-etm@openssh.com, hmac-ripemd160-etm@openssh.com, umac-128-etm@openssh.com, hmac-sha2-512, hmac-sha2-256, hmac-ripemd160, umac-128@openssh.com

```
adminaccess ssh option set ciphers cipher-list
```

Sets the ciphers to be used by SSH daemon.

```
# adminaccess ssh option set ciphers "chacha20-
poly1305@openssh.com,aes256-gcm@openssh.com,aes128-gcm@openssh.com"
Adminaccess ssh option "ciphers" set to "chacha20-
poly1305@openssh.com,aes256-gcm@openssh.com,aes128-gcm@openssh.com".
```

```
adminaccess ssh option set macs MAC-list
```

Sets the MACs to be used by SSH daemon.

```
# adminaccess ssh option set macs "hmac-sha2-512-etm@openssh.com,hmac-
sha2-256-etm@openssh.com,hmac-ripemd160-etm@openssh.com"
Adminaccess ssh option "macs" set to "hmac-sha2-512-
etm@openssh.com,hmac-sha2-256-etm@openssh.com,hmac-ripemd160-
etm@openssh.com".
```

```
adminaccess ssh option set server-port port-number
```

Set the SSH server port. The default port number is 22. Role required: admin, limited-admin.

```
adminaccess ssh option set session-timeout timeout-in-secs
```

Set the SSH client timeout options. The timeout range is 60 to 31,536,000 seconds. The default setting never times out a session. Role required: admin, limited-admin.

Example 8

Set the SSH session timeout period to 10 minutes:

```
# adminaccess ssh option set session-timeout 600
```

```
adminaccess ssh option show
```

Display the SSH option configuration. Role required: admin, limited-admin.

Example 9

Example 9 (continued)

```
# adminaccess ssh option show
Option                               Value
-----
session-timeout                       default (infinite)
server-port                           default (22)
ciphers                               chacha20-poly1305@openssh.com,aes256-gcm@openssh.com,aes128-gcm@openssh.com
macs                                   hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com,hmac-ripemd160-
etm@openssh.com
```

adminaccess telnet

```
adminaccess telnet add host-list
```

Add one or more hosts to the Telnet list. You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address. Host entries cannot include a space. Multiple entries may be separated by commas, spaces, or both. Role required: admin, limited-admin.

```
adminaccess telnet delete host-list
```

Remove one or more hosts (IP addresses, hostnames, or asterisk) from the Telnet list. Host entries may be separated by commas, spaces, or both. Role required: admin, limited-admin.

```
adminaccess telnet option reset [session-timeout]
```

Reset the client session timeout period to the default value, which does not time out sessions. Role required: admin, limited-admin.

```
adminaccess telnet option set session-timeout timeout-in-secs
```

Set the client session timeout period to the specified number of seconds. If no data is received from a Telnet client within the timeout period, and if the client does not respond to a subsequent prompt message, the session terminates. The valid range is from 60 to 31536000 (365 days).

To configure the Data Domain system to prevent sessions from timing out, use the `adminaccess telnet option reset` command. Role required: admin, limited-admin.

Example 10

To set the SSH session timeout period to 10 minutes:

```
# adminaccess telnet option set session-timeout 600
```

```
adminaccess telnet option show
```

Display the Telnet configuration. Role required: admin, limited-admin.

adminaccess trust

```
adminaccess trust add host hostname [type mutual]
```

Establishes the (mutual) trust with the specified host. Role required: admin, limited-admin.

```
adminaccess trust copy {source | destination} hostname
```

Copy all trust to or from the specified host. Role required: admin, limited-admin.

```
adminaccess trust del host hostname [type mutual]
```

Remove the mutual trust from the specified host. Role required: admin, limited-admin.

```
adminaccess trust show [hostname]
```

Show the list of trusted Certificate Authorities (CAs). Role required: admin, limited-admin, security, user, backup-operator, or none.

Example 11 Establishing mutual trust

Establishing mutual trust between two nodes requires the following steps. The example below uses the nodes dd-system-1 and dd-system-2.

1. On dd-system-1, run the `adminaccess certificate show` command to display the pre-created trusted CA certification.
2. On dd-system-1, run the `adminaccess trust add host <dd-system-2-hostname> type mutual` command to create trust with the specified host.
3. On dd-system-1, run the `adminaccess trust show hostname` command to verify the mutual trust relationship is created between the two nodes.
4. On dd-system-2, run the `adminaccess trust show hostname` command to verify the mutual trust relationship is created between the two nodes.

adminaccess web

```
adminaccess web option reset [http-port | https-port | session timeout]
```

Reset the Web options to default values. Role required: admin, limited-admin.

```
adminaccess web option set http-port port-number
```

Set the HTTP access port for the Web client. Default is port 80. Role required: admin, limited-admin.

```
adminaccess web option set https-port port-number
```

Set the HTTPS access port for the Web client. Default is port 443. Role required: admin, limited-admin.

```
adminaccess web option set session-timeout timeout-in-secs
```

Set the Web client session timeout. Range is 300 to 31536000 seconds; the default is 10800 seconds. Role required: admin, limited-admin.

```
adminaccess web option show
```

Show the current values for Web options. Role required: admin, limited-admin.

CHAPTER 2

alerts

The `alerts` command manages current alerts, alert notification groups, and alerts history. When a user logs in, a message is shown indicating the presence of alerts and instructions on how to proceed.

Command options enable sending email to a designated recipient or notification group when an event occurs within the Data Domain system. Depending on the option, information includes alert type, date posted, and resulting action. More than three months of alert history is retained.

The default alert notification group (“default”) is configured to send alerts for any event class with severity level of Warning or above. Email notifications are sent to Data Domain Support at `autosupport-alert@autosupport.datadomain.com`. The default alert notification group can only be reset to default values: it cannot be destroyed.

Some event types, such as those in the environment class that pertain to temperature sensors within the chassis, are detected repeatedly if the underlying condition is not corrected.

This chapter contains the following topics:

- [alerts change history](#) 44
- [alerts clear](#) 44
- [alerts notify-list](#) 44
- [alerts show](#) 47
- [alerts test events](#) 49

alerts change history

There have been no changes to this command in this release.

alerts clear

```
alerts clear alert-id alert-id-list
```

Clear an active alert or list of alerts. Role required: admin, limited-admin.

Argument Definitions

alert-id-list

List of alert identification numbers. To display the alert ID numbers, enter `alerts show all`.

alerts notify-list

```
alerts notify-list add group-name {[class class-list [severity severity]] [emails email-addr-list]}
```

Modify a notification group by adding an event class, severity level, or recipient email address. The system does not accept EMC or Data Domain email IDs. Role required: admin, limited-admin, tenant-admin.

Argument Definitions

class class-list

List of event classes: `cifs`, `cloud`, `cluster`, `environment`, `filesystem`, `firmware`, `ha`, `hardwareFailure`, `network`, `replication`, `security`, `storage`, `syslog`, and `systemMaintenance`.

emails email-addr-list

Email addresses of members in an alert notification group.

group-name

Name of alert notification group.

severity severity

Severity level of event class. The severity levels in decreasing order of severity are: `emergency`, `alert`, `critical`, `error`, `warning`, `notice`, `info`, and `debug`. Default is `warning`.

Example 12

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

```
alerts notify-list create group-name {class class-list [severity severity] | tenant-unit tenant-unit | tenant <tenant> }
```

Subscribe to a notification list, add a class and a severity level to an existing list, or add members to a notification group on a Data Domain system or tenant unit. Role required: admin, limited-admin.

Argument Definitions

class *class-list*

List of event classes: cifs, cloud, cluster, environment, filesystem, firmware, ha, hardwareFailure, network, replication, security, storage, syslog, and systemMaintenance.

emails *email-addr-list*

Email addresses of members in an alert notification group.

group-name

Name of alert notification group.

severity *severity*

Severity level of event class. The severity levels in decreasing order of severity are: emergency, alert, critical, error, warning, notice, info, and debug. Default is warning.

Example 13

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

```
alerts notify-list del group-name { [class class-list] [emails
email-addr-list]}
```

Modify an alert notification group by deleting event classes, email recipients, or both. The system does not accept EMC or Data Domain email IDs for deletion.

Security officer authorization is required only if the *group-name* severity level is set to Warning or above and the command is run on a Retention Lock Compliance system. See the *Data Domain Operating System Administration Guide* for details on alerts. Role required: admin, limited-admin, tenant-admin.

Argument Definitions

class *class-list*

List of event classes: cifs, cloud, cluster, environment, filesystem, firmware, ha, hardwareFailure, network, replication, security, storage, syslog, and systemMaintenance.

emails *email-addr-list*

Email addresses of members in an alert notification group.

group-name

Name of alert notification group.

Example 14

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

```
alerts notify-list destroy group-name
```

Delete an alert notification group. Note that the default alert notification group cannot be destroyed.

Security officer authorization is required only if the *group-name* severity level is set to Warning or above and the command is run on a Retention Lock Compliance system.

See the *Data Domain Operating System Administration Guide* for details on alerts. Role required: admin, limited-admin.

Argument Definitions

group-name

Name of alert notification group.

```
alerts notify-list reset
```

Remove all user-created alert notification groups and restore the default notification group email list to factory defaults.

Role required: admin, limited-admin. Security officer authorization is required for systems with Retention Lock Compliance.

```
alerts notify-list show [group group-name | email email-addr] |
tenant-unit tenant-unit | tenant <tenant>]
```

Display the configuration of all notification lists, or display the lists associated with the specified group, email list, or tenant unit. The system does not accept EMC or Data Domain email IDs. Role required: admin, limited-admin, tenant-admin, security, user, tenant-user, backup-operator, or none.

Argument Definitions

emails email-addr-list

Email addresses of members in an alert notification group.

group-name

Name of alert notification group.

tenant unit

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system. See the *Data Domain Operating System Administration Guide* for more information on SMT.

Example 15

```
# alerts notify-list show eng_lab mlee@yourcompany.com
```

```
alerts notify-list test {group group-name | email email-addr}
```

Send a test notification to an alert notification group or email address. Role required: admin, limited-admin, tenant-admin, security, user, backup-operator, or none.

Argument Definitions

emails email-addr-list

Email addresses of members in an alert notification group.

group-name

Name of alert notification group.

Example 16

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

alerts show

```
alerts show all [local]
```

Display details on all alert notification groups. The `local` argument only applies to cluster configurations, which are not supported in this release. The `local` argument will be removed in a future release. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
alerts show current [local] [tenant-unit tenant-unit]
```

Display a list of currently active alerts on a Data Domain system or tenant unit. The `local` argument only applies to cluster configurations, which are not supported in this release. The `local` argument will be removed in a future release. Role required: admin, limited-admin, tenant-admin, security, user, tenant-user, backup-operator, or none.

Argument Definitions

tenant unit

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system. See the *Data Domain Operating System Administration Guide* for more information on SMT.

```
alerts show current-detailed [local] [alert-id alert-id-list ]
[ tenant-unit tenant-unit]
```

Display detailed information about currently active alerts on a Data Domain system or tenant unit. Role required: admin, limited-admin, tenant-admin, security, user, tenant-user, backup-operator, or none.

Argument Definitions

alert-id-list

List of alert identification numbers. To display the alert ID numbers, enter `alerts show all`.

local

The `local` argument only applies to cluster configurations, which are not supported in this release. The `local` argument will be removed in a future release.

tenant unit

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system. See the *Data Domain Operating System Administration Guide* for more information on SMT.

```
alerts show daily [local]
```

Display daily alert report, including current alerts and 24-hour alert history. The `local` argument only applies to cluster configurations, which are not supported in this release. The `local` argument will be removed in a future release. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
alerts show history [local] [tenant-unit tenant-unit] [last n
{hours | days | weeks | months}] [start MMDDhhmm [[CC]YY] end
MMDDhhmm [[CC]YY]
```

Display alert history on a Data Domain system or tenant unit. Default duration spans the last three months. Role required: admin, limited-admin, tenant-admin, security, user, tenant-user, backup-operator, or none.

Argument Definitions

alert-id-list

List of alert identification numbers. To display the alert ID numbers, enter `alerts show all`.

CC

Use with `start` or `end` arguments to `show` option. Specify first two digits of year. Default is 20.

end MMDDhhmm

Use with `show` option to display alerts for specific interval.

The argument *MMDD* indicates month and day of end date.

The argument *hhmm* indicates hours and minutes of end time (24-hour format).

To specify midnight between Sunday night and Monday morning, use `mon 0000`.

To specify noon on Monday, use `mon 1200`.

last n {hours | days | weeks | months}

Use with `show` option to display alerts for most recent number of *n* (hours, days, weeks, months).

local

The `local` argument only applies to cluster configurations, which are not supported in this release. The `local` argument will be removed in a future release.

start MMDDhhmm

Use with `show` option to display alerts for specific interval.

The argument *MMDD* indicates month and day of start date.

The argument *hhmm* indicates hours and minutes of `start` time (24-hour

format). To specify midnight between Sunday night and Monday morning, use

`mon 0000`. To specify noon on Monday, use `mon 1200`.

tenant unit

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system.

See the *Data Domain Operating System Administration Guide* for more information on SMT.

YY

Use with `start` or `end` arguments to `show` option. Specify last two digits of year.

```
alerts show history-detailed [local] [tenant-unit tenant-unit]
[last n {hours | days | weeks | months}] [start MMDDhhmm
[[CC]YY] end MMDDhhmm [[CC]YY]
```

Display detailed information about historic alerts on a Data Domain system or tenant unit. Default duration spans the last three months. The argument descriptions are the same as for `alerts show history`. Role required: admin, limited-admin, tenant-admin, security, user, tenant-user, backup-operator, or none.

alerts test events

```
alerts test events event-id-list [prompt-for-input]
```

Generate a test alert on the Data Domain system. The `prompt-for-input` prompts for user input as the command runs. This option requires interactive mode. Role required: admin, limited-admin.

Argument Definitions

event-id-list

List of event identification numbers, for example: EVT-STORAGE-00007.

alerts

CHAPTER 3

alias

The `alias` command creates, deletes, and displays command aliases for the Data Domain system command set. Users can manage aliases for only those commands permitted for the user's access role.

This chapter contains the following topics:

- [alias change history](#)52
- [alias add](#) 52
- [alias del](#) 53
- [alias reset](#)53
- [alias show](#)53

alias change history

There have been no changes to this command in this release.

alias add

```
alias add alias-name "command"
```

Add a command alias. Enter the alias name and command, and enclose the command name in quotation marks. The new alias is available only to the user who created it. Role required: admin, limited-admin, security, user, backup-operator, or none.

Default Command Aliases

The following command aliases are included with the system and available to all users.

date

system show date

df

filesystem show space

hostname

net show hostname

ifconfig

net config

iostat

system show stats

netstat

net show stats

nfsstat

nfs show stats

passwd

user change password

ping

net ping

poweroff

system poweroff

reboot

system reboot

sysstat

system show stats

traceroute

route trace

uname
system show version

uptime
system show uptime

who
user show active

alias del

```
alias del alias-name
```

Delete an alias by name. Role required: admin, limited-admin, security, user, backup-operator, or none.

alias reset

```
alias reset
```

Remove user-created aliases and restore defaults. Role required: admin, limited-admin, security, user, backup-operator, or none.

alias show

```
alias show
```

Display all aliases and command definitions. Role required: admin, limited-admin, security, user, backup-operator, or none.

alias

CHAPTER 4

archive

The `archive` command is used only on systems licensed to run the Data Domain Extended Retention software option (formerly Data Domain Archiver). Extended Retention command options enable the feature and configure policies. See the *Data Domain Operating System Administration Guide* for details on functionality, installation, and configuration.

This chapter contains the following topics:

- [archive change history](#) 56
- [archive data-movement](#) 56
- [archive disable](#) 57
- [archive enable](#) 58
- [archive option](#) 58
- [archive report](#) 58
- [archive show](#) 59
- [archive space-reclamation](#) 59

archive change history

There have been no changes to this command in this release.

archive data-movement

```
archive data-movement policy reset age-threshold mtrees mtree-list
```

Reset the age threshold for specified MTrees (*mtree-list* is a colon-separated list). Only files modified in the past (beyond the age threshold) are moved to the retention tier during the next data movement. Role required: admin, limited-admin.

```
archive data-movement policy reset default-age-threshold
```

Reset the age threshold to the default value, which is set with `archive data-movement policy set default-age-threshold`. The default age threshold applies to new MTrees and to MTrees for which the age threshold has not been set. Role required: admin, limited-admin.

```
archive data-movement policy set age-threshold {days | none}
mtreees mtree-list
```

Set the age threshold for specified MTrees (*mtree-list* is a colon-separated list). The value for *days* must be from 14 days to 18250 days (50 years). Role required: admin, limited-admin.

```
archive data-movement policy set default-age-threshold {days |
none}
```

Set the default age threshold. The argument *days* must be from 14 days to 18250 days (50 years). Role required: admin, limited-admin.

```
archive data-movement policy show [mtreees mtree-list]
```

View the data-movement policy for the specified MTrees (*mtree-list* is a colon-separated list). Role required: admin, limited-admin.

```
archive data-movement schedule reset
```

Reset the data-movement schedule to default values. Role required: admin, limited-admin.

```
archive data-movement schedule set {never | days days time time
[every 2wks]} [no-clean]
```

Set the schedule for data movement. Unless you specify *no-clean*, the file system is cleaned after data movement is completed. Note that the *days* argument checks two ranges (and can be either a space- or comma-separated list, or arbitrary text):

- Weekday (Monday-Sunday)
- Day of the month (1-31, regardless of month, plus “last” and “first”)

Any value outside of the two ranges generates an error message. Role required: admin.

Note

For days, “last” is converted to the value 31. If a schedule is set for the 31st of every month at 10:00 PM, it is not executed on months with fewer than 31 days. This is a known issue.

Example 17

Example 17 (continued)

To schedule data movement to occur each Tuesday at 6:00 a.m., enter:

```
# archive data-movement schedule set days "tue" time "06:00"
```

Example 18

To schedule data movement to occur on alternate Tuesdays at 6:00 a.m., enter:

```
# archive data-movement schedule set days "tue" time "06:00" every 2wks
```

```
archive data-movement schedule show
```

Display the data-movement schedule. Role required: admin, limited-admin, security, user, backup-operator.

```
archive data-movement start
```

Start data movement from the active tier to the retention tier. All files that satisfy the age-threshold value are moved to the retention tier. Data movement comprises five phases: seeding, scanning, verifying, packing, and installing. Role required: admin, limited-admin.

```
archive data-movement status
```

Display immediate, one-time view of output from the `archive data-movement start` command. Output shows the point at which data movement has progressed as of the time the command is issued. Role required: admin, limited-admin, security, user, backup-operator.

```
archive data-movement stop
```

Stop data movement to the retention tier. Role required: admin, limited-admin.

```
archive data-movement throttle reset
```

Reset the throttle value to 100 percent (no throttle). The throttle value will take effect without restarting file migration if it is running. Role required: admin, limited-admin.

```
archive data-movement throttle set {25 | 50 | 75 | 100 }
```

Set the throttle value to 25, 50, 75, or 100, where 25 is the slowest, and 100 is the fastest. The throttle value will take effect without restarting file migration if it is running. Role required: admin, limited-admin.

```
archive data-movement throttle show
```

Show the actual throttle value. Role required: admin, limited-admin.

```
archive data-movement watch
```

View data movement progress while the operation is running. If the operation has completed or is not running, output shows current status only. Role required: admin, limited-admin, security, user, backup-operator.

archive disable

```
archive disable
```

Disable the Extended Retention software option. Note that the file system must be destroyed (and all data lost) before Extended Retention can be disabled. Role required: admin, limited-admin.

archive enable

```
archive enable
```

Enable the Extended Retention software option. The file system must be disabled before Extended Retention can be enabled. Role required: admin, limited-admin.

Note

MTree replication is supported from Extended Retention systems to non-Extended Retention systems if both are running DD OS 5.5 or later.

archive option

```
archive option reset local-compression-type
```

Reset the local compression algorithm to the default value `gz` for subsequent data movement to the retention tier. You must restart the file system for this change to take effect. Role required: admin, limited-admin.

```
archive option set local-compression-type {none | lz | gzfast | gz}
```

Set the local compression algorithm for subsequent data movement to the retention tier. You must restart the file system for this change to take effect. Role required: admin, limited-admin.

```
archive option show [local-compression-type | data-movement-packing]
```

Display the local compression algorithm for the retention tier or the progress of the file migration process. Role required: admin, limited-admin, security, user, backup-operator.

archive report

```
archive report generate file-location [path {path-list | all}]
[output-file filename]
```

Create a report showing the name and location of each file in a directory, an MTree, or the entire namespace. If you specify `output-file filename`, the report is saved in this file under the fixed directory `/ddvar`. If the output file argument is not specified, the report is displayed in standard output. The command returns before the entire report is generated, and a footer indicates that the report is complete. Each line in the report contains a file name and its location. The location is shown as *Active* if the file completely resides in the active tier. If the file resides partially or completely in the retention tier, the retention unit name is shown for its location. An asterisk is appended to the line if the file contents span the active tier and retention unit. Role required: admin, limited-admin.

Example 19

To report files in a directory:

```
# archive report generate file-location path /backup/dir1 output-file
report.txt
```

or

Example 19 (continued)

```
# archive report generate file-location path /data/col1/mtree-2/dir3
output-file report.txt
```

To report files in an MTree:

```
# archive report generate file-location path /data/col1/mtree-2
output-file report.txt
```

To report files in the entire namespace:

```
# archive report generate file-location path all output-file
report.txt
```

or

```
# archive report generate file-location output-file report.txt
```

archive show

```
archive show config
```

Displays the Extended Retention configuration. Role required: admin, limited-admin, security, user, backup-operator.

Example 20

```
# archive show config
Enabled                               Yes
Data movement Schedule                Run on day(s) "tue" at "06:00" hrs
Data movement throttle                100 percent
Default age threshold data movement policy 0 days
Run filesys clean after archive data movement No
Archive Tier local compression        gz
Packing data during archive data movement disabled
Space Reclamation                     enabled (one-cycle)
Space Reclamation Schedule            Run on day(s) "fri" at "06:00" hrs
```

archive space-reclamation

```
archive space-reclamation resume
```

Resume the retention tier space reclamation process. Role required: admin, limited-admin.

```
archive space-reclamation schedule reset
```

Reset the space-reclamation schedule to the factory default. Role required: admin, limited-admin.

```
archive space-reclamation schedule set {never | days days time
time [every 2wks]}
```

Set the starting time for space-reclamation. Role required: admin, limited-admin.

```
archive space-reclamation schedule show
```

Show the space-reclamation schedule. Role required: admin, limited-admin, security, user, backup-operator.

```
archive space-reclamation start [one-cycle]
```

Start the space-reclamation process. You may optionally run space-reclamation for only one cycle. Role required: admin, limited-admin.

```
archive space-reclamation status
```

Show the status of the retention tier space-reclamation process. Role required: admin, limited-admin.

Example 21

```
# archive space-reclamation status
```

```
Space-reclamation was started on March 9 2014 10:25 and is currently running.
```

```
archive space-reclamation status-detailed
```

Show the detailed status of the retention tier space-reclamation process. Role required: admin, limited-admin.

```
archive space-reclamation stop
```

Stop the retention tier space-reclamation process. Role required: admin, limited-admin.

```
archive space-reclamation suspend
```

Suspend the retention tier space-reclamation process. Role required: admin, limited-admin.

CHAPTER 5

authentication

The `authentication` command manages NIS users, domains, groups and servers. Command options enable the Data Domain system to participate in an active Network Information Service (NIS) domain, which maintains a centralized repository of users, groups, and server names. NIS adds a global directory that authenticates users from any host on the network.

This chapter contains the following topics:

- [authentication change history](#) 62
- [authentication kerberos](#) 63
- [authentication ldap](#) 64
- [authentication nis](#) 67

authentication change history

New in DD OS 6.1.2

`authentication ldap base reset`

Reset the LDAP base suffix. Role required: admin, limited-admin.

`authentication ldap base set basename`

Set the LDAP base suffix. Role required: admin, limited-admin.

`authentication ldap client-auth reset`

Reset the LDAP client authentication configuration. Role required: admin, limited-admin.

`authentication ldap client-auth set binddn dn-name`

Set the account name to authenticate with the LDAP server. Role required: admin, limited-admin.

`authentication ldap disable`

Disable the LDAP client. Role required: admin, limited-admin.

`authentication ldap enable`

Enable the LDAP client. Role required: admin, limited-admin.

`authentication ldap reset`

Delete the LDAP configuration and set it to the default. Role required: admin, limited-admin.

`authentication ldap servers add server-list`

Add LDAP servers to the *server-list*. Role required: admin, limited-admin.

`authentication ldap servers del server-list`

Delete LDAP servers from the *server-list*. Role required: admin, limited-admin.

`authentication ldap servers reset`

Reset the LDAP servers to their default settings. Role required: admin, limited-admin.

`authentication ldap show`

Display the LDAP configuration. Role required: admin, limited-admin, security, user, backup-operator, or none.

`authentication ldap ssl disable`

Disable SSL for LDAP. Role required: admin, limited-admin.

`authentication ldap ssl enable [method {ldaps | start_tls}]`

Enable SSL for LDAP, and set the SSL method as ldaps or start_tls. The default is ldaps. Role required: admin, limited-admin.

`authentication ldap ssl reset tls_reqcert`

Reset the *tls_reqcert* option to the default value of demand. Role required: admin, limited-admin.

```
authentication ldap ssl set tls_reqcert {never | demand}
```

Specify the checks to perform on the server certificate. The default is demand.
Role required: admin, limited-admin.

```
authentication ldap status
```

Display the LDAP status. Role required: admin, limited-admin, security, user, backup-operator, or none.

authentication kerberos

```
authentication kerberos keytab import
```

Imports the krb5.keytab file from /ddvar to /ddr/etc. If the file is not present in /ddvar, then the command returns an error. Role required: admin, limited-admin.

```
authentication kerberos reset
```

Resets the realm, KDC, and so on, to default configuration. Unjoins a DDR from the domain and unjoins a DDR from a Linux KDC and a Windows KDC. Role required: admin, limited-admin.

Note

To make a DDR part of Windows AD, first unjoin the DDR from the Linux KDC using this command.

```
authentication kerberos set realm home-realm kdc-type {windows  
[kdc kdc-list] | unix kdc kdc-list}
```

Sets the realm for a system and enables Kerberos authentication on the realm. Role required: admin, limited-admin.

Argument definitions

home-realm

The Kerberos realm.

kdc-type

Key Distribution Center type - Windows or UNIX.

kdc-list

List of KDCs.

```
authentication kerberos show config
```

Displays Kerberos configuration. Role required: admin, limited-admin.

Example 22

```
authentication kerberos show config
```

```
Home Realm:      abc.com
KDC List:        10.10.10.10 10.10.10.11
KDC Type:        windows
```

Output definitions

Home Realm

The Kerberos Realm.

KDC List

List of KDCs.

KDC Type

Key Distribution Center type - Windows or UNIX.

authentication ldap

authentication ldap base reset

Reset the LDAP base suffix. Role required: admin, limited-admin.

```
# authentication ldap base reset
LDAP base-suffix reset to empty.
```

authentication ldap base set *basename*

Set the LDAP base suffix. Role required: admin, limited-admin.

```
# authentication ldap base set "dc=anvil,dc=team"
LDAP base-suffix set to "dc=anvil,dc=team".
```

authentication ldap client-auth reset

Reset the LDAP client authentication configuration. Role required: admin, limited-admin.

```
# authentication ldap client-auth reset
LDAP client authentication configuration reset to empty.
```

authentication ldap client-auth set binddn *dn-name*

Set the account name to authenticate with the LDAP server. Role required: admin, limited-admin.

```
# authentication ldap client-auth set binddn "cn=Manager,dc=u2,dc=team"
Enter bindpw:
LDAP client authentication binddn set to "cn=Manager,dc=u2,dc=team".
```

NoteIf the `bindpw` is not specified, the system requests unauthenticated access.

Argument Definitions**binddn**

Account name to use to authenticate with the LDAP server.

authentication ldap disable

Disable LDAP authentication. Role required: admin, limited-admin.

authentication ldap enable

Enable LDAP authentication. A file system restart is required the first time LDAP is enabled. LDAP and NIS cannot be enabled at the same time. Role required: admin, limited-admin.

authentication ldap reset

Delete the LDAP configuration and set it to the default. Role required: admin, limited-admin.

```
authentication ldap servers add server-list
```

Add LDAP servers to the *server-list*. Role required: admin, limited-admin.

Note

If you add an invalid server name and enable LDAP authentication, the system will continue to use the first valid server in the list, based on the set domain.

```
# 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
-----
```

```
authentication ldap servers del server-list
```

Delete LDAP servers from the *server-list*. Role required: admin, limited-admin.

```
# authentication ldap servers del 10.26.16.251:400
LDAP server(s) deleted.
LDAP Servers: 1
# Server
-----
1 10.26.16.250 (primary)
-----
```

```
authentication ldap servers reset
```

Reset the LDAP server configuration to an empty server list. Role required: admin, limited-admin.

```
# authentication ldap servers reset
LDAP server list reset to empty.
```

```
authentication ldap show
```

Display the LDAP configuration. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
# 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.

authentication

```
authentication ldap ssl disable
Disable SSL for LDAP. Role required: admin, limited-admin.
```

```
# authentication ldap ssl disable
Secure LDAP is disabled.
```

```
authentication ldap ssl enable [method {ldaps | start_tls}]
Enable SSL for LDAP, and set the SSL method as ldaps or start_tls. The default is
ldaps. Role required: admin, limited-admin.
```

Note

Run the `adminaccess certificate import` command to import the LDAP CA certificate before enabling LDAP SSL.

```
# authentication ldap ssl enable
Secure LDAP is enabled with 'ldaps' method.
```

```
authentication ldap ssl reset tls_reqcert
Reset the tls_reqcert option to the default value of demand. Role required: admin,
limited-admin.
```

Note

Resetting `tls_reqcert` to demand requires importing the LDAP CA certificate before enabling LDAP SSL. Run the `adminaccess certificate import` command to import the LDAP CA certificate.

```
# 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.
```

```
authentication ldap ssl set tls_reqcert {never | demand}
Specify the checks to perform on the server certificate. The default is demand. Role
required: admin, limited-admin.
```

Note

The default value for `tls_reqcert` is demand, which requires importing the LDAP CA certificate before enabling LDAP SSL. Run the `adminaccess certificate import` command to import the LDAP CA certificate.

```
# authentication ldap ssl set tls_reqcert never
"tls_reqcert" set to "never". LDAP server certificate will not
be verified.
```

```
authentication ldap status
Display the LDAP status. Role required: admin, limited-admin, security, user, backup-
operator, or none.
```

```
# authentication ldap status
Status: good (enabled)
```

authentication nis

```
authentication nis disable
```

Disable the NIS client. Role required: admin, limited-admin.

```
authentication nis domain reset
```

Reset the NIS domain name. Role required: admin, limited-admin.

```
authentication nis domain set domain [servers server-list]
```

Set the NIS domain name and optionally add NIS servers to the *server-list* by specifying the server hostnames. Role required: admin, limited-admin.

```
authentication nis domain show
```

Display the configured NIS domain name. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
authentication nis enable
```

Enable the NIS client. Role required: admin, limited-admin.

```
authentication nis groups add group-list role {user | admin |
backup-operator | limited-admin}
```

Add a role-based access control (RBAC) role for NIS users in the *group-list*. You cannot add an existing tenant-admin group or tenant-user group to an NIS group. See the *Data Domain Operating System Administration Guide* for role definitions. Role required: admin, limited-admin.

Example 23

```
# authentication nis groups add "tul_user group1" role admin
**** "tul_user group1" is currently an tenant-user group.
```

```
authentication nis groups del group-list role {user | admin |
backup-operator | limited-admin}
```

Delete a role-based access control (RBAC) role for NIS users in the *group-list*. See the *Data Domain Operating System Administration Guide* for role definitions. Role required: admin, limited-admin.

```
authentication nis groups reset
```

Delete all added NIS groups. Role required: admin, limited-admin.

```
authentication nis groups show [role {user | admin | backup-
operator | limited-admin}]
```

Display lists of NIS user groups and NIS admin groups. Role required: admin, limited-admin, security, user, backup-operator, or none.

Example 24

```
# authentication nis groups show
NIS Group   Role
-----
group1      user, tenant-admin
group2      user, tenant-user
-----
```

```
authentication nis reset
```

Delete the NIS configuration and set it to the default. Role required: admin, limited-admin.

```
authentication nis servers add server-list
```

Add NIS servers to the *server-list*. Role required: admin, limited-admin.

Note

If you add an invalid server name and enable NIS authentication, the system will continue to use the last valid NIS server name, based on the set domain.

`authentication nis servers del server-list`

Delete NIS servers from the *server-list*. Role required: admin, limited-admin.

`authentication nis servers reset`

Reset the NIS servers to their default settings. Role required: admin, limited-admin.

`authentication nis servers show`

Display a list of NIS servers. Role required: admin, limited-admin, security, user, backup-operator, or none.

`authentication nis show`

Display the NIS configuration. Role required: admin, limited-admin, security, user, backup-operator, or none.

`authentication nis status`

Display the NIS status. Role required: admin, limited-admin, security, user, backup-operator, or none.

CHAPTER 6

authorization

The `authorization` command, which is available only to security officers, establishes or modifies runtime authorization policy. Command options enable security-based functions such as managing filesystem encryption and enabling or disabling authorization policy.

All authorization tasks are logged automatically. The log file includes a timestamp, the identities of the security officer and administrative user, and the Data Domain system on which the task was performed. This log file serves as the audit trail, or “authorization history,” for each action.

This chapter contains the following topics:

- [authorization change history](#) 70
- [authorization guidelines and restrictions](#) 70
- [authorization policy](#) 70
- [authorization show](#) 70

authorization change history

There have been no changes to this command in this release.

authorization guidelines and restrictions

- Procedures requiring authorization must be dual-authenticated by the security officer and the user in the admin role. For example, to set encryption, the admin enables the feature and the security officer enables runtime authorization.

authorization policy

```
authorization policy reset security-officer
```

Reset runtime authorization policy to defaults. Resetting authorization policy is not allowed on Retention Lock Compliance systems. Role required: security.

```
authorization policy set security-officer {enabled | disabled}
```

Enable or disable runtime authorization policy. Disabling authorization policy is not allowed on Retention Lock Compliance systems. Role required: security.

Example 25

```
# authorization policy set security-officer enabled
```

```
authorization policy show
```

Show the current authorization policy configuration. Role required: security.

authorization show

```
authorization show history [last n { hours | days | weeks }]
```

View or audit past authorizations according to the interval specified. Role required: security.

CHAPTER 7

autosupport

The `autosupport` command manages system reports. Command options enable administrative users to manage two reports that describe the state of a Data Domain system: the `autosupport` report and the daily alert summary. By default, both reports are emailed to the Support address only, but users with admin role permissions may configure additional addresses and designate a subject tag keyword to bypass filtering that may block email delivery. For details on configuring `autosupport` notifications, see the *Data Domain Initial Configuration Guide*.

This chapter contains the following topics:

- [autosupport change history](#) 72
- [autosupport guidelines and restrictions](#) 72
- [autosupport add](#) 72
- [autosupport del](#) 72
- [autosupport reset](#) 73
- [autosupport send](#) 73
- [autosupport set](#) 74
- [autosupport show](#) 75
- [autosupport test](#) 75

autosupport change history

There have been no changes to this command in this release.

autosupport guidelines and restrictions

- Use the up and down arrow keys to move through the log. Use the q key to exit. Enter a forward slash and a pattern to search for dates.

autosupport add

```
autosupport add {alert-summary | asup-detailed} emails email-list
```

Add entries to the email list for the daily alert summary or the autosupport report. The system does not accept EMC or Data Domain email IDs. Role required: admin, limited-admin.

Example 26

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

Argument Definitions

alert-summary

Adds the specified emails to the list for daily alert distribution.

asup-detailed

Adds the specified emails to the list for daily autosupport report distribution.

email-list

Specifies the emails to be added to the specified list. Separate the list items with commas, spaces, or both.

autosupport del

```
autosupport del {alert-summary | asup-detailed} emails email-list
```

Delete entries from the email list for the daily alert summary or the autosupport report. The system does not delete EMC or Data Domain email IDs. Role required: admin, limited-admin.

Argument Definitions

alert-summary

Deletes the specified emails from the list for daily alert distribution.

asup-detailed

Deletes the specified emails from the list for daily autosupport report distribution.

email-list

Specifies the emails to be deleted from the specified list. Separate the list items with commas, spaces, or both.

autosupport reset

```
autosupport reset {alert-summary | asup-detailed}
```

Reset asup-detailed email list or alert-summary email list to the default value. Role required: admin, limited-admin.

```
autosupport reset all
```

Reset all autosupport command options to the default values. Output includes details on where autosupport reports and alert summaries are sent and the related schedules. Role required: admin, limited-admin.

```
autosupport reset schedule [alert-summary | asup-detailed]
```

Reset the schedules of the daily alert summary and the autosupport report to the default values.

- By default, the schedule for the daily alert summary is configured with the daily and 0600 options.
- By default, the schedule for the autosupport report is configured with the daily and 0800 options.

Role required: admin, limited-admin.

```
autosupport reset subject-tag
```

Clear the configured subject tag for the autosupport report and daily alert summary. Role required: admin, limited-admin.

autosupport send

```
autosupport send [email-addr] [cmd "cmd"]
```

Email an autosupport report or execute a command and send the output to the autosupport report email list or to the address specified. Role required: admin, limited-admin, security, user, backup-operator, or none.

Example 27

To run the net show stats command and email the results to djones@yourcompany.com:

```
# autosupport send djones@yourcompany.com cmd "net show stats"
```

Argument Definitions**“cmd”**

Run the specified DD OS command. Enclose the command in double quotation marks.

email-addr

Enter the email address to which you want to send the report.

autosupport set

```
autosupport set schedule {alert-summary | asup-detailed}
{[{daily | day(s)} time] | never}
```

Schedule the daily alert summary or the autosupport report. For either report, the most recently configured schedule overrides the previously configured schedule. Role required: admin, limited-admin.

Argument Definitions

alert-summary

Specifies that the schedule defined in the command is for the alert summary.

asup-detailed

Specifies that the schedule defined in the command is for the detailed autosupport report.

daily

Specifies that the selected report is sent daily.

day(s)

Specifies the days on which the report is sent. Valid entries are Mon, Tue, Wed, Thu, Fri, Sat, and Sun. Enter multiple days as needed and separate the days with spaces or commas.

time

Specifies the scheduled time in the HHMM format.

Example 28

To schedule the daily alert summary for 2 p.m. Monday and Friday:

```
# autosupport set schedule alert-summary mon,fri 1400
```

Example 29

To schedule the autosupport report for Tuesday at 4 a.m.:

```
# autosupport set schedule asup-detailed tue 0400
```

Example 30

To schedule the autosupport report for Tuesday at 3 p.m.:

```
# autosupport set schedule asup-detailed tue 1500
```

```
autosupport set subject-tag tag
```

Specify the text that is inserted in the subject line of autosupport report and daily alert summary emails. This allows the recipients to filter the emails based on subject. Maximum number of characters is 64. Role required: admin, limited-admin.

autosupport show

```
autosupport show {all | alert-summary | asup-detailed}
```

Displays the entire autosupport configuration when the `all` is specified. The `alert-summary` option lists the emails that receive the alert summaries, and the `asup-detailed` option lists the emails that receive detailed autosupport reports. Role required: `admin`, `limited-admin`, `security`, `user`, `backup-operator`, or `none`.

Example 31

```
# autosupport show all
The Admin email is:
Detailed autosupport and alert summary to Data Domain currently
enabled.
Detailed autosupport is scheduled to run "daily" at "0600".
Detailed autosupport is sent to:
  myemail1@abc.com
  myemail2@abc.com
  autosupport@autosupport.datadomain.com

Alert summary is scheduled to run "daily" at "0800".
Alert summary is sent to:
  myemail1@abc.com
  myemail2@abc.com
  autosupport@autosupport.datadomain.com
```

```
autosupport show history
```

Display the event history file, which includes the date for each autosupport report. Message system logs are retained for 10 weeks. Role required: `admin`, `limited-admin`, `security`, `user`, `backup-operator`, or `none`.

```
autosupport show report
```

Generate an autosupport report without sending the results to the autosupport report email list. Role required: `admin`, `limited-admin`, `security`, `user`, `backup-operator`, or `none`.

```
autosupport show schedule [alert-summary | asup-detailed]
```

Displays the email schedule for either the alert summary or the detailed autosupport report. Role required: `admin`, `limited-admin`, `security`, `user`, `backup-operator`, or `none`.

autosupport test

```
autosupport test {alert-summary | asup-detailed | support-
notify} | email email-addr
```

Send a test email to the email addresses in the specified list or to a specific email address. Role required: `admin`, `limited-admin`, `security`, `user`, `backup-operator`, or `none`.

Argument Definitions

alert-summary

Specifies the alert summary email list, which is created with the `autosupport add alert-summary` command.

asup-detailed

Specifies the detailed autosupport email list, which is created with the `autosupport add asup-detailed` command.

email-addr

Specifies an email address to which the test email is sent.

support-notify

Specifies the support notify email list, which is created with the `alerts notify-list add` command.

CHAPTER 8

boostfs

The `boostfs` command enables you perform operations using the Data Domain BoostFS software option. BoostFS provides a virtual filesystem that allows you to increase the number of backup applications that integrate with Data Domain systems. It can also help improve backup time, load balancing, and in-flight encryption. These capabilities allow BoostFS to reduce bandwidth usage.

Using the open-source software Filesystem in User Space (FUSE) and the DD Boost SDK, BoostFS allows you to mount your own filesystem implementations.

The `boostfs lockbox set` command allows you to store user credentials in RSA Lockbox. The `boostfs lockbox` includes arguments that let you specify the hostname of the Data Domain system, the storage-unit name, and the storage-unit user.

The `boostfs mount` command allows you to mount storage units on a Data Domain system as a filesystem and access files from the mounted filesystem. This command also allows you to add and manage credentials for access.

- [boostfs change history](#)..... 78
- [boostfs mount](#)..... 78
- [boostfs lockbox](#)..... 79
- [boostfs kerberos](#)..... 79

boostfs change history

There have been no changes to this command in this release.

boostfs mount

The `boostfs mount` command allows you to establish the BoostFS FUSE mount.

```
boostfs mount [-d|--data-domain-system] <data-domain-system>
[-s|--storage-unit] <storage-unit>
[[-o|--option <param>=<value>] ...] <mount-point>
```

Mount the BoostFS file system. Role required: none.

```
boostfs umount <mount-point>
```

Unmount the BoostFS file system. Role required: none.

Argument Definitions

mount-point

The mount-point for the BoostFS system.

storage-unit

The target storage-unit on the Data Domain system.

Command options

The following options are valid for the `boostfs mount` command:

Option	Description
<code>-o directory-name=path/to/subdir</code>	Subdirectory within the storage-unit you select for mounting (default: root of the storage unit). You must create the subdirectory after mounting at the root path, unmounting, and adding the parameter to the subsequent mount command or config file.
<code>-o security=<krb5 lockbox></code>	Security option used for authentication (default: lockbox)
<code>-o allow-others=<true false></code>	Allow users other than the owner of the mount to access the mount
<code>-o log-enabled=<true false></code>	Enable logging (default: true)
<code>-o log-level=<debug info warning error></code>	Log level (default: info)
<code>-o log-dir=/path/to/log</code>	Directory for log files (default: <code>/opt/emc/boostfs/log</code>)
<code>-o log-file=output.log</code>	Log file name (default: <code>/opt/emc/boostfs/log/ddboostfs_<uid>_<gid>.log</code>)
<code>-o log-maxsize=100</code>	Maximum log size in MB (default: 100 MB)
<code>-o log-rotate-num=8</code>	Number of log files to save (default: 8)

Option	Description
<code>-o app-info="text string"</code>	Text string describing the application using boostfs (version, etc) (default: FUSE version)

boostfs lockbox

The `boostfs lockbox` command allows you to set the RSA lockbox values.

```
boostfs lockbox add-hosts hostname [hostname]
```

Adds clients that can access the shared lockbox. When you are adding and removing access to the shared lockbox, you must do so from the machine where the lockbox was initially created. Role required: admin.

```
boostfs lockbox delete-hosts all
```

Deletes all client access to the shared lockbox. Role required: admin.

```
boostfs lockbox delete-hosts hostname [hostname]
```

Deletes specific client access to the shared lockbox. Role required: admin.

```
boostfs lockbox {remove | query} [-d | --data-domain-system]
data-domain-system
[-s | --storage-unit] storage-unit-name
```

If the credentials have been stored in an RSA lockbox, this command returns the username after the query is submitted with the specified Data Domain hostname and storage-unit. Role required: admin.

```
boostfs lockbox set [-d | --data-domain-system] data-domain-
system
```

```
[-u | --storage-unit-username] storage-unit-username
[-s | --storage-unit] storage-unit-name
```

To store credentials in an RSA lockbox, the user specifies the Data Domain hostname, the storage-unit name, and the storage-unit user. After providing that information, the user is prompted for the password. Role required: admin.

Note

The command `boostfs lockbox set` fails if there is an existing Lockbox file in the same location. This includes Lockbox files generated with older versions of BoostFS. For example, the existence of a BoostFS 1.1 Lockbox causes the creation of a Lockbox with BoostFS 1.2 to fail.

```
boostfs lockbox remove [-d | --data-domain-system] data-domain-
system
[-s | --storage-unit] storage-unit-name
```

Removes the stored RSA lockbox credentials in the specified Data Domain system and storage-unit. Role required: admin.

```
boostfs lockbox show-hosts
```

Shows all clients that can access the shared lockbox. Role required: admin.

boostfs kerberos

The `boostfs kerberos` command allows you to add, verify, and remove Kerberos credentials.

```
boostfs kerberos set [-u | --storage-unit-username] <storage-
unit-username>
```

```
[-s |--storage-unit] <storage-unit-name>  
[[-r |--kerberos-realm] <kerberos-realm>]  
[[-m | --kerberos-username] <kerberos-username>]
```

Allows you to add Kerberos credentials. Role required: admin.

```
boostfs kerberos query [[-u |--storage-unit-username] <storage-  
unit-username>
```

```
[-s |--storage-unit] <storage-unit-name>]  
[[-m |--kerberos-username] <kerberos-username>]
```

Checks for Kerberos credentials. Role required: admin.

```
boostfs kerberos remove [-u | --storage-unit-username]  
<storage-unit-name>
```

Removes Kerberos credentials. Role required: admin.

CHAPTER 9

cifs

The `cifs` command manages CIFS data access between a Data Domain system and Windows clients. Command options enable and disable access to a Data Domain system from media servers and other Windows clients that use the CIFS protocol. The `cifs` command sets the authentication mode, share management, and administrative access, and displays status and statistics for CIFS clients.

This chapter contains the following topics:

- [cifs change history](#) 82
- [cifs disable](#)..... 82
- [cifs enable](#)..... 82
- [cifs local-group](#)..... 82
- [cifs option](#)..... 83
- [cifs reset](#)..... 83
- [cifs restart](#)..... 84
- [cifs set](#)..... 84
- [cifs share](#)..... 85
- [cifs show](#)..... 86
- [cifs status](#)..... 86
- [cifs troubleshooting](#)..... 86

cifs change history

Modified arguments and modified output in DD OS 6.1.2

`cifs option reset name`

The `support-smb1` and `support-smb2` options provide the ability to enable and disable support for SMB version 1 and SMB version 2. The default value for both is enabled.

`cifs option set name value`

The `support-smb1` and `support-smb2` options provide the ability to enable and disable support for SMB version 1 and SMB version 2. The default value for both is enabled.

`cifs option show [current | all]`

Displays the values for the `support-smb1` and `support-smb2` options.

cifs disable

```
cifs disable
```

The CIFS server stops listening on port 445. Role required: admin, limited-admin.

cifs enable

```
cifs enable
```

The CIFS server starts listening on port 445. Role required: admin, limited-admin.

cifs local-group

```
cifs local-group add group-name members member-list
```

Add a domain user or domain group to the cifs local group using a comma separated list. Role required: admin, limited-admin.

Note

Do not use `cifs local-group add` when an `f5` option has already been set using the `cifs option set f5` command.

```
cifs local-group del group-name members {all | member-list}
```

Delete a domain user or domain group from the cifs local group using the word all or a comma separated list. Role required: admin, limited-admin.

```
cifs local-group show list [group-name]
```

Display brief information about the cifs local group, for example: group name and number of members present in this group. Role required: admin, limited-admin.

```
cifs local-group show detailed [group-name]
```

Display detailed information about the cifs local group, for example: group name, group SID (security identifier), and group ID as well as the group member names and their SIDs. Role required: admin, limited-admin.

cifs option

```
cifs option reset name
```

Reset a CIFS option to default value. Name field will support group name "dd limited-admin group1 " to " dd limited-admin group50." Role required: admin, limited-admin, user.

```
cifs option reset f5 name
```

Reset a CIFS option to default value. For use with Data Domain Systems using F5 Network's tiered storage solution ARX. Role required: admin, limited-admin.

```
cifs option set name value
```

Validates and sets the option name and value only if both are within the supported range. Name field will support group name "dd limited-admin group1 " to " dd limited-admin group50." Role required: admin, limited-admin, user.

```
cifs option set f5 name value
```

Set a CIFS option. For use with Data Domain Systems using F5 Network's tiered storage solution ARX. Role required: admin, limited-admin.

Note

Do not use `cifs option set f5` when a CIFS local group has been set using the `cifs local-group add` command.

```
cifs option show [current | all]
```

Display CIFS options. Role required: admin, limited-admin.

Argument definitions

current

Display only currently set options. Default is current if option not specified.

all

Display all options except the undocumented options that are not set by user.

```
cifs option reset server-signing
```

Resets server signing to disabled, which is the default. Role required: admin, limited-admin.

```
cifs option set server-signing [enabled | disabled | mandatory]
```

Server Message Block (SMB) signing is a security mechanism that improves the security of the SMB protocol. When enabled using the auto option, it is possible for clients that support SMB signing to connect, although it is also possible for clients that do not support SMB signing to connect. When SMB signing is enabled using the mandatory option, both computers in the SMB connection must support SMB signing, and the SMB connection will not be successful if one computer does not support SMB signing. Role required: admin, limited-admin.

cifs reset

```
cifs reset authentication
```

Reset the CIFS authentication to the default: workgroup. Role required: admin, limited-admin.

```
cifs reset nb-hostname
```

Reset the NetBIOS hostname to the default: none. Role required: admin, limited-admin.

```
cifs reset stats
```

Reset cifs statistics. Role required: admin, limited-admin.

cifs restart

```
cifs restart [force]
```

Restart all likewise services. Role required: admin, limited-admin.

Argument definitions

force

Forces the system to restart likewise services.

cifs set

```
cifs set authentication active-directory realm { [dc1  
[dc2 ...]] | * }
```

Set authentication to Active Directory (AD). The realm must be a fully qualified name. Use commas, spaces, or both to separate entries in the domain controller list. Security officer authorization is required for systems with Retention Lock Compliance enabled. Role required: admin, limited-admin.

Note

Data Domain recommends using the asterisk to set all controllers instead of entering them individually.

Argument definitions

realm

The Windows Active Directory realm.

dc1, dc2

Domain Controller 1, Domain Controller 2. You can use the * wildcard character in the string.

When prompted, enter a name for a user account. The type and format of the name depend on if the user is inside or outside the company domain.

- For user “Administrator” inside the company domain, enter the name only: administrator.
- For user “Jane Doe” in a trusted domain, enter the user name and domain: jane.doe@trusteddomain.com. The account in the trusted domain must have permission to join the Data Domain system to your company domain.

The Data Domain system automatically adds a host entry to the DNS server. It is not necessary to create the entry manually.

If you set the NetBIOS hostname using the command `cifs set nb-hostname`, the entry is created for NetBIOS hostname only, not the system hostname. Otherwise, the system hostname is used.

```
cifs set authentication workgroup workgroup
```

Set the authentication mode to workgroup for the specified workgroup name. Role required: admin, limited-admin.

Argument definitions**workgroup**

Workgroup name.

```
cifs set nb-hostname nb-hostname
```

Set the NetBIOS hostname. Role required: admin, limited-admin.

cifs share

```
cifs share create share path path {max-connections max
connections | clients clients | users users | comment comment}
```

Create a new share. Role required: admin, limited-admin.

Note

This command accepts the `/backup` alias for the default (backup) MTree in addition to `/data/coll/backup`. For paths in all other MTrees, use `/data/coll/mtree-name`.

Argument Definitions**share**

A descriptive name for the share.

path

The path to the target directory.

max-connections

The maximum number of connections to the share allowed at one time.

clients

A comma-separated list of clients allowed to access the share. Specify the clients by hostname or IP address. No spaces or tabs are allowed and the list must be enclosed in double quotes. If the clients argument is not specified when creating the share, the share is not accessible by any client. To make the share accessible for all clients, enter the clients argument and precede client name by an ampersand.

users

A comma-separated list of user names. Other than the comma delimiter, spaces (blank or tab) are treated as part of the user name because a Windows user name can have a space in the name.

The user names list can include group names. Group names must be preceded by the symbol for the word *at* (`@`).

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. Separate group and user names by commas only. Spaces may be included within a group name but are not allowed as delimiters for group names.

comment

A descriptive comment about the share.

```
cifs share destroy share
```

Delete a share. Role required: admin, limited-admin.

```
cifs share disable share
```

Disable a share. Role required: admin, limited-admin.

```
cifs share enable share
```

Enable a share. Role required: admin, limited-admin.

```
cifs share modify share {max-connections max connections |
clients clients | users users | comment comment}
```

Modify a share configuration with the same configuration options as the cifs share create option, except for its path. You cannot change the path for an existing share. Modifications apply to new connections only. Role required: admin, limited-admin.

See the share create command option for a description of the command variables. To remove a user list for the share, specify *users*.

```
cifs share show [share]
```

Display share configurations for all shares, or for a specified or custom share, as well as shared access control lists. Role required: admin, limited-admin, user, backup-operator, security, none.

cifs show

```
cifs show active
```

Display all active CIFS clients. The system displays the computer, the user, opens, connection time, and idle time for each session. The system also displays the user, mode, locks, and file for each open file. A summary of sessions and open files is also displayed. Role required: admin, limited-admin, user, backup-operator, security, none.

```
cifs show config
```

Displays the CIFS configuration and whether the DDR is in workgroup mode or active directory mode as well as the maximum open files. Role required: admin, limited-admin, user, backup-operator, security, none.

Note

In the command output, "Max open files per connection" displays the maximum number of open files on a Data Domain system, not the number of open files per connection.

```
cifs show detailed-stats
```

Display detailed statistics on CIFS activity and performance. Role required: admin, limited-admin, user, backup-operator, security, none.

```
cifs show stats
```

Display basic statistics on CIFS activity and performance. Role required: admin, limited-admin, user, backup-operator, security, none.

cifs status

```
cifs status
```

Show status of CIFS: enabled or disabled. Role required: admin, limited-admin, user, backup-operator, security, none.

cifs troubleshooting

```
cifs troubleshooting domaininfo
```

Report domain information; for example, to check the connectivity between the Data Domain system and the domain. Also to confirm if authentication issues are due to domain connectivity. Role required: admin, limited-admin.

```
cifs troubleshooting group groupname | gid | SID
```

List details for a specified group. Role required: admin, limited-admin.

```
cifs troubleshooting list-groups
```

List all CIFS groups. Role required: admin, limited-admin.

```
cifs troubleshooting list-users
```

List all CIFS users. Role required: admin, limited-admin.

```
cifs troubleshooting performance
```

Collect tcpdump and ddfs traces for CIFS performance analysis. Role required: admin, limited-admin.

Example 32

To troubleshoot performance problems:

```
Enter: cifs troubleshooting performance
```

```
Enter: support bundle upload
```

```
cifs troubleshooting user username | uid | SID
```

Display details on a specified user.

CHAPTER 10

client-group

The `client-group` command lets you configure and monitor external clients in groups, independent of protocol used.

Note

The Client Group feature is supported only for clients that use DD Boost, NFS, VTL, or CIFS protocols.

Client Group monitoring provides stream counting, stream limit checks, and access checks against an allowed MTree list. Client Group displays show read and write stream counters, incoming and outgoing byte counts and network usage, and active stream file activity with duration and image name. Client Group history provides the history of every image written or read as well as historical statistics for each group, collected at a specified interval.

Client Group supports both IP and FibreChannel transport types.

This chapter contains the following topics:

• client-group change history	90
• client-group guidelines and restrictions	90
• client-group add	91
• client-group compression show	91
• client-group create	92
• client-group data-access-permit-list	92
• client-group del	93
• client-group destroy	93
• client-group rename	93
• client-group show	93
• client-group stats options	102
• client-group stream-limit	104

client-group change history

New commands in DD OS 6.1.2

```
client-group show host-stats {all | hostname} view {current-hour | last-24-hours | daily-total} [display-unit {Bytes | KiB | MiB | GiB}]
```

Displays data sent and received by a specified host for the last hour, the last 24 hours, or a total daily summary. The system stores 24 hours worth of statistics per host for average daily calculations. Daily statistics tracking for a host starts with the first read/write stream processed by that host, and ends when the host is inactive for 24 hours.

```
client-group show file-history {all | group-name} [last n {mins | hours | days}] [end MMDDhhmm[[CC]YY]] [count n] [host hostname] [mtree mtree-name] [min-size bytes]
```

Display the file history of a specified client group, or all client groups. Filtering options include specifying a time period, a host, an MTree, or a minimum file size.

```
client-group show performance {group-name} {start MMDDhhmm[[CC]YY]} [end MMDDhhmm[[CC]YY]] [display-unit {Bytes | KiB | MiB | GiB}]
```

Display the performance history for a specific client-group, or all client groups in 10 minute intervals within the specified time range.

```
client-group show stream-history {all | group-name} view {log-interval | performance-interval} [last n {mins | hours | days}] [end MMDDhhmm[[CC]YY]] [count n]
```

Display the stream usage by a specific group, or all groups, for a specified time period.

Modified arguments in DD OS 6.1.2

```
client-group compression show accumulated-history {all | group-name} view {log-interval | performance-interval | reset-interval} [end MMDDhhmm[[CC]YY]] [last n {mins | hours | days}] [count n] [display-unit {Bytes | KiB | MiB | GiB}]
```

The `performance-interval` option displays statistics for 10 minute intervals until the specified end time. The `[display-unit {Bytes | KiB | MiB | GiB}]` option provides different measurement options for the output values.

client-group guidelines and restrictions

- A group of hosts is added to a client group in order to monitor and control the group as an entity. The monitoring captures the read/write bytes and concurrent active streams per group. This monitoring information is periodically written to `clients_stats.log`.
- The control takes the form of stream limit checks and permission validation against the allowed MTree (storage-unit).
- For NFS streams, near-line access is assumed when write or read requests are sent with block sizes of less than 4K bytes.

- When Client Group is used with BoostFS with Kerberos authentication, the hostname field in output screens and log files shows the client's IP address, not the hostname.

client-group add

```
client-group add group-name host host-list
```

Add a host or a list of hosts to a client group. A client group must have been created before you can add a host to a client group. To create a client group, use the `client-group create` command.

Note

Host names must be entered using all lowercase characters. Using uppercase characters results in the client being put in the "unassigned" group.

Note

For VTL, the host name is the VTL pool name.

Hosts do not move between groups while they have active streams. A host can only be in one group at any given time, so a host finds its group on its first stream activity. Once a host starts a read/write operation, it remains in the group it started on, independent of any add commands. The add for host only takes effect when the read/write starts. If additional streams are active concurrently, they will all belong to the same group. When you move a host to another group, the host must not be active with any streams, or it will not start in the next group.

The search priority for the *host-list* in identifying a *group-name* is:

1. IP address of the client. The IP address is entered as a range--for example, host 10.30.2.28/32
 - For IPv4, you can select five different range masks, based on network.
 - For IPv6, fixed masks /64 /112 and /128 are available.
2. Hostname. The hostname must match the client-sent name in protocol communication. This can be verified with `ddbboost show connections` for DD Boost or `nfs show active` for NFS.
3. Partial FQDN.

The first match is used.

The "unassigned" group, which is the default group for all clients not added to a specific group, does not allow *host-list* (that is, the command checks and does not allow specification of *host-list*).

client-group compression show

```
client-group compression show accumulated [group-name] [view {both | log-interval | reset-interval}]
```

View accumulated compression statistics and deduplication ratio for groups of clients to detect and isolate poor deduplication.

```
client-group compression show accumulated-history {all | group-name} view {log-interval | performance-interval | reset-interval} [end MMDDhhmm[[CC]YY]] [last n {mins | hours | days}] [count n] [display-unit {Bytes | KiB | MiB | GiB}]
```

View historical group compression statistics for groups of clients over multiple log or reset intervals to detect and isolate poor deduplication.

The value for `count` must be less than or equal to 100.

```
client-group compression show detail-history {all | group-name}
[host hostname] [mtree mtree-name] [min-size bytes] [end
MMDDhhmm[[CC]YY]] [last n {min | hours | days}] [count n]
```

View detailed history per file, including compression statistics and deduplication ratio for specific groups of clients.

The value for `min-size` must be less than 1073741823.

The value for `count` must be less than or equal to 100.

The `performance-interval` option is only supported when a single client-group name is specified.

client-group create

```
client-group create group-name
```

Create a client group.

You can create a maximum of 64 groups; the "unassigned" group is created automatically. Clients are mapped to the "unassigned" group until configured for another group. The `group-name` cannot exceed 24 alphanumeric characters; the hyphen (-) and underscore (_) characters are also allowed. The group-name cannot be any of the following: all, no, none, limit, group, client, view.

client-group data-access-permit-list

```
client-group data-access-permit-list add group-name mtree
mtree-name-list
```

Add a list of MTrees (`/data/coll/mtree-name` OR `mtree-name`) or storage units that should be validated at the start of a read/write operation. Each client group must check against the allowed MTree list. No checking is done when there is no configured MTree or storage unit. If the access check against any specified MTrees or storage unit fails, a permission error is returned. The application then deletes the empty file and lets you correct the MTree used in the policy.

If a group does not have any data-access-permit settings, permission checks will not be performed. The check is only performed one time at the start of a read or write stream operation.

The command accepts a list of MTree names or a single MTree name.

When a storage-unit is deleted or renamed, the client-group MTree is removed or renamed accordingly.

The MTree is not exclusive to a client-group--the same MTree or storage-unit can be specified for multiple client-groups.

The command allows both the full `mtree /data/coll/mtree-name` and the short version `mtree-name`, but all displays show only the short version `mtree-name`.

Note

This command is not supported for VTL.

client-group del

```
client-group del group-name host host-list
```

Delete a host from a client group.

Note

For VTL, the host name is the VTL pool name.

You can delete a host-list from a client group at any time, but the change does not take effect until the all active read/write jobs complete on the group-name started.

When a stream starts, using the priority search order described in `client group add`, the host finds its client-group. A host will not check its client-group association until all its current streams complete and it starts a new read or write stream.

client-group destroy

```
client-group destroy group-name
```

Destroy an empty client group.

Remove all configured hosts from the client group before issuing the `destroy` command. The group must not have active client streams. The "unassigned" group cannot be destroyed.

When a client-group is destroyed, all statistics information for this client-group is cleared, except for the information in the `clients_stats.log` file.

client-group rename

```
client-group rename group-name new-group-name
```

Rename a client group.

You can rename a client group at any time. The *group-name* cannot exceed 24 alphanumeric characters; the hyphen (-) and underscore (_) characters are also allowed. The group-name cannot be any of the following: all, no, none, limit, group, client, view.

The client-group name is also used by both log files, `clients_history.log` and `clients_stats.log`. The new name will be used as new log entries are written.

client-group show

```
client-group show config [ group-name ] [view {summary | hosts  
| stream-limit | mtree}]
```

Show configuration of all client groups, a particular group or selected configuration.

summary

- View summary is part of ASUP
- Shows Boost, NFS, CIFS, and/or VTL start time as applicable
- Shows all the statistics options settings
- Shows the stream limits, number of hosts, and Mtree configured

hosts

- List of every host and its associated group
- Can list only hosts for a particular group
- Host configured may be subnet or a domain

stream-limit**mtree**

- List of every group MTree or storage-unit
- Select a specific group for MTree or storage-unit

The first protocol read or write of a stream will set the protocol started time. This time is refreshed when the file-system restarts such as on an upgrade or reboot process.

Note

ASUPs will have the output of `client-group show config view summary`, which includes the protocol started time.

Example 33**# client-group show config**

```
Group-name  Client      Group
           Stream   Stream  Hosts  MTree
           Limit   Limit   Count  Count
-----
unassigned      6         18     -      0
group2          8         10     3      2
group3          0          0     1      0
-----
```

```
Group-name      Host-name
-----
group2          ddbboost-dl.datadomain.com
group2          10.4.5.0/24
group3          *.emc.com
-----
```

```
Group-name      MTree-name
-----
group2          DDBOOST_STRESS
group2          STU3
-----
```

```
Stats options: log-interval = 60, log-condition = updated, reset-
interval = 60
```

```
Boost started: 2016/05/01 20:48:03
NFS started:   2016/05/01 18:19:08
```

Example 34**# client-group show config view summary**

```
Group-name  Client      Group
           Stream   Stream  Hosts  MTree
           Limit   Limit   Count  Count
```

Example 34 (continued)

```

-----
unassigned      6      18      -      0
group2          8      10      3      2
group3          0      0       1      0
-----
Stats options: log-interval = 60, log-condition = updated, reset-
interval = 60

Boost started: 2016/05/01 20:48:03
NFS started:   2016/05/01 18:19:08

```

```

client-group show file-history {all | group-name} [last n {mins
| hours | days}] [end MMDDhhmm[[CC]YY]] [count n] [host
hostname] [mtree mtree-name] [min-size bytes]

```

Display the file history of a specified client group, or all client groups. Filtering options include specifying a time period, a host, an MTree, or a minimum file size.

Example 35

```

# sysadmin@rtp-ost-arch2# client-group show files-history all host
10.6.109.203 last 2 hours min-size 4096

```

Operation	Client Path	Client	Server	Average
Start	Last	Duration	End	
Group-Name	Hostname	Interface		
Interface	Offset	Offset	hh:mm:ss	Timestamp
Throughput	(MB/s)			
BLUE_GROUP	10.6.109.203	10.25.25.133	10.6.109.203	cifs-
write	BLUE_CIFS/bluemedi	1488831828	C1_F1.1488831828	.img
0	9,037,194,240	01:05:45	2017/03/06 16:26:56	2.18

```

client-group show files [ group-name | host hostname | mtree
mtree-name ]

```

List active stream files.

All the information displayed for this command is placed in `clients_history.log` when a file completes. The `clients_history.log` defaults to rotate at 100M and keeps nine zip files.

The hostname is as reported from the host, so DD Boost and NFS send the FQDN.

Note

The `/data/col1/` is stripped out of the path display.

Group-Name

Client group-name with active streams

Client Hostname

FQDN or short-name sent by client

Client Interface

IP address of Client (or FC indication)

Server Interface

IP of DD side (or FC indication)

Operation

- dsp-write or syn-dsp-write – Boost DSP write with synthetic if specified
- bst-write or nfs-write – non DSP write for either DD Boost or NFS protocol
- bst-read or nfs-read – Read for either DD Boost or NFS protocol
- cifs-write – Write for CIFS protocol
- cifs-read – Read for CIFS protocol
- vtl-write – Write for VTL protocol
- vtl-read – Read for VTL protocol
- vtl-rd-w – Read or write for VTL protocol

Path

Full path not including “/data/col1”

Start Offset

Start offset written or read

Last Offset

Last offset written or read

Duration

Duration active in read/write state

Example 36

client-group show files

Start	Client	Client	Server	Operation	Path	
Group-Name	Last	Duration	Interface	Interface		Offset
Offset	hh:mm:ss					
group5	bst.emc.com	192.168.1.203	192.168.1.98	dsp-write	STU/wr_0	0
515,899,392	00:00:02					
group5	bst.emc.com	192.168.1.203	192.168.1.98	bst-read	STU/wr_0	0
515,899,392	00:01:03					
group2	ddbboost-dl	10.6.109.177	10.26.16.182	nfs-write	MT1/data_1	0
5,287,707,000	00:00:51					
group2	ddbboost-dl	10.6.109.177	10.26.16.182	nfs-read	MT1/data_1	0
5,287,707,000	00:00:51					

Example 37 Log view

Example 37 Log view (continued)

```
# log view debug/clients_history.log
HEADER_V1 < date > < time > group=< groupname > host=< FQDN > cl-if=< IP | FC > dd-if=< IP |
FC > op=< read/write> path=< mtree-name/file > s_off=< bytes > l_off=< bytes > dur=<
hh:mm:ss >
2015/08/28 17:41:09 group=group5 host=ddboost-dl.datadomain.com cl-if=192.168.1.203 dd-
if=192.168.1.98 op=dsp-write path=DDBOOST_STRESS_SU/write_000 s_off=0 l_off=5368709120
dur=00:00:22
2017/04/04 07:13:32 group=High_DP_G host=wangy34-dl.datadomain.com cl-if=128.222.90.222 dd-
if=10.25.17.53 op=nfs-write path=backup/test.0000.0000 s_off=0 l_off=107374182584
size=107374182584 orig=107734373224 g_comp=105588162404 l_comp=48956463624 dur=00:17:38
```

```
client-group show host-stats {all | hostname} view {current-
hour | last-24-hours | daily-total} [display-unit {Bytes | KiB
| MiB | GiB}]
```

Displays data sent and received by a specified host for the last hour, the last 24 hours, or a total daily summary. The system stores 24 hours worth of statistics per host for average daily calculations. Daily statistics tracking for a host starts with the first read/write stream processed by that host, and ends when the host is inactive for 24 hours.

If a display unit is not specified, the output shows bytes by default.

All the information displayed by this command is stored in memory, and is overwritten after 24 hours.

Write

Amount of incoming data

Filtered

Boost DSP mode duplicate data

Post_Ic

Boost DSP mode compressed data

Read

Amount of outgoing data

Network

Amount of data on the network (IP or FC)

Pre-Comp

Data written before compression.

Post-Comp

Storage used after compression.

Global-Comp

Amount of data compressed.

Write Stream

Number of write streams, displayed as the number active or an average per hour.

Write Stream Max

Maximum number of write streams, displayed as the number active or an average per hour.

Read Stream

Number of read streams, displayed as the number active or an average per hour.

Read Stream Max

Maximum number of read streams, displayed as the number active or an average per hour.

Client Stream Reject

Number of client streams rejected by the system.

Example 38

```
# client-group show host-stats all view total-24-hours display-unit
GiB

Hostname      Timestamp      Write      Filtered      Post_lc
Read          Network      Pre      Global      Post      Write
Write         Read        Read      Client
(Boost)
Stream        Stream        Stream    Stream    Stream    (Boost)
GiB           GiB          GiB       GiB       GiB       Comp      Comp
hrs           Max          Avg/hrs   Max       Reject    Bytes    Avg/
-----
-----
-----
h1.emc.com    2017-05-02 06:00:00      84.14      84.14
19.80        0.00        19.80     45.33     31.07     14.43
3/4          6           2/1       6         0
h2.emc.com    2017-05-02 06:00:00      51.32      51.32
16.29        0.00        16.29     76.36     33.73     15.70
10/5         10          0/0       1         0
-----
-----
-----
```

```
client-group show performance {group-name} {start
MMDDhhmm[[CC]YY]} [end MMDDhhmm[[CC]YY]] [display-unit {Bytes |
KiB | MiB | GiB}]
```

Display the performance history for a specific client-group, or all client groups in 10 minute intervals within the specified time range.

Filtered

Boost DSP mode duplicate data

Post_lc

Boost DSP mode compressed data

Read

Amount of outgoing data

Network

Amount of data on the network (IP or FC)

Pre-Comp

Data written before compression.

Post-Comp

Storage used after compression.

Global-Comp

Amount of data compressed.

Client Steam Max

Maximum stream count a single client reaches in the log-interval.

Client Stream Limit

Stream limit that may be configured against each client in the client-group.

Client Stream Reject

Number of streams rejected due to the client exceeding the stream limit.

Group Stream Active

Number of active streams at the time the log history is written.

Group Stream Max

Maximum stream count the client-group reaches in the log-interval.

Group Stream Limit

Stream limit that may be configured against the entire client-group.

Group Stream Reject

Number of streams rejected due to the client exceeding the group limit.

Mtree Stream Reject

Number of streams rejected beyond a DD Boost storage stream limit setting.

```
client-group show stats [ group-name ] [interval seconds ]
```

Shows data sent and received for a specified interval for each client group.

All the information displayed for this command is periodically written to clients_stats.log. The clients_stats.log defaults to rotate at 100M and keeps nine zip files.

Group-Name

Client group-name

Write (GiB)

byte count of incoming data

Filtered (GiB)

Boost DSP mode duplicate data

Post_lc (GiB)

Boost DSP mode compressed data

Read (GiB)

byte count of outgoing data

Network (GiB)

byte count on the network (IP or FC)

Client stream Max-count

Maximum seen for a client

Group Stream Max-count

Maximum seen for group

Recent Stream Rejected

Rejected due to limits exceeded

Client Stream Limit

Configured client stream limit

Group Stream Limit

Configured group stream limit

Group Stream active

Active running streams on group

Limit checks include client limit, group limit, and storage-unit limit (if storage-unit). Limit violations cause the job to fail, and the exceeded count in the log is incremented.

These statistics are reset with the log interval (default 1 hour):

- Recent stream rejected (count of rejected streams due to limit violations)
- Stream max-count per client in group within interval
- Stream max-count for group within interval

The data statistics are reset based on the reset interval (default 12 hours)

Example 39

```
#client-group show stats
```

Group	Write Group	Filtered Group	Post_lc Recent	Read	Network	Client Stream	Client Stream
Group-Name	GiB	GiB	GiB	GiB	GiB	Max-count	Limit
Stream	Stream	Stream	Stream				
Active	Max-count	Limit	Rejected				
unassigned	0.00	0.00	0.00	0.00	0.00	0	
6	0	0	18	0			
group3	0.00	0.00	0.00	0.00	0.00	0	
5	0	0	20	0			
group2	9.32	3.81	0.60	4.92	10.44	1	
0	1	1	0	0			

Example 40

```
# client-group show stats dpstest interval 2
02/01 10:59:09
-----Written MiB/s----- Network MiB/s Read MiB/s -- Stream Counters in Interval -- -
Stream Limits-
Data Filtered Post_lc In/Out Data Max-clnt Max-grp Curr-grp Reject
Client Group
-----
```

Client	Data	Filtered	Post_lc	In/Out	Data	Max-clnt	Max-grp	Curr-grp	Reject
5	9.38	0.00	0.00	9.38	0.00	1	1	1	0
5	5.50	0.00	0.00	5.50	0.00	1	1	1	0
5	12.00	0.00	0.00	12.00	0.00	1	1	1	0
5	15.00	0.00	0.00	15.00	0.00	1	1	1	0
5	5.50	0.00	0.00	5.50	0.00	1	1	1	0

Example 41 Log view

```
# log view debug/clients_stats.log
HEADER_V3 {date time} group_name= write= filtered= post_lc= read= network= max-client-stream=
max-client-limit= active-group-stream= max-group-stream= max-group-limit=
exceed_client_limit= exceed_group_limit= exceed_mtree_limit
2015/12/02 11:51:25 group=group2 write=20789854584 filtered=14530687283 post_lc=986809934
read=0 network=11554606539 max-client-stream=1 max-client-limit=0
active-group-stream=1 max-group-stream=1 max-group-limit=0 exceed_client_limit=0
exceed_group_limit=0 exceed_mtree_limit=0
```

```
client-group show stream-history {all | group-name} view {log-
interval | performance-interval} [last n {mins | hours | days}]
[end MMDDhhmm[[CC]YY]] [count n]
```

Display the stream usage by a specific group, or all groups, for a specified time period.

Client Steam Max

Maximum stream count a single client reaches in the log-interval.

Client Stream Limit

Stream limit that may be configured against each client in the client-group.

Client Stream Reject

Number of streams rejected due to the client exceeding the stream limit.

Group Stream Active

Number of active streams at the time the log history is written.

Group Stream Max

Maximum stream count the client-group reaches in the log-interval.

Group Stream Limit

Stream limit that may be configured against the entire client-group.

Group Stream Reject

Number of streams rejected due to the client exceeding the group limit.

Mtree Stream Reject

Number of streams rejected beyond a DD Boost storage stream limit setting.

Example 42

```
# client-group show stream-history unassigned view performance-
interval end 04242000 last 1 hours count 2
```

Group-Name	End	Duration	Client	Client
Client	Group	Group	Group	MTree
Stream	Stream	Stream	Stream	Stream
Reject	Active	Max	Limit	Reject
Reject	Active	Max	Limit	Reject
unassigned	2017-04-24 19:10:00	00:10:00	0	1
0	0	0	0	0
unassigned	2017-04-24 19:40:00	00:10:00	0	19
0	18	19	0	0

Reached specified count.

Example 42 (continued)

`client-group show streams [group-name | host hostname]`
 List streams for each active group. The Stream limits per group and per client are also shown.

Group-Name

Client group-name with active streams

Client Hostname

FQDN or short-name sent by client

DSP-Write Stream Count

DD Boost only write streams

Data-Write Stream Count

Boost Write or NFS Write

Meta-Read Stream Count

Synthetic Write operation, the read stream

Data-Read Stream Count

Boost or NFS Read

Example 43

```
# client-group show streams
```

Client	Client Group	DSP-Write Stream	Data-Write Stream	Meta-Read Stream	Data-Read Stream
Group-Name	Hostname	Count	Count	Count	Count
Limit	Limit				
group2	ddboost-dl.datadomain.com	2	0	2	
2	0 0				

Example 44

```
# client-group show streams
```

```
** Failed to get list: No active clients in list
   DD_ERR_EMPTY = 5006
```

client-group stats options

`client-group stats options set [log-interval mins] [log-condition {updated | always}] [reset-interval mins]`

Configure optional settings for logging.

- These settings are for all groups. The `reset-interval` needs to be a multiple of the log interval.

- These options control the logging to `debug/clients_stats.log`.
- You can view the log using `log view debug/clients_stats.log`.
- Once over 24 hours of log data is collected, this header is put into the log:
`HEADER_V3 {date time} group= write= filtered= post_lc= read=
network= max-client-stream= max-client-limit= active-group-
stream= max-group-stream= max-group-limit=
exceed_client_limit= exceed_group_limit= exceed_mtree_limit=`

log-interval

- `Debug/clients_stats.log` file update interval
- Default value is 60 min.
- When statistics are written to file, these per-group counters are reset:
 - Client stream Max-count: maximum stream count seen for a client in group during log period
 - Group Stream Max-count: maximum stream count seen on entire group during log period
 - Recent Stream Rejected: Number of streams rejected due to stream limit exceeded in group

The log file shows the exact stream limit exceeded:

```
exceed_client_limit= exceed_group_limit=
exceed_mtree_limit=.
```

log-condition

Only log groups with updated states, or log all groups

reset-interval

- `Debug/clients_stats.log` file reset data counters interval
- Default value is one day
- Maximum value can be set for seven days
- On the reset interval, after statistics are written to file, these per-group counters are reset:
 - Write – byte count of incoming data
 - Filtered – Boost DSP mode filtered duplicate data
 - Post_lc – Boost DSP mode sent compressed incoming data
 - Read – byte count of outgoing data
 - Network – byte count on the network (IP or FC)

When a reset occurs, a record of every group is written to `clients_stats.log`.

```
client-group stats options reset [log-interval] [log-condition]
[reset-interval]
```

Reset logging options to their default values.

log-interval

- `Debug/clients_stats.log` file update interval
- Default value is 60 minutes

log-condition

- Debug/clients_stats.log file update groups with updated states or always
- Default is only when states have been updated

reset-interval

- Debug/clients_stats.log file reset data counters interval
- Default value is one day

client-group stream-limit

`client-group stream-limit set group-name [client n] [group n]`
 Set stream limits per client in the group and/or for the entire group. Stream limits are only checked when the stream starts.

- The client limit must be within the group limit if a group limit is specified. The group limit or client limit cannot exceed the maximum combined limit for the Data Domain type.
- When the stream limit is set to zero (the default), no checking is done.
- Storage-unit stream limits are checked together with client-group limits.
- The first stream limit settings to exceed will fail stream limit check.
- At the start of a read or write, a failed stream limit check results in an error.
- Once a host starts a read or write operation, limits are not rechecked.
- Statistics kept for client-group show the rejected count due to stream limits.

Note

This command is not supported for VTL.

`client-group stream-limit reset group-name [client] [group]`
 Reset stream limit setting to zero.

Note

This command is not supported for VTL.

CHAPTER 11

cloud

The `cloud` command is used only on systems licensed to run the Data Domain Cloud Tier software option. Command options let you enable the feature and configure a profile. See the *Data Domain Operating System Administration Guide* for details on using the DD OS System Manager user interface.

This chapter contains the following topics:

- [cloud change history](#)106
- [cloud clean](#) 106
- [cloud enable](#) 106
- [cloud profile](#)106
- [cloud provider](#) 107
- [cloud status](#) 108
- [cloud unit](#)108

cloud change history

There have been no changes to this command in this release.

cloud clean

`cloud clean frequency reset`

Reset the cloud tier cleaning frequency to the default value. Role required: admin and limited-admin.

`cloud clean frequency set interval`

Set the frequency for cloud tier cleaning. Role required: admin and limited-admin.

`cloud clean frequency show`

Show the cloud tier cleaning frequency. Role required: admin and limited-admin.

`cloud clean show config`

Show the cloud tier cleaning configuration.

`cloud clean start unit-name`

Start cloud tier cleaning on the cloud unit *unit-name*. Role required: admin.

`cloud clean status`

Show cloud tier cleaning status.

`cloud clean stop`

Stop the cloud tier cleaning process. Role required: admin and limited-admin.

`cloud clean throttle reset`

Reset throttle percentage for cloud tier cleaning to the default value. Role required: admin and limited-admin.

`cloud clean throttle set percent`

Set throttle percentage for cloud tier cleaning (100 is fastest, 0 is slowest). Role required: admin and limited-admin.

`cloud clean throttle show`

Show throttle percentage for cloud tier cleaning.

`cloud clean watch`

Monitor the cloud tier cleaning process.

cloud enable

`cloud enable`

Enable the Data Domain Cloud Tier software option. If a file system already exists, you can enable the cloud tier for that system. If you are creating a new file system, you can enable the cloud tier at the time you create the file system. Role required: admin and limited-admin.

cloud profile

`cloud profile add profile-name`

Add a new cloud profile to the system. Role required: admin and limited-admin.

Note

Before adding the first cloud profile, be sure to import root CA certificates of the cloud provider or any proxy using `adminaccess certificate import ca application cloud file file-name`. Certificates need to be stored in `/ddr/var/certificates` before you run the `adminaccess` command.

- ECS requires access key, secret key and endpoint.
- Virtustream requires access key, secret key, storage class, and region. The storage class can be `premium`, `standard`, `premium-infrequent-access`, or `standard-infrequent-access`.
- AWS S3 requires access key, secret key, storage class, and region. For enhanced security, the Cloud Tier feature uses Signature Version 4 for all AWS requests. Signature Version 4 signing is enabled by default.
- Azure requires account name, whether or not the account is an Azure Government account, primary key, secondary key, and storage class.
- S3 Flexible providers require the provider name, access key, secret key, region, endpoint, and storage class.

```
cloud profile del profile-name
```

Delete an existing cloud profile. Role required: admin and limited-admin.

```
cloud profile modify profile-name
```

Modify an existing cloud profile, including provider credentials. The system prompts you to modify individual details of the 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.

Role required: admin and limited-admin.

```
cloud profile show [all | profile-name]
```

Show details for all cloud profiles or for a specific cloud profile.

cloud provider

```
cloud provider verify
```

Validates the DD Cloud Tier configuration by performing the following validation steps:

- Cloud enablement check: Verifies that DD Cloud Tier is enabled on the Data Domain system, and the appropriate license, passphrase, and configuration are set.
- Connectivity check: Verifies the existence of the correct certificate, and tests the connection to the cloud provider endpoint.
- User account validation: Creates a test cloud profile and bucket based on the specified configuration values.
- Cloud provider validation: Verifies the cloud provider supports the S3 operations required for DD Cloud Tier.

All test data, including the test bucket, created by this command is automatically deleted when the cloud provider verification is complete.

Role required: admin and limited-admin.

Depending on the cloud provider configured, this command requires the following information:

- ECS requires access key, secret key and endpoint.
- Virtustream requires access key, secret key, storage class, and region. The storage class can be `premium`, `standard`, `premium-infrequent-access`, or `standard-infrequent-access`.
- AWS S3 requires access key, secret key, storage class, and region. For enhanced security, the Cloud Tier feature uses Signature Version 4 for all AWS requests. Signature Version 4 signing is enabled by default.
- Azure requires account name, whether or not the account is an Azure Government account, primary key, secondary key, and storage class.
- S3 Flexible providers require the provider name, access key, secret key, region, endpoint, and storage class.

cloud status

```
cloud status
```

Displays whether the Cloud Tier is enabled or not.

cloud unit

```
cloud unit add unit-name profile profile-name
```

Add a cloud unit using the specified profile. Role required: admin and limited-admin.

```
cloud unit del unit-name
```

Delete the specified cloud unit. A cloud unit cannot be deleted when the file system is running. Role required: admin and limited-admin.

Note

The `cloud unit del` command is not allowed on Retention Lock Compliance systems.

```
cloud unit disable unit-name
```

Disable the specified cloud unit. Role required: admin and limited-admin.

```
cloud unit enable unit-name
```

Enable the specified cloud unit. Role required: admin and limited-admin.

```
cloud unit list [unit-name]
```

List a specific cloud unit or all cloud units. Output includes the unit's name, profile, and status. No authorization information is listed. Role required: admin and limited-admin.

CHAPTER 12

compression

Physical capacity measurement (PCM) provides space usage information for a sub-set of storage space. From the command line interface you can view space usage information for MTrees, tenants, tenant units, and pathsets.

This chapter contains the following topics:

- [compression change history](#)..... 110
- [compression physical-capacity-measurement](#)..... 110

compression change history

There have been no changes to this command in this release.

compression physical-capacity-measurement

Physical capacity measurement commands.

```
compression physical-capacity-measurement disable
```

Disable physical capacity measurement. Role required: admin and limited-admin.

```
compression physical-capacity-measurement enable [and-initialize]
```

Enable physical capacity measurement. Role required: admin and limited-admin.

```
compression physical-capacity-measurement pathset add pathset-name paths pathlist
```

Add paths to an existing pathset.

Note

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/coll/mtree3` is selected, do not select any subdirectories under `mtree3`.

Role required: admin and limited-admin.

Argument definitions

Note

Argument definitions for the compression commands are interchangeable. If an argument is not defined under the command you are looking at, you will find it under another command in the compression section.

pathset-name

Specify the name of a pathset.

paths pathlist

Specify a list of pathnames.

```
compression physical-capacity-measurement pathset create pathset-name paths pathlist [measurement-retention {days | default}]
```

Add a new pathset. A pathset is a container that holds a collection of directory or file paths.

Note

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/coll/mtree3` is selected, do not select any subdirectories under `mtree3`.

Role required: admin and limited-admin.

Argument definitions**measurement-retention {*days* | default}**

Specify the number of days for retention of measurement reports. The default is 180 days.

```
compression physical-capacity-measurement pathset del pathset-name paths pathlist
```

Delete paths from an existing pathset. Role required: admin and limited-admin.

```
compression physical-capacity-measurement pathset destroy pathset-name
```

Destroy a pathset. Role required: admin and limited-admin.

Argument definitions***pathset-name***

Specify the name of a pathset.

```
compression physical-capacity-measurement pathset modify pathset-name [measurement-retention {days | default}]
```

Modify an existing pathset. Role required: admin and limited-admin.

```
compression physical-capacity-measurement pathset show detailed [all | pathset-name]
```

Detailed list of pathsets. Role required: admin and limited-admin.

Argument definitions**all**

Shows all information about the object specified by the command option.

```
compression physical-capacity-measurement pathset show list [all | pathset-name]
```

List pathsets. Role required: admin and limited-admin.

```
compression physical-capacity-measurement sample show current {all | user user | task-id id | pathsets pathset-list tenants tenant-list | tenant-units tenant-unit-list | mtrees mtree-list}
```

Show the status of specified tasks. Role required: admin and limited-admin.

Argument definitions**user *user***

Specify a username.

task-id *id*

Specify a task id for a physical capacity measurement task.

pathsets *pathset-list*

Specify the name of a list of pathsets.

tenants *tenant-list*

Specify the list of tenant names.

tenant-units *tenant-unit-list*

Specify the list of tenant unit names.

mtrees *mtree-list*

Specify certain MTrees.

```
compression physical-capacity-measurement sample show detailed-
history [all | user user | task-id id | pathsets pathset-list |
tenants tenant-list | tenant-units tenant-unit-list | mtrees
mtree-list] [last n {hours | days | weeks | months |
measurements} | start MMDDhhmm[[CC]YY] [end MMDDhhmm[[CC]YY]]
```

Show the detailed history of compression physical capacity measurement samples. Role required: admin and limited-admin.

Argument definitions

[last *n* {hours | days | weeks | months | measurements} | start *MMDDhhmm* [[*CC*]*YY*][end *MMDDhhmm* [[*CC*]*YY*]]

Specify the timeframe (last hours, days, and so on) to display statistics for. And, or, you can specify a beginning and ending date.

start *MMDDhhmm*[[*CC*]*YY*][end *MMDDhhmm*[[*CC*]*YY*]]

Specify starting and ending dates and times for reporting.

```
compression physical-capacity-measurement sample show error-
history [all | task-id id | pathsets pathset-list | tenants
tenant-list | tenant-units tenant-unit-list | mtrees mtree-
list] [last n {hours | days | weeks | months | errors} | start
MMDDhhmm[[CC]YY] [end MMDDhhmm[[CC]YY]]
```

Show the error history for compression physical capacity measurement samples. Role required: admin and limited-admin.

Argument definitions

[last *n* {hours | days | weeks | months | errors} | start *MMDDhhmm* [[*CC*]*YY*][end *MMDDhhmm* [[*CC*]*YY*]]

Specify the timeframe (last hours, days, and so on) to display statistics for. And, or, you can specify beginning and ending dates.

```
compression physical-capacity-measurement sample show history
[all | user user | task-id id | pathsets pathset-list | tenants
tenant-list | tenant-units tenant-unit-list | mtrees mtree-
list] [last n {hours | days | weeks | months | measurements} |
start MMDDhhmm[[CC]YY] [end MMDDhhmm[[CC]YY]]
```

Show the history of compression physical capacity measurement samples. Role required: admin and limited-admin.

Note

The DD System prunes the historical physical capacity measurement samples on a daily basis and keeps the following distribution of historical samples: for MTrees, tenant units, and tenants, no more than 1 sample per hour for the last 90 days, then no more than 1 per day for the last year, then no more than 1 per week for the last 10 years. For pathsets, the historical samples are kept according to the measurement-retention specified (the default is 180 days) when the pathset was created or modified.

```
compression physical-capacity-measurement sample start
[pathsets pathset-list | tenants tenant-list | tenant-units
```



```
tenant-unit-list | mtrees mtree-list] [priority {normal | urgent}]
```

Start the compression physical capacity measurement sample tasks. Role required: admin and limited-admin.

⚠ WARNING

When multiple objects are specified by this command, the system attempts to start measurement samples for each object. However, the command succeeds only if measurement samples actually start for all specified objects. Otherwise, the system does not start samples for any of the specified objects and the command fails.

Argument definitions

priority {normal | urgent}

Specify normal or urgent. Normal priority submits a measurement task to the processing queue. Urgent priority submits a measurement task to the front of the processing queue.

```
compression physical-capacity-measurement sample stop {all | task-id id | pathsets pathset-list | tenants tenant-list} [tenant-units tenant-unit-list | mtrees mtree-list]
```

Stop the compression physical capacity measurement sample tasks. Role required: admin and limited-admin.

⚠ CAUTION

The system stops the specified object's measurement task(s), but not measurement tasks active for any objects contained in the specified object.

```
compression physical-capacity-measurement schedule add name {pathsets pathset-list | tenants tenant-list | tenant-units tenant-unit-list | mtrees mtree-list}
```

Add objects to a compression physical capacity measurement schedule. Role required: admin and limited-admin.

NOTICE

If multiple objects are present in a schedule, the system attempts to start a measurement sample for each object. When the system cannot start a measurement sample for an object, an alert is generated.

Argument definitions

name

Specify a schedule name.

```
compression physical-capacity-measurement schedule create name [pathsets pathset-list | tenants tenant-list | tenant-units tenant-unit-list | mtrees mtree-list] [priority {normal | urgent}] time time [day days |monthly days, last-day]
```

Add a new compression physical capacity measurement schedule. Role required: admin and limited-admin.

Example 45

Example 45 (continued)

Add the schedule `sched1` to pathsets `ps1`, `ps2`, and `ps3`.

```
# compression physical-capacity-measurement schedule add sched1
pathsets ps1 ps2 ps3
```

Argument definitions**time** *time*

Specify the time for a schedule using the following 24-hour formats: 0000 or 00:00

day *days*

With the keyword `day`, specify *days* as the days of the week using either lowercase, three letter abbreviations for the days of the week: `mon`, `tue`, `wed`, `thu`, `fri`, `sat`, `sun`, or as integers: 0 = Sunday, 1 = Monday, 2 = Tuesday, 3 = Wednesday, 4 = Thursday, 5 = Friday, 6 = Saturday.

monthly *days, last-day*

Specify the days of the month using integers (1-31) and, optionally, use the word "last-day" to include the last day of every month in the year.

Example 46

Create a schedule named `sched1` at 4 p.m.

```
# compression physical-capacity-measurement schedule create sched1
pathset ps1 16:00
```

```
compression physical-capacity-measurement schedule del name
{pathsets pathset-list | tenants tenant-list | tenant-units
tenant-unit-list | mtrees mtree-list}
```

Delete objects from a compression physical capacity measurement schedule. Role required: admin and limited-admin.

```
compression physical-capacity-measurement schedule disable name
```

Disable a physical capacity measurement schedule. Role required: admin and limited-admin.

```
compression physical-capacity-measurement schedule destroy
name]
```

Destroy a physical capacity measurement schedule. Role required: admin and limited-admin.

Example 47

Destroy the schedule named `sched1`.

```
# compression physical-capacity-measurement schedule destroy sched1
```

```
compression physical-capacity-measurement schedule enable name
```

Enable a physical capacity measurement schedule. Role required: admin and limited-admin.

```
compression physical-capacity-measurement schedule modify name
[priority {normal | urgent}] time time [day days |monthly days,
last-day]
```

Add a new compression physical capacity measurement schedule. Role required: admin and limited-admin.

```
compression physical-capacity-measurement schedule show [all |
name | pathsets pathset-list | tenants tenant-list | tenant-
units tenant-unit-list | mtrees mtree-list]
```

Show the compression physical capacity measurement schedules. Role required: admin and limited-admin.

```
compression physical-capacity-measurement status
```

Show the physical capacity measurement status (enabled or disabled). Role required: admin and limited-admin.

```
compression physical-capacity-measurement throttle reset
```

Reset the throttle percentage for physical capacity measurement to the default value: 20 percent. Role required: admin and limited-admin.

```
compression physical-capacity-measurement throttle set 1-100
```

Set the throttle percentage for physical capacity measurement. Role required: admin and limited-admin.

```
compression physical-capacity-measurement throttle show
```

Show the throttle percentage for physical capacity measurement. Role required: admin and limited-admin.

Note

The throttle default setting is 20 percent.

compression

CHAPTER 13

config

The `config` command manages Data Domain system configuration settings. Command options include changing individual configuration parameters and viewing the configuration setup. For information on how to configure the system, see the *Data Domain Operating System Initial Configuration Guide* and the *Data Domain Operating System Administration Guide*.

This chapter contains the following topics:

- [config change history](#) 118
- [config reset](#) 118
- [config set](#) 118
- [config setup](#) 119
- [config show](#) 119

config change history

Modified arguments in DD OS 6.1.2

```
config set mailserver host [user user]
```

Command is updated to add the `[user]` option, to specify a username and password on the mail server to improve security.

Modified output in DD OS 6.1.2

```
config show mailserver
```

Command is updated to display the mail server username if one is configured.

config reset

```
config reset location
```

Delete the location configured with the `config set location` command. Role required: admin and limited-admin.

```
config reset mailserver
```

Delete the mail server configured with the `config set mailserver` command. Role required: admin and limited-admin.

```
config reset timezone
```

Reset the time zone to the default, which is US/Pacific. Role required: admin and limited-admin. This command option requires security officer authorization if Retention Lock Compliance is enabled on any MTrees.

config set

```
config set admin-email email-addr
```

Set the email address for the administrator who should receive system alerts and autosupport reports. The system requires one administrative email address. Use the `autosupport` and `alerts` commands to add other email addresses. To check the current setting, use `config show admin-email`. Role required: admin and limited-admin.

```
config set admin-host host
```

Set the machine from which you can log in to the Data Domain system to view system logs and use system commands. The hostname can be a simple or fully qualified hostname or an IP address. The specified host is also added to the FTP and Telnet lists configured with the `adminaccess` command and to the CIFS and NFS lists created with the `cifs share create` and `nfs add` commands. This command provides a quick way to add authentication privileges to multiple lists. To check the current setting, use `config show admin-host`. Role required: admin and limited-admin.

Example 48

```
# config set admin-host admin12.yourcompany.com
```

```
config set location "location"
```

Configure a description of a Data Domain system's location. A description of a physical location helps identify the machine when viewing alerts and autosupport emails. If the

description contains one or more spaces, the description must be in double quotation marks. To check the current setting, use `config show location`. Role required: admin and limited-admin.

Example 49

```
# config set location "row2-num4-room221"
```

```
config set mailserver host [user user]
```

Configure the SMTP mail server used by the Data Domain system. The `user` parameter supports specifying a username and password on the mail server to improve security. To check the current setting, use `config show mailserver`. Role required: admin and limited-admin.

Example 50

```
# config set mailserver mail.yourcompany.com
```

```
config set timezone zonename
```

Set the system clock to a specific time zone. The default setting is US/Pacific. Do any of the following to see the time zone name options.

- Enter `config set timezone ?` to display a list of regional zone names.
- Enter `config set timezone region_zonename` to display zone names for cities and areas in the specified region.
- Enter `config set timezone etc` to display valid GMT zone names.
- See Appendix A in the *Data Domain Operating System Command Reference*.

Note

For additional time zone names that are not displayed in DD OS, see the "Miscellaneous" section of Appendix A in the *Data Domain Operating System Command Reference*.

Changes to the time zone require a system reboot. Role required: admin. This command option requires security officer authorization if any MTrees are enabled with Retention Lock Compliance.

config setup

```
config setup
```

Launch a utility program that prompts you to configure settings for the system, network, filesystem, CIFS, NFS, and licenses. Press Enter to cycle through the selections and confirm any changes when prompted. Choices include **Save**, **Cancel**, and **Retry**.

This command option is unavailable on Retention Lock Compliance systems. Use System Manager to change configuration settings. Role required: admin and limited-admin.

config show

```
config show admin-email
```

Display the administrator email address to which alert summaries and autosupport reports are sent. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
config show admin-host
```

Display the administrative host from which you can log into the Data Domain system to view system logs and use system commands. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
config show all
```

Display all config command settings. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
config show location
```

Display the Data Domain system location description, if you set one. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
config show mailserver
```

Display the name of the mail server that the Data Domain system uses to send email, and the mail server username if one is configured. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
config show timezone
```

Display the time zone used by the system clock. Role required: admin, limited-admin, security, user, backup-operator, or none.

CHAPTER 14

data-movement

The `data-movement` command is used only on systems licensed to run the Data Domain Cloud Tier software option. Command options let you configure data movement policies and options. See the *Data Domain Operating System Administration Guide* for details on using the DD OS System Manager user interface.

This chapter contains the following topics:

- [data-movement change history](#) 122
- [data-movement policy](#) 122
- [data-movement recall](#) 122
- [data-movement resume](#) 122
- [data-movement schedule](#) 123
- [data-movement start](#) 123
- [data-movement status](#) 124
- [data-movement stop](#) 124
- [data-movement suspend](#) 124
- [data-movement throttle](#) 124
- [data-movement watch](#) 125

data-movement change history

There have been no changes to this command in this release.

data-movement policy

```
data-movement policy reset {age-range | age-threshold} mtrees
mtree-list
```

Resets the data-movement policy for the specified MTrees. Role required: admin and limited-admin.

```
data-movement policy set age-range min-age days max-age days
to-tier cloud cloud-unit unit-name mtrees mtree-list
```

Set the age range policy for the specified MTrees. The value for *days* must be from 14 days to 18250 days (approximately 50 years). The upper limit is optional. Role required: admin and limited-admin.

```
data-movement policy set age-threshold days to-tier cloud
cloud-unit unit-name mtrees mtree-list
```

Set the age threshold policy for the specified MTrees. The value for *days* must be from 14 days to 18250 days (approximately 50 years). Role required: admin and limited-admin.

```
data-movement policy show [all | to-tier cloud [cloud-unit
unit-name] | mtrees mtree-list]
```

View the data-movement policy for all MTrees, the specified MTrees, or the specified cloud units.

data-movement recall

```
data-movement recall path pathname
```

Recall a file from the cloud tier to the active tier. The maximum number of files that can be recalled at one time depends on the system memory configuration:

- Systems with more than 256 GB of memory can recall up to 16 files at one time.
- Systems with less than 256 GB of memory can recall up to 8 files at one time.
- DD VE instances can recall up to 4 files at one time.

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. Role required: admin or limited-admin.

Note

If a cloud file is present in a snapshot and not in the active MTree, it cannot be recalled. The only way to recall cloud files that are in snapshots and not in active MTrees is to do a fastcopy operation to copy the files from the snapshot to the active MTree.

data-movement resume

```
data-movement resume
```

Resume data movement to the cloud. Role required: admin and limited-admin.

data-movement schedule

```
data-movement schedule set to-tier cloud {never | days days
time time [every n wks]}
```

Set the schedule for data movement to the cloud. Note that the `days` argument accepts two ranges (and can be either a space- or comma-separated list, or arbitrary text):

- Weekday (Monday-Sunday)
- Day of the month (1-31, regardless of month, plus “last” and “first”)

Any value outside of the two ranges generates an error message. Data movement occurs no more frequently than weekly. Role required: admin and limited-admin.

Note

For `days`, “last” is converted to the value 31. If a schedule is set for the 31st of every month at 10:00 PM, it is not executed on months with fewer than 31 days. This is a known issue.

Example 51

To schedule data movement to occur each Tuesday at 6:00 a.m., enter:

```
# data-movement schedule set to-tier cloud days "tue" time "06:00"
```

Example 52

To schedule data movement to occur on alternate Tuesdays at 6:00 a.m., enter:

```
# data-movement schedule set to-tier cloud days "tue" time "06:00"
every 2 wks
```

```
data-movement schedule show
```

Show the schedule for data movement to the cloud.

data-movement start

```
data-movement start [[to-tier cloud [cloud-unit unit-name]] |
[mtrees mtree-list]]
```

Start data movement to the cloud. You can start data movement to the specified cloud unit, for all MTrees with configured data-movement policies, or for specified MTrees. Role required: admin and limited-admin.

Example 53

```
# data-movement start
Data-movement started.
Run "data-movement watch" to monitor progress.
```

data-movement status

```
data-movement status [path {pathname | all | [queued] [running]
[completed] [failed]} | to-tier cloud [detailed] | all]
Show the status of data movement as well as recall.
```

Note

Starting in DD OS version 6.1.1.5, the data recall status displays **File Size and Logical Bytes Moved instead of Bytes Copied (Pre-comp) and Bytes Copied (Post-comp)**.

The **detailed** option displays data movement statistics for the source and destination MTrees, and information about the files for which data movement is currently in progress.

Note

If a Data Domain system has been upgraded to DD OS version 6.1.1.5 or later from an older version of the software, this command will display 0 files inspected and 0 files eligible until the next data movement operation starts.

This command is not supported when the DD Cloud Tier cloud seeding feature is active.

Example 54

```
# data-movement status
Data-movement to cloud tier:
-----
Data-movement:
  100% complete; time:  0:01:18
Moved (post-comp): 94.24 MiB, (pre-comp): 1.33 GiB,
Files inspected: 29, Files eligible: 29, Files moved: 29, Files failed: 0
```

data-movement stop

```
data-movement stop [path pathname | to-tier {active | cloud} |
all]
```

Stop data movement as specified. Stop a running recall job or remove a queued recall job. All data movement to the active tier or to the cloud tier, or for the specified file can be stopped. Role required: admin and limited-admin.

data-movement suspend

```
data-movement suspend
Suspend data movement to the cloud. Role required: admin and limited-admin.
```

data-movement throttle

```
data-movement throttle reset
```

Reset the throttle value to 100 percent (no throttle). The throttle value takes effect without restarting data movement if it is running. Role required: admin and limited-admin.

```
data-movement throttle set {25 | 50 | 75 | 100 }
```

Set the throttle value to 25, 50, 75, or 100, where 25 is the slowest, and 100 is the fastest. The throttle value takes effect without restarting data movement if it is running. Role required: admin and limited-admin.

Note

The throttle is for adjusting resources for internal Data Domain processes; it does not affect network bandwidth.

```
data-movement throttle show
```

Show the current throttle value. Role required: admin and limited-admin.

data-movement watch

```
data-movement watch
```

View data movement progress while the operation is running. If the operation has completed or is not running, output shows current status only.

Example 55

```
# data-movement watch
Data-movement:
  97% complete; time: 0:02:03
  Moved (post-comp): 94.26 MiB, (pre-comp): 1.33 GiB,
  Files inspected: 29, Files eligible: 29, Files moved: 29, Files failed: 0

Data-movement was started on Nov 15 2017 06:02 and completed on Nov 15 2017 06:05
Moved (post-comp): 94.26 MiB, (pre-comp): 1.33 GiB,
Files inspected: 29, Files eligible: 29, Files moved: 29, Files failed: 0
```

data-movement

CHAPTER 15

ddboost

The `ddboost` command manages the integration of Data Domain systems and disk backup devices. Command options create and delete storage units on the storage server, and display the disk space usage of each storage unit. The Data Domain Boost software option also supports advanced load balancing and failover, distributed segment processing, encryption, and low-bandwidth optimization.

Quotas provision Data Domain system storage among different backup applications. Quotas restrict the logical (uncompressed and unduplicated) storage capacity for each storage unit. DD Boost storage unit quota limits (hard or soft) can be set or removed dynamically. Quotas may also be used to provision various DD Boost storage units with different sizes, enabling an administrative user to monitor the usage of a particular storage unit over time. Note that it is possible to configure quotas on a system and run out of storage before quota limits are reached.

Like MTree quota limits, the `ddboost storage-unit create` command includes optional arguments to specify quota limits at the time the storage unit is created. Output of the `ddboost storage-unit show` command indicates if a quota is defined for the storage unit.

Fibre Channel transport is available for DD Boost via the DD Boost over Fibre Channel service and Automatic Image Replication (AIR) is also supported.

The Multiuser Storage Unit Access Control feature enhances the user experience by supporting multiple usernames for the DD Boost protocol, providing data isolation for multiple users sharing a Data Domain system. Using the DD Boost protocol, the backup application connects to the Data Domain system with a username and password to support this feature. Both the username and password are encrypted using public key exchange. The `tenant-unit` keyword is introduced to the `ddboost storage-unit` command for integration with the Secure Multi-Tenancy feature. One storage unit must be configured for each tenant unit. Each tenant unit can be associated with multiple storage units.

See the *Data Domain Boost for OpenStorage Administration Guide* and the *Data Domain Operating System Administration Guide* for details.

This chapter contains the following topics:

• ddboost change history	129
• ddboost guidelines and restrictions	129
• ddboost association	129
• ddboost clients	131
• ddboost destroy	133
• ddboost disable	133
• ddboost enable	133
• ddboost event	133
• ddboost fc	135
• ddboost file-replication	136
• ddboost ifgroup	140
• ddboost option	143

- [ddboost reset](#)..... 144
- [ddboost set](#)..... 144
- [ddboost show](#)..... 144
- [ddboost status](#)..... 147
- [ddboost storage-unit](#)..... 148
- [ddboost streams](#)..... 151
- [ddboost user](#)..... 152

ddboost change history

Modified output in DD OS 6.1.2

```
ddboost option show [distributed-segment-processing | virtual-synthetics | fc | global-authentication-mode | global-encryption-strength]
```

Output for DD Boost over FC does not display when the command is run on a DD VE instance.

```
ddboost fc show detailed-stats
```

Output for detailed fibre channel statistics now uses the units MB, GB, and TB instead of millions (M), billions (B), and trillions (T) of bytes.

ddboost guidelines and restrictions

- DD Boost is a licensed software option. If basic options do not work, verify that the proper licensing has been implemented on your Data Domain system.
- Quota limits are enforced only if MTree quotas are enabled. A message displays in the output notifying users if the quota feature is disabled.
- When a storage unit is created, quota limits are set to the default MTree quota size.
- If MTree quotas are enabled, backups are stopped if a hard limit is reached.
- Enabling quotas may cause OpenStorage backup applications to report non-intuitive sizes and capacities. See Knowledge Base article 85210, available on the Support portal, for details.
- Only one Automatic Image Replication (AIR) association is allowed for a specified storage unit and the target Data Domain system and the target storage unit.
- Do not use DD Boost Fibre Channel server names to create AIR associations. Use IP server names only.
- DD Boost-over-Fibre Channel operation is expected to continue without user intervention when the Fibre Channel endpoints failover.

ddboost association

```
ddboost association create local-storage-unit {{replicate-to} remote-hostname remote-storage-unit | {replicate-from} remote-hostname remote-storage-unit [{import-to} client-hostname]}
```

Create a storage unit association between the specified storage unit and the target Data Domain system storage unit and optionally, the import-to target client. The import-to client option is applicable when the AIR filecopy feature has been configured to run, and it lets you specify a single NetBackup media server client to receive AIR filecopy events from the target Data Domain system. Role required: admin and limited-admin.

Note

When the AIR import-to client option is used in an environment where multiple media servers could be used to connect to the AIR filecopy target Data Domain system, the import-to client specified in this command is the one and only NetBackup media server client able to receive AIR filecopy event notifications.

If the connected client has a hostname that does not match the defined hostname in the registry, this `ddfs.info` log message is generated: `DDBoost association replication-from has import-to client [<client-name>] specified and a mismatch was found for client requesting events: <client>`

Example 56

```
# ddboost association create feature2 replicate-to kuma-ost11.emc.com feature2
DDBoost association created.
```

Example 57

```
# ddboost association create TEST_LSU7a replicate-from localhost TEST_LSU8a import-to nbu-
client.datadomain.com
DDBoost association created.
```

```
ddboost association destroy local-storage-unit {replicate-to |
replicate-from} remote-hostname remote-storage-unit
```

Destroy the storage unit association for the specified storage unit. Role required: admin and limited-admin.

Note

This command option deletes unprocessed events in the local storage unit if the association specified is `{replicate-from}`. It does not delete user data in the local storage unit.

```
ddboost association show [all | storage-unit storage-unit]
```

Show the storage unit association list for a specified local storage unit or all local storage units with an association. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 58

```
# ddboost association show
Local Storage Unit   Direction           Remote Host         Remote Storage Unit  Import To
-----
TEST_LSU1a.         replicate-from      localhost           TEST_LSU2a..         -
TEST_LSU2a..        replicate-to        localhost           TEST_LSU1a.          -
-----
```

Example 59

```
# ddboost association show
Local Storage Unit   Direction           Remote Host         Remote Storage Unit  Import To
-----
TEST_LSU1a.         replicate-from      localhost           TEST_LSU2a..         -
TEST_LSU2a..        replicate-to        localhost           TEST_LSU1a.          -
TEST_LSU7a          replicate-from      localhost           TEST_LSU8a           nbu-
```

Example 59 (continued)

```
client.datadomain.com
-----
-----
```

ddboost clients

```
ddboost clients add client-list [encryption-strength {none |
medium | high} authentication-mode {one-way | two-way | two-
way-password | anonymous | kerberos}]
```

Note

The maximum length for a client hostname is 63 characters.

Add clients to the DD Boost client list and enable encryption for those clients. Use the `authentication-mode` option to configure the minimum authentication requirement. A client attempting to connect using a weaker authentication setting is blocked. The global authentication mode and the global encryption strength options override any client-specific values if the global values are higher. In that case, a client attempting to connect using a weaker authentication or encryption setting than the global settings is blocked. Both one-way and two-way authentication require the client to be knowledgeable of certificates. If the encryption strength is *none*, only kerberos authentication-mode is supported. Kerberos is only available to clients running BoostFS. Role required: admin and limited-admin.

Argument definitions

These are the valid encryption strength values, which are shown in order of increasing strength:

none

This encryption strength can only be applied if the authentication mode is Kerberos.

medium

Use the AES128-SHA cipher suite.

high

Use the AES256-SHA cipher suite.

These are the valid authentication mode values, which are shown in order of increasing strength:

none

This is the least secure but most backwards-compatible option.

You can select none if your system has crucial performance requirements and you do not need protection from man-in-the-middle (MITM) attacks. Your system can operate in the same manner and avoid suffering performance degradation due to TLS. However, pre-shared key (PSK) authentication is not employed.

When encryption is set to none, authentication must also be set to none.

anonymous

No certificates are exchanged. After the SSL handshake, the communication channel between the DD Boost client and the Data Domain server is encrypted.

one-way

The DD Boost client requests authentication from the Data Domain server, and the Data Domain server sends the appropriate certificate to the DD Boost client. The DD Boost client verifies the certificate. The communication channel between the DD Boost client and the Data Domain server is encrypted.

two-way-password

Two-way authentication is provided by having each side encrypt a value provided by the other side. Both sides prove knowledge of the user password, resulting in each proving its identity to the other side.

Authentication is completed using an SSL handshake. Once the handshake is completed, the communication channel between the DD Boost client and the Data Domain server is encrypted.

two-way

The DD Boost client requests authentication from the Data Domain server using the server certificate. The Data Domain server also requests authentication from the DD Boost client using the client certificate. After authentication through an SSL handshake, the communication channel between the DD Boost client and the Data Domain server is encrypted.

This is the most strongest level of authentication.

Kerberos authentication is as strong as two way authentication, but Kerberos authentication in DD Boost does not currently support encryption.

kerberos

The DD Boost client requests a Kerberos Ticket Granting Ticket (TGT) from the Key Distribution Center (KDC). When the client credentials are verified, the KDC sends a TGT to the DD Boost client. The DD Boost client then requests a Kerberos Ticket Granting Service (TGS) for the desired service. The KDC grants the TGS for the requested service to the DD Boost client.

```
ddboost clients del client-list
```

Delete clients from DD Boost client list. Role required: admin and limited-admin.

```
ddboost clients modify client-list [encryption-strength {none | medium | high} authentication-mode {one-way | two-way | two-way-password | anonymous | kerberos}]
```

Modify clients in the DD Boost client list. See the `ddboost clients add` command for a description of the command variables. If the encryption strength is *none*, only kerberos authentication-mode is supported. Kerberos is only available to clients running BoostFS. Role required: admin and limited-admin.

```
ddboost clients reset
```

Reset DD Boost client list to factory default. Role required: admin and limited-admin.

```
ddboost clients show active [all | client hostname | storage-unit storage-unit | tenant-unit tenant-unit]
```

Show DD Boost client activity. Information displayed includes client hostname, client and server interface IP addresses, operation (`read` or `write`), mode, storage-unit, and tenant-unit. For read operations, mode can be `compressed`. For write operations, mode can be `dsp` and/or `synthetic`. Role required: admin, limited-admin, tenant-admin, security, user, tenant-user, backup-operator, or none.

Example 60

Example 60 (continued)

```
# ddboost clients show active all
Client          Client          Server          Operation    Mode    Storage-unit
Hostname        Interface       Interface
-----
ddboost-dl.datadomain.com  192.168.1.220  192.168.1.224  write       dsp     NEW
-----
```

```
ddboost clients show config
```

Show DD Boost client list. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 61

```
# ddboost clients show config
Client          Encryption Strength  Authentication Mode
-----
*              none                 none
bu_srv_1       high                 one-way
bu_srv_2       high                 two-way
bu_srv_3       high                 anonymous
-----
```

See the `ddboost clients add` command for a description of Encryption Strength and Authentication Mode.

ddboost destroy

```
ddboost destroy
```

Delete all storage units from the Data Domain system. The command permanently removes all data (files) contained in the storage units. You must also manually remove (expire) corresponding catalog entries in the backup software. Role required: admin.

ddboost disable

```
ddboost disable
```

Disable DD Boost. During the process of disabling DD Boost, all file replication transfers and read/write jobs are also disabled. Role required: admin, limited-admin, and security.

ddboost enable

```
ddboost enable
```

Enable DD Boost. If the user, user ID (UID), or group ID (GID) changes, the Data Domain system updates all files and storage units the next time this command is run. Role required: admin and limited-admin.

ddboost event

```
ddboost event show [all | storage-unit storage-unit]
```

Show the event list for the specified local storage unit or all local storage units with a {replicate-from} association. Role required: admin, limited-admin, security, user, backup-operator, none.

Events formatted with the suffix `.event.nnnnnnnn` have been processed but not yet deleted. Events formatted with the suffix `.imgset` have not yet been processed.

Example 62

```
# ddboost event show
DDBoost events:
test2:    bluemedi.a.emc.com_31234_6589_1.event.000000000000000006
192:rtp-ost-sparc1.emc.com_rtp-ost-dd670c2.emc.com_1328637954_1.imgset
```

Output definitions (event.nnnnnnn)

first media server

Hostname of the media server to which the event is first delivered. In the example: `bluemedi.a.emc.com`.

proc_id

Process identifier on the first media server to which the event is initially delivered. In the example: `31234`.

thread_id

Thread identifier on the first media server to which the event is first delivered. In the example: `6589`.

images

Number of images contained in event. In the example: `1`. Typically this number is `1` because only the IM image file is contained in an event.

event

Image set identifier. In the example: `000000000000000006`.

Output definitions (imgset)

job

NetBackup duplication job identifier. In the example: `192`.

source server

Hostname of the NetBackup server. In the example: `rtp-ost-sparc1.emc.com`.

source Data Domain system

Hostname of the Data Domain system from where event originated. In the example: `rtp-ost-dd670c2.emc.com`.

image date

NetBackup image time stamp. In the example: `1328637954`.

images

Number of images contained in event. In the example: `1`. Typically this number is `1` because only the IM image file is contained in an event.

imgset

Image set identifier.

ddboost fc

```
ddboost fc dfc-server-name reset
```

Reset DD Boost Fibre Channel server name. Role required: admin and limited-admin.

```
ddboost fc dfc-server-name set server-name
```

Set DD Boost Fibre Channel server name. The default dfc-server-name has the format DFC-<base hostname>. Role required: admin and limited-admin.

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

Show DD Boost Fibre Channel server name. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost fc dump start logfile-id logfile-id [formatted]
[snaplen bytes] [logfile-count-limit count] [logfile-size-limit
bytes] [virtual-connection virtual-connection-id] [client-
hostname hostname] [initiator initiator] [target-endpoint
endpoint] [destination-tcp-port tcp-port]
```

Start DD Boost Fibre Channel message tracing. Role required: admin and limited-admin.

```
ddboost fc dump status
```

Show DD Boost Fibre Channel message tracing. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost fc dump stop
```

Stop DD Boost Fibre Channel message tracing. Role required: admin and limited-admin.

```
ddboost fc group add group-name initiator initiator-spec
```

Add one or more initiators to a DD Boost Fibre Channel group. Role required: admin and limited-admin.

```
ddboost fc group add group-name device-set [count count]
[endpoint {all | none | endpoint-list}] [disk]
```

Add one or more DD Boost devices to a DD Boost Fibre Channel group; if "disk" is not specified, then the default DFC device type ("Processor") is added. Valid range for count argument is 1-64. Role required: admin and limited-admin.

```
ddboost fc group create group-name
```

Create a DD Boost Fibre Channel group. Role required: admin and limited-admin.

```
ddboost fc group del group-name initiator initiator-spec
```

Remove one or more initiators from a DD Boost Fibre Channel group. Role required: admin and limited-admin.

```
ddboost fc group del group-name device-set {count count | all}
```

Remove one or more DD Boost devices from a DD Boost Fibre Channel group. Role required: admin and limited-admin.

```
ddboost fc group destroy group-name
```

Destroy a DD Boost Fibre Channel group. Role required: admin and limited-admin.

```
ddboost fc group modify group-name device-set [count count]
[endpoint {all | none | endpoint-list}]
```

Modify a device set for a DD Boost Fibre Channel group. Role required: admin and limited-admin.

```
ddboost fc group rename src-group-name dst-group-name
```

Rename a DD Boost Fibre Channel group. Role required: admin and limited-admin.

```
ddboost fc group show detailed group-spec [initiator initiator-name]
```

Show details of DD Boost Fibre Channel groups. Output includes information on device names, system addresses, LUNs, and endpoints. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost fc group show list [group-spec] [initiator initiator-name]
```

Display a list of configured DD Boost Fibre Channel groups. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost fc show detailed-stats
```

Show DD Boost Fibre Channel detailed statistics. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost fc show stats [endpoint endpoint-spec | initiator initiator-spec] [interval interval] [count count]
```

Show DD Boost Fibre Channel detailed statistics periodically based on filter. The interval is an optional number of seconds with a minimum of 1 and a maximum of 4294967295. The count is an optional ordinal value with a minimum of 1 and a maximum of 4294967295. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost fc status
```

Show DD Boost Fibre Channel status. Output includes information on admin state and process state. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost fc trace disable [component {all | component-list}]
```

Disable DD Boost Fibre Channel tracing. Role required: admin and limited-admin.

```
ddboost fc trace enable [component {all | component-list}]  
[level {all | high | medium | low}]
```

Enable DD Boost Fibre Channel tracing. Role required: admin and limited-admin.

```
ddboost fc trace show [component {all | component-list}]
```

Show DD Boost Fibre Channel trace status. Role required: admin, limited-admin, security, user, backup-operator, none.

ddboost file-replication

```
ddboost file-replication option reset {low-bw-optim |  
encryption | ipversion}
```

Reset to default file-replication options. Reset low-bandwidth optimization or encryption to the default value (disabled). Reset IP version to the default value (ipv4). Role required: admin and limited-admin.

Note

Low-bandwidth optimization is not supported on DD Extended Retention systems.

```
ddboost file-replication option set encryption {enabled |  
disabled}
```

Enable or disable encrypted data transfer for DD Boost file-replication. This command must be entered on both systems—the source system and the destination (target) system. Role required: admin and limited-admin.

```
ddboost file-replication option set ipversion {ipv4 | ipv6}
```

Set the preferred IP version for DD Boost file-replication. If the ipversion option is ipv6, IPv6 is the preferred IP address type for managed file-replication. If the ipversion

option is `ipv4`, then IPv4 is the preferred IP address type for managed file-replication. If a preferred IP address is not specified, the default is IPv4. Role required: admin and limited-admin.

Note

If necessary, check the IP version of the destination system to ensure that it is set as desired.

```
ddboost file-replication option set low-bw-optim {enabled | disabled}
```

Enable or disable low bandwidth optimization for DD Boost. This command must be entered on both systems—the source system and the destination (target) system. Default setting is disabled. Role required: admin and limited-admin.

Note

Low-bandwidth optimization is not supported on DD Extended Retention systems.

```
ddboost file-replication option show [encryption]
```

Show state of encryption: enabled or disabled. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost file-replication option show [ipversion]
```

Show IP version options: IPv4 or IPv6. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost file-replication option show [low-bw-optim]
```

Show state of low bandwidth optimization: enabled or disabled. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost file-replication reset stats
```

Clear file-replication statistics when DD Boost is enabled. Role required: admin and limited-admin.

```
ddboost file-replication show active [all | storage-unit storage-unit | tenant-unit tenant-unit]
```

Show the status of a DD Boost file-replication to destination Data Domain systems. Output for low-bandwidth optimization shows the function as enabled and running, or as enabled/off which means a configuration mismatch with the other side. In addition, the output shows the filenames for both the source and destination. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

```
ddboost file-replication show detailed-file-history [all | storage-unit storage-unit | tenant-unit tenant-unit] [duration duration{day | hr}]
```

Show a detailed, file-based replication history. Data for each file name is organized by date, time, and direction (outbound or inbound). The remote hostname is included in the output. The duration of the day and hour must be entered without a space, for example, `10day` or `5hr`. In DD OS Release 5.4.1.0, a new column —`Post-synthetic-optim`— is added to the command outputs to provide statistics for synthetic replication. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

```
ddboost file-replication show detailed-history [all | storage-unit storage-unit] [duration duration{day | hr}] [interval interval {hr}]
```

Show a detailed, cumulative view of file-replication history. Data is organized by date, time, and direction (outbound or inbound). The duration of the day and hour must be

entered without a space, for example, 10day or 5hr. In DD OS Release 5.4.1.0, a new column `—Post-synthetic-optim—` is added to the command outputs to provide statistics for synthetic replication. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

```
ddboost file-replication show file-history [all | storage-unit
storage-unit | tenant-unit tenant-unit] [duration duration{day
| hr}]
```

Show the data-transfer history of inbound and outbound traffic for files in the Data Domain system `/backup` directory. The remote hostname is included in the output. The duration of the day and hour must be entered without a space, for example, 10day or 5hr. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

Note

There is a discrepancy between the network bytes output of the CLI and the network bytes output of the GUI. This is because the CLI reports only one direction of network bytes for `Network (KB)`, and the SMS report used by the GUI reports the sum of `network_bytes_in + network_bytes_out` for `Network Bytes (MiB)`.

```
ddboost file-replication show history [all | storage-unit
storage-unit] [duration duration{day | hr}] [interval
interval{hr}]
```

Show the data transfer history between the source and destination Data Domain systems. The duration of the day and hour must be entered without a space, for example, 10day or 5hr. The output shows the filenames for both the source and destination. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost file-replication show performance [interval sec] [count
count]
```

Show in real time the amount of pre-compressed outbound and inbound data compared to network throughput or post-compressed data. The count displays the number of lines equal to the count value. Output is shown for the specified interval. If no count is specified, the output continues indefinitely until interrupted by `^C`. If no interval is specified, the output is updated every 2 seconds. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 63

The following example displays the output of a Data Domain system which is the source and is replicating to another Data Domain system. Therefore, the source is set only for outgoing replication. For some of the time periods shown, no data was actually sent. If inbound replication traffic was set, then the columns under inbound would be filled in similarly. The `ddboost file-replication show performance` command is used without any options:

```
# ddboost file-replication show performance
01/14 08:48:05
      Outbound
Pre-comp  Network      Pre-comp  Network
(KB/s)    (KB/s)    (KB/s)    (KB/s)
-----
      165521      442          -          -
      2245638     6701          -          -
```

Example 63 (continued)

-	-	-	-
906275	2677	-	-
1280554	3801	-	-
-	-	-	-
2386719	7046	-	-
294790	899	-	-
-	-	-	-
2491855	7383	-	-
-	-	-	-
-	-	-	-
1459252	4323	-	-
-	-	-	-
-	-	-	-
2556742	7575	-	-

Note

Managed file replication statistics are populated under the `ddboost file-replication show performance` command and kept separate from the backup and restore stream stats.

```
ddboost file-replication show stats
```

Monitor outbound and inbound traffic on a Data Domain system during replication. Compression ratio increases when low-bandwidth optimization is enabled. Role required: admin, security, user, backup-operator, none.

In DD OS Release 5.4.1.0, a new row —Bytes after synthetic optimization — is added to the command outputs to provide statistics for synthetic replication.

Note

Managed file replication statistics are populated under the `ddboost file-replication show stats` command and kept separate from the backup and restore stream stats.

Example 64

The following example displays the output of a Data Domain system which is logging the inbound replication traffic:

```
# ddboost file-replication show stats
Direction:                                Outbound
Network bytes sent:                        0
Pre-compressed bytes sent:                 0
Bytes after synthetic optimization:         0
Bytes after filtering:                     0
Bytes after low bandwidth optimization:     0
Bytes after local compression:             0
Compression ratio:                        0

Direction:                                Inbound
Network bytes received:                    292,132,082,144
Pre-compressed bytes received:             36,147,804,001,570
Bytes after synthetic optimization:        36,147,804,001,570
Bytes after filtering:                     1,664,434,329,750
```

Example 64 (continued)

```
Bytes after low bandwidth optimization: 1,664,434,329,750
Bytes after local compression:        179,887,419,678
Compression ratio:                    123.7
```

ddboost ifgroup

```
ddboost ifgroup add group_name {interface {ipaddr | ipv6addr} |
client host}
```

Note

This command is deprecated. An alternative command is `ifgroup add group_name`.

Add an interface, client, or both to *group-name* or to the default group. Prior to adding an interface you must create the *group_name* unless the group name is the default group. Role required: admin, limited-admin.

This command provides full ifgroup 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. The default group behaves in the same manner as any other group.

- The group-name “default” is created during an upgrade of a fresh install and is always used if *group_name* is not specified.
- You can 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*.

- An ifgroup client is a member of a single ifgroup *group-name* and may consist of a fully qualified domain name (FQDN) such as `ddboost.datadomain.com`, wild cards such as `*.datadomain.com` or `*`, a short name such as `ddboost`, or IP range of the client (`xx.xx.xx.0/24` for IPv4 or `xxxx::0/112` for IPv6, for example). When a client's source IP address is evaluated for access to the ifgroup, the order of precedence is:
 1. IP address of the connected Data Domain system
 2. Connected client IP range. This host-range check is useful for separate VLANs with many clients where there isn't a unique partial hostname (domain). For IPv4, 16, 20, 24, 28, and 32 bit masks are supported. For IPv6, 64, 112, and 128 bit masks are supported.

3. Client Name: `abc-11.d1.com`
4. Client Domain Name: `*.d1.com`
5. All Clients: `*`

If none of these checks find a match, ifgroup interfaces are not used for this client.

For detailed information about this order of precedence, see the *Data Domain Boost for Partner Integration Administration Guide* or the *Data Domain Boost for OpenStorage Administration Guide*.

- By default, the maximum number of groups is eight. It is possible to increase this number by editing the system registry and rebooting.

Additionally, the IP address must be configured on the Data Domain system and its interface must be enabled. You can add public or private IP addresses for data transfer connections. After adding an IP address as an interface, you can enable advanced load balancing and link failover.

See the *Data Domain Boost for Partner Integration Administration Guide* or the *Data Domain Boost for OpenStorage Administration Guide*, and the *Data Domain Operating System Administration Guide* for more information on interface groups.

```
ddboost ifgroup create group-name
```

Note

This command is deprecated. An alternative command is `ifgroup create group_name`.

Create a group with the name *group-name* for the interface. Group names may contain alphanumeric characters, hyphens, and underscores. System hostnames, fully qualified hostnames, and wildcard hostnames indicated by an asterisk may also be used. Reserved group names that cannot be used are `default`, `all`, or `none`. Role required: `admin`, `limited-admin`.

```
ddboost ifgroup del group_name {interface {ipaddr | ipv6addr} | client host}
```

Note

This command is deprecated. An alternative command is `ifgroup del group_name`.

Remove an interface, client, or both from *group_name* or default group. Deleting the last IP address interface disables the ifgroup. If this is the case, you have the option of terminating this command option. Role required: `admin`, `limited-admin`.

```
ddboost ifgroup destroy group-name
```

Note

This command is deprecated. An alternative command is `ifgroup destroy group_name`.

Destroy the group name. Only empty groups can be destroyed. Interfaces or clients cannot be destroyed but may be removed sequentially or by running the command option `ddboost ifgroup reset group-name`. Role required: `admin`, `limited-admin`.

Note

The group-name “default” cannot be destroyed.

```
ddboost ifgroup disable group-name
```

Note

This command is deprecated. An alternative command is `ifgroup disable group_name`.

Disable a specific group by entering the *group-name*. If *group-name* is not specified, the command applies to the default group. Role required: admin, limited-admin.

```
ddboost ifgroup enable group-name
```

Note

This command is deprecated. An alternative command is `ifgroup enable group_name`.

Enable a specific group by entering the *group-name*. If *group-name* is not specified, the command applies to the default group. Role required: admin, limited-admin.

```
ddboost ifgroup rename source-group-name destination-group-name
```

Note

This command is deprecated. An alternative command is `ifgroup rename`.

Rename the ifgroup *source-group-name* to *destination-group-name*. This command option does not require disabling the group. The default group cannot be renamed. Role required: admin, limited-admin.

```
ddboost ifgroup reset group-name
```

Note

This command is deprecated. An alternative command is `ifgroup reset group_name`.

Reset a specific group by entering the *group-name*. If *group-name* is not specified, the command applies to the default group. Role required: admin, limited-admin.

```
ddboost ifgroup show config {interfaces | clients | groups | all} [group-name]
```

Note

This command is deprecated. An alternative command is `ifgroup show config`.

Display selected configuration options. If no selection is made, all information about the specified *group-name* is shown. Role required: admin, limited-admin, security, user, backup-operator, none.

If *group-name* is not specified, information for all the groups is shown. Select the all argument to view configuration options of all groups. All users may run this command option.

```
ddboost ifgroup status group-name
```

Note

This command is deprecated.

Show status of the specified *group-name*: enabled or disabled. Role required: admin, limited-admin, security, user, backup-operator, none.

If *group-name* is not specified, status for the default group is shown. All users may run this command option.

ddboost option

```
ddboost option reset {distributed-segment-processing | virtual-synthetics | fc | global-authentication-mode | global-encryption-strength}
```

Reset distributed segment processing and virtual synthetics to the default option of enabled. Reset Fibre Channel to the default option of disabled. Reset global authentication mode and global encryption strength to the default value "none." Due to dependencies between these two global values, both global values are reset to "none" when either is reset. Role required: admin, limited-admin.

```
ddboost option set distributed-segment-processing {enabled | disabled}
```

Enable or disable the distributed segment processing feature on DD Boost. Distributed-segment-processing feature is supported on single-node configurations and systems with Extended Retention only. Role required: admin, limited-admin.

```
ddboost option set fc {enabled | disabled}
```

Enable or disable Fibre Channel for DD Boost. Fiber Channel features are supported on single-node configurations and systems with Extended Retention. Role required: admin, limited-admin.

```
ddboost option set global-authentication-mode {none | two-way | two-way-password} global-encryption-strength {none | medium | high}
```

Enables and sets global authentication and encryption mode and strength. Due to dependencies, both options must be set at once. Encryption can be "none" only if authentication is "none."

Setting the global authentication mode and encryption strength establishes minimum levels of authentication and encryption that all connection attempts by all clients must meet or surpass. Clients attempting to connect at levels less than either of these values will be blocked.

You should also note the following:

- The global authentication mode and the global encryption strength options override any client specific values if the global values are higher.
- If client settings are defined for a particular client or group of clients and the client settings are stronger than the global settings, the client settings will be the minimum required in order to make a connection.
- If the client settings are lower than the global values, the global values will be required in order to make a connection.

```
ddboost option set virtual-synthetics {enabled | disabled}
```

Enable or disable the virtual synthetics feature on the DD Boost. Virtual synthetics features are supported on single-node configurations and systems with Extended Retention. Role required: admin, limited-admin.

```
ddboost option show [distributed-segment-processing | virtual-synthetics | fc | global-authentication-mode | global-encryption-strength]
```

Show status of distributed segment processing, virtual synthetics, Fibre Channel, global authentication, or global encryption. If no argument is specified, status for all

arguments are shown. Status is enabled or disabled. Default is enabled for distributed segment processing and virtual synthetics. Default is disabled for Fibre Channel. All users may run this command. Default is "none" for both global authentication mode and global encryption strength. Role required: admin, limited-admin, security, user, backup-operator, none.

```
ddboost option show global-authentication-mode | global-encryption-strength
```

Shows global authentication and encryption mode and strength.

ddboost reset

```
ddboost reset stats
```

Reset statistics when DD Boost is enabled, or as a network recovery procedure to clear job connections after the network connection is lost. Role required: admin, limited-admin.

```
ddboost reset user-name user-name
```

This command is deprecated. Use `ddboost user unassign` command instead. Role required: admin, limited-admin.

ddboost set

```
ddboost set user-name user-name
```

This command is deprecated. Use `ddboost user assign` command instead. Role required: admin, limited-admin.

Note

The `ddboost user assign` command will be the default command to create users for DD Boost storage units.

ddboost show

```
ddboost show connections [detailed]
```

Show DD Boost active clients and client connections. Client information includes name, idle status, plug-in version, OS version, application version, encryption, DSP, and transport. Using the `detailed` option provides CPU and memory data. Role required: admin, limited-admin, security, user, backup-operator, none.

Note

- When DD Boost Fibre Channel is enabled, connections are listed in the category Interfaces and are named DDBOOST_FC. The ifgroup Group Name category does not apply to DD Boost Fibre Channel; therefore, the group name is listed as n/a.
 - AIR replication job count will be displayed as a `Src-repl` job in the output of a `ddboost show connections` command, the same as other NetBackup and Backup Exec optimized duplication jobs.
 - Both the control connection for file replication and the actual replication interfaces are displayed.
-

```
ddboost show histogram
```


Display a DD Boost histogram for the Data Domain system. Role required: admin, security, user, backup-operator, none.

The DD Boost histogram table lists the set of protocol requests sent from a DD Boost client to the Data Domain system. The table shows how many of each request were sent from the client, how many were responded to with an error code, and the processing time broken down in specific intervals. This table is used primarily by Data Domain field engineers to isolate configuration or performance issues.

Note

A protocol message with an error count may not indicate a problem. For example, a DDP_LOOKUP with a high error count may just mean that the application issued a request to find a file before creating it simply to verify that the file didn't already exist.

Output definitions

mean

The mathematical mean time for completion of the operations, in milliseconds.

std-dev

The standard deviation for time to complete operations, derived from the mean time, in milliseconds.

<1ms

The number of operations that took less than 1 millisecond.

<5ms

The number of operations that took between 1 milliseconds and 5 milliseconds.

<10ms

The number of operations that took between 5 milliseconds and 10 milliseconds

<100ms

The number of operations that took between 10 milliseconds and 100 milliseconds.

<1s

The number of operations that took between 100 milliseconds and 1 second.

<10s

The number of operations that took between 1 second and 10 seconds.

>10s

The number of operations that took more than 10 seconds.

total

The total time taken for a single operation, in milliseconds.

max

The maximum time taken for a single operation, in milliseconds.

min

The minimum time taken for a single operation, in milliseconds.

Example 65

Example 65 (continued)

```
# ddboost show histogram
07/23 09:43:16

-----
OPER          mean      std-dev    <1ms    <5ms    <10ms    <100ms  ...
-----
...
DDP_GETATTR  0.00ms    0.00ms     0        0        0         0 ...
DDP_LOOKUP   3.34ms    29.17ms   13389    126      137       617 ...
DDP_WRITE    0.30ms    4.98ms   125...   5048    1776     1502 ...
...

```

```
ddboost show stats [ interval seconds ] [count count]
```

Show DD Boost statistics. The interval is an optional number of seconds with a minimum of 1 and a maximum of 4294967295. The count is an optional ordinal value with a minimum of 1 and a maximum of 4294967295. Role required: admin, limited-admin, security, user, backup-operator, none.

Output varies depending on which options are specified.

Example 66

This example shows the default output when neither `interval` nor `count` is specified:

```
# ddboost show stats
07/08 14:54:09

DD Boost statistics:

OPER                                     Total      Failed
DDP_GETATTR                             :           0      [0]
DDP_LOOKUP                               :           0      [0]
DDP_ACCESS                               :           0      [0]
DDP_READ                                 :           0      [0]
DDP_WRITE                                :           0      [0]
DDP_CREATE                               :           0      [0]
DDP_REMOVE                               :           0      [0]
DDP_REaddir                              :           0      [0]
DDP_FSSTAT                               :           0      [0]
DDP_REPL_START                           :           0      [0]
DDP_REPL_STOP                            :           0      [0]
DDP_REPL_STATUS                          :           0      [0]
DDP_QUERY                                :          12      [0]
.
.
.
.

-----
Count      Errors
-----
Image creates      0          0
Image deletes     0          0
Pre-compressed bytes received  0          -
Bytes after filtering  0          -
Bytes after local compression  0          -
Network bytes received  0          -
Compression ratio  0.0        -
Total bytes read   0          0
Token Access

```

Example 66 (continued)

```

Connected using secure token 0      0
Exceptions
  Key Not Found                0      -
  Failed to Decrypt            0      -
  Version Failure              0      -
  UID Failure                  0      -
  GID Failure                  0      -
  Invalid Start Time           0      -
  Expired Token                0      -
  Invalid Client               0      -
  Invalid Serial Number        0      -
  Path Failure                  0      -
-----

```

Example 67

This example shows the output when both `interval` and `count` are specified:

```

# ddboost show stats interval 10 count 5
08/11 06:13:34
Backup      Post-comp   Network     Restore     Network     Backup      Restore
KB/s        Written KB/s  In KB/s    KB/s        Out KB/s    Conn        Conn
-----
          0          0          0          0          0          0          0
          0          0          0          0          0          0          0
          0          0          0          0          0          0          0
          0          0          0          0          0          0          0
          0          0          0          0          0          0          0

```

```
ddboost show user-name
```

This command is deprecated. Use `ddboost user show` command instead. Role required: admin, limited-admin, security, user, backup-operator, none.

The output will display the default DD Boost user if one is configured, otherwise, the output will display that there is no default user.

ddboost status

```
ddboost status
```

Display status of DD Boost: enabled or disabled. Role required: admin, limited-admin, security, user, backup-operator, none.

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 output of the `ddboost status` command is: DD Boost status: enabled for ProtectPoint only. If both licenses are installed, the output is unchanged: DD Boost status: enabled.

ddboost storage-unit

```
ddboost storage-unit create storage-unit user user-name
[tenant-unit tenant-unit] [quota-soft-limit n {MiB|GiB|TiB|
PiB}] [quota-hard-limit n {MiB|GiB|TiB|PiB}] [report-physical-
size n {MiB|GiB|TiB|PiB}] [write-stream-soft-limit n] [read-
stream-soft-limit n] [repl-stream-soft-limit n] [combined-
stream-soft-limit n] [combined-stream-hard-limit n]
```

Create a storage unit, assign tenant, and set quota and stream limits. Role required: admin, limited-admin.

Note

If the quota feature is not enabled, the quota is created but a message appears stating the feature is disabled and quota limits are not enforced.

Note

The `tenant-unit` option is introduced for integration with the Secure Multi-Tenancy (SMT) feature. If a `tenant-unit` is specified, and the storage-unit user has a role other than none, the command fails. To remove a user's association with a tenant-unit, use the `ddboost storage-unit modify` command and set the `tenant-unit` value to none.

For more information about SMT, refer to the *Data Domain Operating System Administration Guide*.

Storage unit names can be up to 50 characters. Naming conventions for creating storage units include uppercase and lowercase letters—A-Z and a-z, numbers 0–9, embedded space, comma, period, exclamation mark, hash mark, dollar sign, percent sign, plus sign, at sign, equal sign, ampersand, semi colon, caret, tilde, left and right parentheses, left and right brackets, left and right braces.

You can assign four types of soft stream warning limits against each storage-unit (read, write, replication, and combined), and you can assign a combined hard stream limit. Assigning a hard stream limit per storage-unit enables you to fail new DD Boost streams when the limit is exceeded, including read streams, write streams, and replication streams. The hard stream limit is detected before the stream operation starts. The hard stream limit cannot exceed the capacity of the Data Domain system model, and it cannot be less than any other single limit (read, write, replication, or combined).

The following example shows how to create a storage unit with stream limits:

```
# ddboost storage-unit create NEW_STU0 user user2 write-stream-soft-limit 5
read-stream-soft-limit 1 repl-stream-soft-limit 2 combined-stream-hard-limit 10
Created storage-unit "NEW_STU0" for "user2".
Set stream warning limits for storage-unit "NEW_STU0".
```

Quotas may cause OpenStorage backup applications to report unexpected sizes and capacities. See Knowledge Base article 85210, available on the Online Support website.

```
ddboost storage-unit delete storage-unit
```

Delete a specified storage unit, its contents, and any DD Boost associations. The deleted storage-unit retains its old name and is shown as deleted in the MTree list. Role required: admin.

This command fails to delete the specified storage-unit if there is a data movement policy that is configured on the storage-unit, or a data movement operation is in progress when the delete command is issued.

Note

You must also manually remove (expire) corresponding catalog entries from the backup application.

```
ddboost storage-unit modify storage-unit [user user-name]
[tenant-unit {tenant-unit | none}] [quota-soft-limit {n {MiB|
GiB|TiB|PiB} | none}] [quota-hard-limit {n {MiB|GiB|TiB|PiB} |
none}] [report-physical-size {n {MiB|GiB|TiB|PiB} | none}]
[write-stream-soft-limit {n | none}] [read-stream-soft-limit {n
| none}] [repl-stream-soft-limit {n | none}] [combined-stream-
soft-limit {n | none}] [combined-stream-hard-limit {n | none}]
```

Modify storage-unit user, tenant, and quota and stream limits. Specifying *none* for any parameter disables that parameter. Role required: admin, limited-admin.

Note

If a tenant-unit value other than none is specified, and the storage-unit user has a role other than none, the command fails. To remove a user's association with a tenant-unit, set the `tenant-unit` value to `none`.

The following example shows how to modify the stream limits for a storage unit:

```
# ddboost storage-unit modify NEW_STU1 write-stream-soft-limit 3
read-stream-soft-limit 2 repl-stream-soft-limit 1 combined-stream-hard-limit 8
NEW_STU1: Stream soft limits: write=3, read=2, repl=1, combined=none
```

If DD Boost storage units are replicated with MTree or collection replication, each storage unit on the target must have the DD Boost user added with the command `ddboost storage-unit modify` before being accessed by the Boost backup software. The following example sets the user of the storage unit `STU1` to `ostuser`.

```
# ddboost user show
ostuser
Assuming storage-unit STU1
# ddboost storage-unit modify STU1 user ostuser
```

The example below shows that a nonexistent storage-unit cannot be modified:

```
# ddboost storage-unit modify hello user user5
**** Failed to find storage-unit
```

```
ddboost storage-unit rename storage-unit new-storage-unit
```

Rename a storage-unit while maintaining its:

- Username ownership
- Stream limit configuration
- Capacity quota configuration and physical reported size
- AIR association on the local Data Domain system

Role required: admin, limited-admin.

Note

- A DD Boost association on a remote host must be modified manually.
 - You cannot use this command to rename an Mtree.
-

The example below shows the renaming of a storage-unit:

```
# ddboost storage-unit rename task1 tasking1
storage-unit "task1" renamed to "tasking1".
```

```
ddboost storage-unit show [compression] [storage-unit] [tenant-
unit tenant-unit]
```

List storage-units assigned to tenant-unit and images in a storage-unit. Displays the compression for all storage units (the original byte size, global and local compression) or the filenames in a specified storage unit. The list of files in a storage unit is shown in the output only if a storage unit name is specified. This command can filter on a specific storage-unit or tenant-unit. Deleted storage-units have a status of D. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

The example below displays the list of storage-units:

```
# ddboost storage-unit show
Name                Pre-Comp (GiB)  Status  User      Report Physical
                  Size (MiB)
-----
backup              3.0           RW      sysadmin  -
DDBOOST_STRESS_SU  60.0           RW      sysadmin  -
task2               0.0           RW      sysadmin  -
tasking1           0.0           RW      sysadmin  -
DD1                0.0           RW      sysadmin  -
D6                 5.0           RW      sysadmin  -
TEST_DEST          0.0           D       sysadmin  -
STU-NEW            0.0           D       ddul      -
getevent           0.0           RW      ddul      -
DDP-5-7            120.0          RW      sysadmin  -
TESTME             150.0          RW      sysadmin  -
DDP-5-7-F          100.0          RW      sysadmin  -
testSU              0.0           RW      sysadmin  200
-----
D      : Deleted
Q      : Quota Defined
RO     : Read Only
RW     : Read Write
RD     : Replication Destination
```

```
ddboost storage-unit undelete storage-unit
```

Recover a 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

Role required: admin, limited-admin.

Note

- Deleted storage units are available until the next `filesystem clean` command is run.
 - You cannot use this command to restore an Mtree.
 - You cannot use this command to restore a storage unit that is deleted using the DD Boost SDK. To recover a storage unit that is deleted using the DD Boost SDK:
 1. Type `mtree show` to determine which deleted MTree is the storage unit to be undeleted. To help identify it, the first 32 characters of the original storage unit name are included in MTree name, in the following format: *deleted-original-su-name-xxxxxxxx*, where *xxxxxxxx* is a timestamp in ms. If this is not sufficient, look in `messages.engineering` to find the renamed value.
 2. Type `ddboost storage-unit undelete deleted-mtree-name`, using the mtree name identified in step 1.
 3. Type `ddboost rename deleted-mtree-name original-su-name`.
-

The example below shows the recovery of a deleted storage-unit:

```
# ddboost storage-unit undelete task1
Storage-unit "task1" undeleted successfully.
```

ddboost streams

`ddboost streams show active [all | storage-unit storage-unit | tenant-unit tenant-unit]`

Display active streams per storage-unit. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

```
# ddboost streams show active storage-unit STU-1
```

Name	----- Active Streams -----				----- Soft Limits -----				----- Hard Limit -----
	Read	Write	Repl-out	Repl-in	Read	Write	Repl	Combined	Combined
STU-1	0	0	0	0	-	-	-	-	25

DD System Stream Limits: read=30 write=90 repl-in=90 repl-out=82 combined=90

Note

The DD system stream limits above are based on the type of the DD system.

You can assign four types of soft stream warning limits against each storage-unit (read, write, replication, and combined), and you can assign a combined hard stream limit. Assigning a hard stream limit per storage-unit enables you to fail new DD Boost streams when the limit is exceeded, including read streams, write streams, and replication streams. The hard stream limit is detected before the stream operation starts. The hard stream limit cannot exceed the capacity of the Data Domain system model, and it cannot be less than any other single limit (read, write, replication, or combined).

When any stream count exceeds the warning limit quota, an alert is generated. The alert automatically clears once the stream limit returns below the quota for over 10 minutes.

Note

DD Boost users are expected to reduce the workload to remain below the stream warning quotas or the system administrator can change the warning limit configured to avoid exceeding the limit.

To create a storage unit with stream limits, enter:

```
# ddboost storage-unit create NEW_STU0 user user2 write-stream-soft-limit 5
read-stream-soft-limit 1 repl-stream-soft-limit 2 combined-stream-hard-limit 10
Created storage-unit "NEW_STU0" for "user2".
Set stream warning limits for storage-unit "NEW_STU0".
```

To modify the stream limits for storage units, enter:

```
# ddboost storage-unit modify NEW_STU1 write-stream-soft-limit 3
read-stream-soft-limit 2 repl-stream-soft-limit 1 combined-stream-hard-limit 8
NEW_STU1: Stream soft limits: write=3, read=2, repl=1, combined=none
```

```
ddboost streams show history {storage-unit storage-unit |
tenant-unit tenant-unit} [interval {1 | 10 | 60 | 1440} ]
[lastrn {hours | days | weeks | months} | start MMDDhhmm[[CC]YY]]
[end MMDDhhmm[[CC]YY]]
```

Display streams history per storage-unit or a list of storage-units associated with a tenant-unit. The interval is expressed in minutes, and it must be 1, 10, 60, or 1440. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

Example 68

```
# ddboost streams show history storage-unit stu1 interval 10 last
1hours
INTERVAL: 10 mins
 "-" indicates that the data is not available for the intervals

Storage-Unit: "stu1"
Date          Time      read      write      repl-out    repl-in
YYYY/MM/DD   HH:MM    streams   streams    streams     streams
-----
2013/08/29   12:00     0         0         0           0
2013/08/29   12:10     0         0         0           0
2013/08/29   12:20     0         1         0           0
2013/08/29   12:30     0         2         0           0
2013/08/29   12:40     0         2         0           0
2013/08/29   12:50     0         1         0           0
2013/08/29   13:00     0         0         0           0
-----

# ddboost streams show history storage-unit stu2
Storage-unit /data/coll/stu2 not configured
```

ddboost user

```
ddboost user assign user-name-list
```

Assign Data Domain system users to the list of recognized DD Boost users. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

This command is typically used for applications that create storage-units through the DD Boost SDK APIs.

Note

When a storage-unit is created with a valid Data Domain system local user that is not assigned to DD Boost, the user is automatically added to the DD Boost user list.

Example 69

```
# ddboost user assign user1 user2
User "user1" assigned to DD Boost.
User "user2" assigned to DD Boost.

# ddboost user show
DD Boost user
-----
user1
user2
-----

# ddboost user unassign user1
User "user1" unassigned from DD Boost.
```

```
ddboost user option reset user-name [default-tenant-unit]
```

Unassign DD Boost user *user-name* from default tenant-unit. This command removes DD Boost user *user-name* from the list of valid users for DD Boost. However, this command does not unassign a user if the user is still the owner of a storage-unit. Role required: admin, limited-admin.

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

Set the default tenant-unit for the specified DD Boost user. When a storage-unit is created with a user, the tenant-unit is automatically associated with that user. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

Note

The role of the user must be none.

Example 70

The following output displays how to set the tenant-units when working with SMT. (The default tenant-unit is displayed only when SMT is enabled.)

```
# smt tenant-unit create tu2
Tenant-unit "tu2" created.
# ddboost user option set user2 default-tenant-unit tu2
Default-tenant-unit is set to "tu2" for user "user2".

# ddboost user show
DD Boost user   Default tenant-unit
-----
user2           tu2
-----

# ddboost user option reset user2
Default-tenant-unit is reset for user "user2".
```

```
ddboost user revoke token-access user-name-list
```

Revoke the token-key for users on the list. Role required: admin, limited-admin.

Example 71

```
# ddboost user revoke token-access boostuser1 boostuser2
Revoked token access for user "boostuser1".
Revoked token access for user "boostuser2".
```

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

List DD Boost users and, if SMT is enabled, their default tenant-units. This command also shows whether users have token access. You can use the `ddboost user assign` command or the `ddboost storage-unit create user` command to assign users. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

Example 72

```
# ddboost user show
```

DD Boost user	Default tenant-unit	Using Token Access
ddbu1	Unknown	Yes
ddbu2	Unknown	-
ddbu3	Unknown	Yes
ddbu4	Unknown	-
ddbu5	Unknown	-
ddbu6	Unknown	-
ddbu7	Unknown	Yes
ddbu8	Unknown	-

```
ddboost user unassign user-name
```

Unassign a user from the DD Boost user list. This command deletes the user from the DD Boost users list. A user can only be deleted when it does not own any storage-units. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

Note

The `ddboost user unassign` command does not validate the DD Boost user, it only looks to see if the user has been previously assigned to the DD Boost users list.

Example 73

If the administrator is trying to delete a specific user named `user4`, who has not been previously assigned to the users list, by either using the `ddboost user unassign` or `ddboost user option reset` command, then the following output will display that this user is not assigned to the DD Boost users list:

```
# ddboost user unassign user4
*** User "user4" is not assigned to DD Boost.

# ddboost user option reset tenant4
*** User "tenant4" is not assigned to DD Boost.
```

CHAPTER 16

disk

The `disk` command manages disks and displays disk locations, logical (RAID) layout, usage, and reliability statistics. Each Data Domain system reports the number of disks in the system. For a Data Domain system with one or more Data Domain external disk shelves, commands also include entries for enclosures and disks.

This chapter contains the following topics:

- [disk change history](#) 156
- [disk beacon](#) 156
- [disk fail](#) 156
- [disk multipath](#) 157
- [disk port](#) 157
- [disk release](#) 158
- [disk rescan](#) 158
- [disk reset](#) 158
- [disk set](#) 159
- [disk show](#) 159
- [disk status](#) 164
- [disk unfail](#) 165

disk change history

There have been no changes to this command in this release.

disk beacon

```
disk beacon {enclosure-id.disk-id | serialno}
```

Cause the LEDs associated with the specified disk to flash. Use this command to verify communications with a disk or to identify which physical disk corresponds to a disk ID.

The LEDs that flash are the LEDs that signal normal operation on the target disk and the IDENT LEDs for the enclosure and the controller. The power supply IDENT LEDs also flash on DS60 enclosures. Press Ctrl-C to stop the flash. To display disk identification information, enter `disk show hardware`. To beacon all disks in an enclosure, type `enclosure beacon`. Role required: admin, limited-admin.

disk fail

```
disk fail enclosure-id.disk-id
```

Fail a disk and force reconstruction. To display disk identification information, enter `disk show hardware`. Role required: admin, limited-admin.

disk multipath

```
disk multipath option reset {monitor}
```

Disable multipath configuration monitoring. When disabled, failures in paths to disk devices do not generate alerts. Multipath configuration monitoring is disabled by default. Role required: admin, limited-admin.

```
disk multipath option set monitor {enabled | disabled}
```

Enable multipath configuration monitoring. When enabled, failures in paths to disk devices generate alerts and log multipath events. If monitoring is disabled, multipath event logging is not performed, meaning disk multipath show history is not updated. Multipath configuration monitoring is disabled by default. Role required: admin, limited-admin.

```
disk multipath option show
```

Show the configuration of multipath monitoring. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
disk multipath reset stats
```

Clear statistics of all disk paths in expansion shelves. Role required: admin, limited-admin.

```
disk multipath resume port port
```

Allow I/O on specified initiator port. Use `disk multipath status` to display the available ports. Role required: admin, limited-admin.

```
disk multipath show history
```

Show the history of multipath events. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
disk multipath show stats [enclosure enc-id]
```

Show statistics for all disk paths or for the specified enclosure only. To view the enclosure IDs, enter `enclosure show summary`.

```
disk multipath status [port-id]
```

Show multipath configurations and runtime status. To view the port IDs, enter the command without the port ID. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
disk multipath suspend port port
```

Disallow I/O on the specified initiator port, and stop traffic on particular ports during scheduled maintenance of a SAN, storage array, or system. This command does not drop a Fibre Channel link. To view the ports, enter `disk multipath status`. Role required: admin, limited-admin.

disk port

```
disk port enable port-id
```

Enable the specified initiator port. To display the disk ports, enter `disk port show summary`. Role required: admin, limited-admin.

```
disk port show {stats | summary}
```

Show disk port statistics or configuration and status. When the `disabled` status appears without an asterisk, the port is administratively disabled. When the `disabled` status appears with an asterisk and an error message below the report, the system has disabled the port. Role required: admin, limited-admin, security, user, backup-operator, or none.

Example 74

```
# disk port show stats
Port      Command Target      Bus      Host      Device      Device
          Aborts  Resets      Resets   Resets    Additions   Removals
-----
1a        0        0           0         1         0           0
1b        0        0           0         1         0           0
1c        0        0           0         1         0           0
1d        0        0           0         1         0           0
2a        0        0           0         1        16           0
2b        0        0           0         1        16           0
2c        0        0           0         1        32           0
2d        0        0           0         1        32           0
3a        0        0           0         1        16           0
3b        0        0           0         1        16           0
3c        0        0           0         1        32           0
3d        0        0           0         1        32           0
-----
```

Example 75

```
# disk port show summary
Port      Connection Link      Connected      Status
          Type      Speed     Enclosure IDs
-----
1a        SAS                12 Gbps      2              offline
1b        SAS                12 Gbps      2              offline
1c        SAS                12 Gbps      2              offline
1d        SAS                12 Gbps      2              offline
2a        SAS                12 Gbps      2              online
2b        SAS                12 Gbps      2              offline
2c        SAS                12 Gbps      3              online
2d        SAS                12 Gbps      4              online
3a        SAS                12 Gbps      2              online
3b        SAS                12 Gbps      2              offline
3c        SAS                12 Gbps      3              online
3d        SAS                12 Gbps      4              online
-----
```

disk release

```
disk release persistent-id {persistent-id | all}
```

Releases the disk's persistent ID and enables persistent ID to be assigned on the next boot. Role required: admin, limited-admin.

disk rescan

```
disk rescan [enclosure-id.disk-id]
```

Rescan all disks or a specified disk to look for newly removed or inserted disks or LUNs or power on a drive. To view disk IDs with both enclosure ID and disk ID, enter `disk show hardware`. Role required: admin, limited-admin.

disk reset

```
disk reset performance
```

Reset disk performance statistics to zero. Role required: admin, limited-admin.

disk set

```
disk set dev disk-id spindle-group 1-16
```

Assign a LUN group to the disk. To display the disk IDs, enter `disk show hardware`. You must restart the file system after adding the LUN. Role required: admin, limited-admin.

disk show

```
disk show failure-history
```

Display a list of disk failure events, which include the date, time disk ID, enclosure serial number, and disk serial number. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
disk show hardware
```

Display disk hardware information. The output includes a column for slot identification. The identification displayed in the Slot column is based on the type of enclosure that contains each disk.

- Slot numbering begins at 0 for DD2500 and ES30 enclosures.
- Slot numbering begins at 0 for DD4200, DD4500, and DD7200 controllers, however slot 0 is unused and does not appear in output.
- Slot numbering begins at 1 for DD990 controllers and earlier.
- Slot numbering for DS60 enclosures uses a letter and a number to define the row and column location of each disk.

Role required: admin, limited-admin, security, user, backup-operator, or none.

Output Definitions (Disk Information)

Disk (enc/disk)

The enclosure and disk ID numbers.

Slot

The slot number for the disk.

Manufacturer/Model

The manufacturer model designation.

Firmware

The firmware revision on each disk.

Serial No.

The manufacturer serial number for the disk.

Capacity

The data storage capacity of the disk when used in a Data Domain system. The Data Domain convention for computing disk space defines one gigabyte as 230 bytes, giving a different disk capacity than the manufacturer's rating.

Type

The type of disk drive.

Output Definitions (System Information)

Disk

Each LUN accessed by the Data Domain system as a disk.

LUN

The LUN number given to a LUN on the third-party physical disk storage system.

Port WWN

The world-wide number of the port on the storage array through which data is sent to the Data Domain system.

Manufacturer/Model

A label that identifies the manufacturer. The display may include a model ID, 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 No.

The serial number from the third-party physical disk storage system for a volume that is sent to the Data Domain system.

Capacity

The amount of data in a volume sent to the Data Domain system. GiB = Gibibytes, the base-2 equivalent of Gigabytes. MiB = Mebibytes, the base-2 equivalent of Megabytes. TiB = Tebibytes, the base-2 equivalent of Terabytes.

```
disk show performance [interval {5min [count 1-12] | 1hour
[count 1-24] | 1day [count 1-7] | cumulative}] [enclosure-id |
enclosure-id.disk-id | devn]
```

Display disk performance statistics for each disk. Each column displays statistics averaged since the last `disk reset performance` command or the last system power cycle. Role required: admin, limited-admin, security, user, backup-operator, or none.

Argument Definitions

interval

Use the interval argument to display performance data that can indicate trends in disk performance. The data interval, 5 minutes, 1 hour, 1 day, or cumulative, defines the period for which data is reported. The count defines how many of the most recent intervals you want to display.

enclosure-id

Specify an enclosure number to display the performance data for only the disks in that enclosure. To display the enclosure IDs, enter `enclosure show summary`.

enclosure-id.disk-id

Specify an enclosure number and disk number to display the performance data for a specific disk. To display the available disks, enter `disk show performance`.

devn

Specify a SCSI target or vdisk device to display the performance data for the device. To display the available devices, enter `disk show performance`.

Output Definitions**Disk (enc/disk)**

The enclosure and disk numbers.

Read

Disk access statistics for read operations.

Read+Write

Disk access statistics for total (read +write) operations.

Write

Disk access statistics for write operations.

KiB/sec

The average data transfer speed in KiB/second.

IOPs

The average number of read, write, or total (read +write) input/output operations per second (IOPs).

Resp(ms)

The average response time in milliseconds.

Ops >1s

The number of operations that required more than 1 second for processing.

MiB/sec

The average number of mebibytes per second (MiB/s) written to storage. Mebibytes are the base-2 equivalent of Megabytes.

Random

The percentage of random IOPs.

Busy

The average percent of time that at least one command is queued for storage access.

```
disk show reliability-data
```

View details of the hardware state of each disk. Output also includes the operational state of drives and if the drive is present or absent. Output is typically used by Data Domain Support for troubleshooting assistance. Role required: admin, limited-admin, security, user, backup-operator, or none.

Output Definitions**Disk**

The enclosure.disk-id disk identifier.

Slot

The disk slot number.

ATA Bus CRC Err

The uncorrected raw UDMA CRC errors.

Reallocated Sectors

The number of mapped-out defective sectors.

Temperature

The current temperature of each disk in Celsius and Fahrenheit. The allowable case temperature range for disks is from 5 degrees centigrade to 55 degrees centigrade.

```
disk show reservation
```

Display all existing reservation information within attached disks.

```
disk show state
```

Display state information for all disks in a Data Domain system. If a RAID disk group reconstruction is underway, columns for the disk identifier, progress, and time remaining are included in command output. Role required: admin, limited-admin, security, user, backup-operator, or none.

Disk State Definitions

The following describes the symbols that define the state of each disk:

- .
(Period) In-Use. The disk is being used for backup data storage.
- (Dash) Not installed. The enclosure firmware has determined that no disk is installed.
- A**
Absent. DD OS does not detect a disk in the indicated location, and no firmware status is available. The disk may be absent, or there may be some condition that makes the disk appear to be absent.
- C**
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.
- d**
Destination. The disk is in use as the destination for storage migration.
- E**
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.
- K**
Known. The disk is a supported disk that is ready for allocation.
- m**
Migrating. The disk is in use as the source for storage migration.
- O**
Foreign. The disk has been assigned to a tier, but the disk data indicates the disk may be owned by another system.
- P**
Powered Off. The disk power has been removed by Data Domain Support.

R
Reconstruction. The disk is reconstructing in response to a `disk fail` command or by direction from RAID/SSM.

s
Spare. The disk is available for use as a spare.

U
Unknown. An unknown disk is not allocated to the active or retention tier. It might have been failed administratively or by the RAID system.

v
Available. An available disk is allocated to the active or retention tier, but it is not currently in use.

Y
System. System disks store DD OS and system data. No backup data is stored on system disks.

Example 76

The following example shows the output for 15 disk enclosure such as the ES30.

```
# disk show state
Enclosure  Disk
-----
1          . . .
2          s . . . . . . . . . . .
-----

Legend    State              Count
-----
.         In Use Disks       59
s         Spare Disks         4
-         Not Installed   15
-----
Total 18 disks
```

Example 77

The following example shows the output for a system with 2 15-disk enclosures and 2 60-disk enclosures.

```
# 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
   |-----|-----|-----|-----|
   | Pack 1 | Pack 2 | Pack 3 | Pack 4 |
   E(49-60) |. . s |. . s |. . s | - - - |
   D(37-48) |. . . |. . . |. . . | - - - |
   C(25-36) |. . . |. . . |. . . | - - - |
   B(13-24) |. . . |. . . |. . . | - - - |
   A( 1-12) |. . . |. . . |. . . | - - - |
   |-----|-----|-----|-----|
3
4
   |-----|-----|-----|-----|
   | Pack 1 | Pack 2 | Pack 3 | Pack 4 |
   E(49-60) |. . s |. . s |. . s | - - - |
   D(37-48) |. . . |. . . |. . . | - - - |
   C(25-36) |. . . |. . . |. . . | - - - |
   B(13-24) |. . . |. . . |. . . | - - - |
   A( 1-12) |. . . |. . . |. . . | - - - |
   |-----|-----|-----|-----|

Legend  State          Count
-----  -----
.       In Use Disks    59
s       Spare Disks   4
-       Not Installed 15
-----

Total 138 disks
```

`disk show stats`

Provides a dynamic display of the default output for the `disk show performance` command. Press `Ctrl + C` to terminate the dynamic display. Role required: `admin`, `limited-admin`, `security`, `user`, `backup-operator`, or `none`.

disk status

`disk status`

View details on the Data Domain system disk status. Output includes the number of disks in use and failed, the number of spare disks available, and if a RAID disk group reconstruction is underway.

Note

The RAID portion of the display could show one or more disks as Failed while the Operational portion of the display could show all drives operating nominally. A disk can be physically functional and available, but not in use by RAID, possibly because of user intervention.

Reconstruction is done per disk. If more than one disk is to be reconstructed, the disks queued for reconstruction show as spare or hot spare until reconstruction begins.

In the first line of output, disk status is indicated by one of the following, high-level states.

Destination

The disk is in use as the destination for storage migration.

Error

A new head unit is in this state when Foreign storage is present. For a system configured with some storage, the error indicates that some or all of its own storage is missing.

Migrating

The disk is in use as the source for storage migration.

Normal

The system is operational and all disks are available and ready for use.

Warning

One or more of the following conditions require user action.

- RAID system degraded
- Foreign storage
- Failed or Absent disks

Role required: admin, limited-admin, security, user, backup-operator, or none.

disk unfail

```
disk unfail enclosure-id.disk-id
```

This command attempts to make a disk previously marked Failed or Foreign usable to the system. To display the enclosure and disk ID for each disk, enter `disk show hardware`. Role required: admin, limited-admin.

disk

CHAPTER 17

elicense

The `elicense` command manages electronic licenses generated by the Electronic License Management System (ELMS)

This chapter contains the following topics:

- [elicense change history](#) 168
- [elicense reset](#) 168
- [elicense show](#) 168
- [elicense update](#) 169

elicense change history

Modified arguments in DD OS 6.1.2

`elicense reset [restore-evaluation]`

The `[restore-evaluation]` parameter resets a DD VE instance to its factory default licenses.

`elicense show [licenses | locking-id | software-id | scheme | all]`

The `scheme` parameter shows the licensing used on the Data Domain system or DD VE instance.

Modified behavior in DD OS 6.1.2

`elicense update [check-only] [license-file]`

The command checks the validity of the start and end dates to ensure that the end date is not in the past.

elicense reset

`elicense reset [restore-evaluation]`

- Use `elicense reset` to delete all existing licenses.
- Use `restore-evaluation` to reset DD VE to its factory default licenses.

Role required: admin, limited-admin.

Note

When features such as HA, vdisk, VTL, or DD Boost are in use, `elicense reset` returns an error message. Disable any such features before issuing the `elicense reset` command.

elicense show

`elicense show [licenses | locking-id | software-id | scheme | all]`

Show current license information. Specify `locking-id` to display the serial number, `licenses` to display all licenses installed, and `all` to display licenses, locking-ID, and last modified. The `software-id` option is only available for DD VE. Issuing `elicense show` is the same as issuing `elicense show all`. The `scheme` parameter shows the licensing used. For example, schemes for a physical DDR include `Unknown` and `Data Domain licensing`. Supported schemes for DD VE include `node-locked mode` and `served mode`. Role required: admin, limited-admin, security, backup-operator, user.

Example 78

The following output shows `elicense show`.

Example 78 (continued)

```
# elicense show
> System locking-id:
1CEXYV7RNJ55ZVR9R5GHBA14KJ8DEK8RJTJ3NB7DT34A21TX6P5ECMTDHGYDGR9AJZFZUJ
GTG3UZXF5PZ32G62HBMZL5JEF2FMWA67RXBFYG
>
> System software-id:   ELMDDV01188LXW
> Instance software-id: Not available
>
> Licensing scheme: EMC Electronic License Management System (ELMS)
node-locked mode
>
> Capacity licenses:
> ##   Feature      Capacity      Type          State          Expiration
> ---  -
> 1    CAPACITY    18.18 TiB    subscription  grace
2018-01-28    License expired
> ---  -
>
> Other possible note:
> License sever unreachable (N/A for expiration date)
> Locking-id mismatch (N/A for expiration date)
> Unable to validate license (generic case) (N/A for expiration date)
```

elicense update

elicense update [check-only] [*license-file*]

- Use `elicense update` to cut and paste elicense codes. When finished pasting, enter **CTRL+D**.
- Use `elicense update filename` to transfer a `.lic` file to `/ddvar`.
- Use the `check-only` option to validate the evaluation license file content, including:
 - Signature
 - Feature name
 - Capacity values
 - Capacity units
 - Expiration date

Role required: admin, limited-admin.

Note

All licenses for the system have to be put in a single file. Every time that elicensess are updated, the previous licenses are overwritten.

elicense

CHAPTER 18

enclosure

The `enclosure` command identifies and displays information about Data Domain system enclosures and attached expansion shelves. Beginning with 5.4, output from the `enclosure show` command option includes device VPD information for enclosures on newer Data Domain appliances (DD4500 and DD7200). VPD information enables users to monitor systems more efficiently.

This chapter contains the following topics:

- [enclosure change history](#) 172
- [enclosure guidelines and restrictions](#) 172
- [enclosure beacon](#) 172
- [enclosure release](#) 172
- [enclosure show](#) 172
- [enclosure test](#) 177

enclosure change history

There have been no changes to this command in this release.

enclosure guidelines and restrictions

- Enclosure numbers are not static and may change when the system is rebooted. (Numbers are generated according to when the shelves are detected during system startup.)
- If a Data Domain system or a previously installed shelf, or both, require spare disks and none are available, disks from a newly installed shelf are allocated to the existing RAID groups (disk groups) when the new shelf is recognized by the `disk rescan` command. The shelf allocating the disks requires at least 14 disks available for its own RAID group.

enclosure beacon

```
enclosure beacon enclosure
```

Cause the LEDs associated with the specified enclosure to flash. Use this command to verify communications with an enclosure or to identify which physical enclosure corresponds to an enclosure ID.

The LEDs that flash are the LEDs that signal normal operation on all enclosure disks and the IDENT LEDs for the enclosure and the controller. The power supply IDENT LEDs also flash on DS60 enclosures. Press Ctrl-C to stop the flash. Role required: admin, limited-admin.

enclosure release

```
enclosure release persistent-id {serialno | persistent-id | all}
```

Remove the persistent ID assignment of an enclosure from the system. This removal becomes complete when the system is restarted. Use `enclosure show persistent-id` to display the serial numbers and IDs for persistent enclosures. Role required: admin, limited-admin.

Example 79

```
# enclosure release persistent-id US1V4100200097
```

enclosure show

```
enclosure show all [enclosure]
```

Display detailed information about the installed components and component status for all enclosures. The controller is enclosure 1. To see the IDs for all enclosures, enter `enclosure show summary`. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
enclosure show chassis [enclosure]
```

Show part numbers, serial numbers, and component version numbers for one or all enclosures. To see the IDs for all enclosures, enter `enclosure show summary`.
Role required: admin, limited-admin, security, user, backup-operator, or none.

```
enclosure show controllers <enclosure>
```

Display the controller model number and related information for a system controller or expansion shelf. To see the IDs for all enclosures, enter `enclosure show summary`.
Role required: admin, limited-admin, security, user, backup-operator, or none.

```
enclosure show persistent-id
```

The command enables persistent enclosure numbering management method. Role required: admin, limited-admin, security, user, backup-operator, or none.

Controller Definitions (Physical Enclosure Shell)

Enclosure

The number listed here is the enclosure number assigned by the Data Domain OS. (Enclosure 1 is the system controller.) This number is the argument passed to the command.

Model

The product name, such as DD860 or ES30.

Capacity

The number of usable drive slots in the enclosure.

Serial No.

The serial number of the physical enclosure. As with the WWN, this describes the enclosure and does not change if components are swapped. Depending on when the enclosure was manufactured, this may be the same value as the WWN. This value matches the serial number printed on the label on the back of the enclosure.

Number of Controllers

The number of shelf controllers currently inserted into the enclosure.

Output Definitions (Controller Modules)

Controller 1

Identifies which shelf controller module the block of information is for. If both shelf controllers are installed, there are blocks for Controller 1 and Controller 2.

Firmware

The revision level of the firmware that resides on the shelf controller. This value can be different for each shelf controller.

Serial

The serial number for the shelf controller. The serial number is different for each shelf controller and differs from the enclosure serial number.

Part

The part number for the shelf controller.

Status

The current status of the shelf controller.

Type

The type of shelf controller.

```
enclosure show cpus [enclosure]
```

Display CPU information, such as the number of CPUs, type, and speed for one or all enclosures. To see the IDs for all enclosures, enter `enclosure show summary`.

Role required: admin, limited-admin, security, user, backup-operator, or none.

```
enclosure show fans [enclosure]
```

Display the current status of fans in one or all enclosures. To see the IDs for all enclosures, enter `enclosure show summary`. Role required: admin, limited-admin, security, user, backup-operator, or none.

Output Definitions**Enclosure**

The enclosure number, starting from 1, for the Data Domain system.

Description

The ID for each power or cooling unit.

Level

The fan speed. This value depends on the internal temperature and amount of cooling required.

Status

The fan status: OK or Failed.

```
enclosure show io-cards [enclosure]
```

Display I/O card information such as the device type, firmware revision, and address for one or all enclosures. To see the IDs for all enclosures, enter `enclosure show summary`. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
enclosure show memory [enclosure]
```

Show the current DIMM inventory, speed, size, and ID numbers for one or all enclosures. To see the IDs for all enclosures, enter `enclosure show summary`. Memory size is calculated in base-2 Mib. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
enclosure show misconfiguration [enclosure]
```

Displays any misconfigurations detected on a DD9500 controller. This command is not supported on other controllers and expansion shelves in this release.

Example 80

```
# enclosure show misconfiguration
Memory DIMMs:
  No misconfiguration found.
IO Cards:
Slot   Device           Status
----   -
0      NVRAM              missing
9      NVRAM              misplaced
-----
CPUs:
  No misconfiguration found.
Disks:
  No misconfiguration found.
```

```
enclosure show nvram [enclosure]
```

Displays NVRAM ID information, component temperatures, and locations for one or all enclosures. If output indicates one or more component errors, an alerts notification is sent to the designated group and the Daily Alert Summary email includes an entry citing details of problem. To see the IDs for all enclosures, enter `enclosure show summary`. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
enclosure show persistent-id
```

Display the online enclosures with persistent identification numbers. To identify offline enclosure numbers, enter `enclosure show summary`.

Persistent enclosure IDs are featured in DD OS Release 5.7 and later. A persistent enclosure ID is assigned to the enclosure serial number and remains assigned after power cycles, reboots, and cable changes. If the enclosure is removed, the persistent ID remains assigned until cleared with the `enclosure release persistent-id` command. The system controller is always assigned enclosure ID 1. Each new enclosure is assigned to the lowest unreserved ID number.

Example 81

```
# enclosure show persistent-id
Serial No.      Model No.      Persistent ID
-----
APM00120502639  ES30           2
APM00120502638  ES30           3
APM00120600566  ES30           4
APM00120503381  ES30           5
APM00120600565  ES30           6
APM00120503377  ES30           7
APM00120600563  ES30           8
APM00120503378  ES30           9
US1V4100200126  ES30          10
US1V4100200097  ES30           -
-----
9 enclosure(s) persisted.

(-) Persistent id will be assigned on next boot.
(*) Does not match existing enclosure id until next boot.
```

```
enclosure show powersupply [enclosure]
```

Displays power supply ID and status information for one or all enclosures. To see the IDs for all enclosures, enter `enclosure show summary`. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
enclosure show summary
```

List enclosures, model and serial numbers, state, OEM names and values, and capacity (number of disks in the enclosure). The serial number for an expansion shelf is the same as the chassis serial number, which is the same as the enclosure WWN and the OPS panel WWN. Role required: admin, limited-admin, security, user, backup-operator, or none.

Enclosure states may be one of the following:

Offline

No connectivity to shelf. Shelf was connected previously. Also occurs if there is no power to the enclosure following startup).

Online

Operating as expected. No problems detected.

Fault

Applies to ES20 only. Indicates no communication with firmware.

Found

Enclosure detected. This transition state is very brief and rarely seen.

Error

Hardware or software error.

Software Error

Typically means busy. Try again later.

```
enclosure show temperature-sensors [enclosure]
```

Lists the temperatures for monitored components in one or all enclosures. To see the IDs for all enclosures, enter `enclosure show summary`.

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. If the system temperature drops below or rises above the parameters defined in the profile, the system shuts down. For example, the temperature profile for some system models shuts down the system when the temperature drops below 0 degrees Celsius or rises above 80 degrees Celsius.

Role required: admin, limited-admin, security, user, backup-operator, or none.

Output Definitions

Enclosure

The enclosure number, starting from 1, for the Data Domain system.

Description

The ID for each monitored component. The components listed depend on the model and are often shown as abbreviations. Some examples are:

- 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)
- FP temp (front panel)

C/F

Ambient readings are displayed as positive numbers and indicate the approximate component temperature in Celsius and Fahrenheit. CPU temperatures may be shown in relative or ambient readings. Relative readings are displayed as negative numbers and indicate the difference between the current temperature and the CPU throttling point, when the CPU reduces its power consumption.

Status

If temperature thresholds are defined for a component, the Status column displays the component status determined by the threshold configuration. If the component temperature is within the configured thresholds, the status is OK. Warning status indicates the temperature is above the acceptable threshold, and Critical status indicates the temperature is above the shutdown threshold. When no thresholds are defined for a component, the Status column displays a dash (-).

Example 82

```
# enclosure show temperature-sensors
Enclosure  Description      C/F      Status
-----
1          MLB TEMP 1          43/109   OK
          MLB TEMP 2          33/91    OK
          FP TEMP           30/86    OK
          BP LEFT TEMP      34/93    OK
          BP MIDDLE TEMP     31/88    OK
          BP RIGHT TEMP     30/86    OK
          LP TEMP           42/108   OK
          FHFL TEMP         46/115   OK
          CPU 0 TEMP        42/108   OK
-----
```

Example 83

```
# enclosure show temperature-sensors 2
Enclosure  Description      C/F      Status
-----
2          LCC A              26/79    -
          LCC B              27/81    -
          Internal ambient  21/70    OK
          PSU A Temp #1     27/81    -
          PSU A Temp #2     22/72    -
          PSU B Temp #1     26/79    -
          PSU B Temp #2     22/72    -
-----
```

```
enclosure show topology
```

Show the layout of the SAS enclosures attached to a system. Role required: admin, limited-admin, security, user, backup-operator, or none.

enclosure test

```
enclosure test topology port duration minutes
```

Test communications with the specified port for the specified number of minutes. To display the port IDs, enter `enclosure show io-cards`. Role required: admin, limited-admin.

enclosure

CHAPTER 19

filesys

The `filesys` command displays statistics, capacity, status, and use of the filesystem. Command options also clear the statistics file, and start and stop filesystem processes. The `filesys clean` command options reclaim physical storage within the filesystem. Command output for disk space or the amount of data on disks is computed using base-2 calculations. See the *Data Domain Operating System Administration Guide* for details.

The `filesys archive` command option is specific to a Data Domain system with Extended Retention, and enables administrative users to provision the filesystem with tiered storage. See the *Data Domain Operating System Administration Guide* for details on command usage and examples.

This chapter contains the following topics:

• filesys change history	180
• filesys archive	180
• filesys clean	181
• filesys create	183
• filesys destroy	183
• filesys disable	183
• filesys enable	183
• filesys encryption	184
• filesys expand	193
• filesys fastcopy	193
• filesys option	194
• filesys report	197
• filesys restart	198
• filesys show	198
• filesys status	202
• filesys sync	202

filesys change history

New commands in DD OS 6.1.2

```
filesys encryption key-manager keys create
```

Creates a new active key on the KeySecure key manager.

```
filesys encryption key-manager set key-rotation-policy [every
<n> {weeks | months} | none]
```

Sets a key rotation policy on the KeySecure key manager.

```
filesys encryption key-manager reset [key-rotation-policy]
```

Change the existing key rotation policy on the KeySecure key manager.

```
filesys encryption key-manager keys modify [<key-id> | muid
<key-muid>]
```

Modifies the state of an existing key in the KeySecure key manager.

filesys archive

```
filesys archive unit add
```

Create a retention unit from the retention tier of the file system, change the state of disks or LUNs from Available to In Use, and add the new retention unit to the file system. The system creates a retention unit using all available storage, but fails the operation if the combined size of all available storage in the retention tier is larger than the maximum supported size. Role required: admin.

Note

If a system already has archive units configured, it will fail with the following error message: "Cannot add an archive unit to this system: The system already has an archive unit."

```
filesys archive unit del archive-unit
```

Delete a specific retention unit and change the state of disks or LUNs to available. A retention unit can be deleted from the tier only when the Data Domain system file system is disabled. Disabling the file system stops all Data Domain system operations, including `filesys clean`. Role required: admin.

This command destroys all data in the retention unit. Files within the retention unit must be deleted to remove them from the namespace.

Note

This command option is not available on a Retention Lock Compliance system.

```
filesys archive unit expand archive-unit
```

Expand the size of the specified retention unit and the size of a sealed unit. The system expands the storage space in the retention tier to the maximum supported size on the platform. Role required: admin, limited-admin.

```
filesys archive unit list [archive-unit | all]
```

List all retention units or a specific retention unit. Role required: admin, limited-admin, user, backup-operator, security, none.

```
filesys archive unit unseal [archive-unit-name]
```

Unseals the last sealed unit and only works when there is no target unit. This command can only be run when the file system is offline. Role required: admin, limited-admin.

filesys clean

```
filesys clean reset {schedule | throttle | all}
```

Reset the clean schedule to the default of Tuesday at 6 a.m. (tue 0600), the default throttle of 50 percent, or both. Role required: admin, limited-admin.

```
filesys clean set schedule { never | daily time | <day(s)> time  
| biweekly day time | monthly <day(s)> time }
```

Set schedule for the clean operation to run automatically. Data Domain recommends running the clean operation once a week to maintain optimal availability of the file system. However, if there is no shortage of disk space you may clean less often. Role required: admin, limited-admin.

Argument Definitions

never

Turn off the clean schedule.

daily

Run command every day at the set time.

time

Time is 24-hour format and must be specified by four digits. The time mon 0000 is midnight between Sunday night and Monday morning. 2400 is not a valid time. A new set schedule command cancels the previous setting.

biweekly

Run command on alternate weeks. Bi-weekly cleaning is recommended for file migration on systems with Extended Retention.

monthly

Starts command on the day or days specified at the set time. Days are entered as integers from 1 to 31.

day(s)

Runs on the day or days specified. Days are entered as integers from 1 to 31.

Example 84

To run the clean operation automatically every Tuesday at 4 p.m.: # `filesys clean set schedule tue 1600`

Example 85

To set file system cleaning to run on alternate Tuesdays at 6:00 a.m., enter: # `filesys clean set schedule biweekly "tue" "06:00"`

Example 86

Example 86 (continued)

To run the operation more than once in a month, set multiple days in a single command. For example, to clean the file system on the first and fifteenth day of the month at 4 p.m., enter: `# filesys clean set schedule monthly 1,15 1600`

```
filesys clean set throttle percent
```

Set clean operations to use a lower level of system resources when the Data Domain system is busy. At zero percent, cleaning runs slowly or not at all, depending on how busy the system is. At 100 percent, cleaning uses system resources in the standard way. Default is 50 percent. When the Data Domain system is not running backup or restore operations, cleaning runs at 100 percent. Range: max: 100, min: 0. Role required: admin.

Example 87

To set the clean operation to run at 30 percent of its potential speed: `# filesys clean set throttle 30`

```
filesys clean show config
```

Display settings for file system cleaning. All users may run this command option. Role required: admin, limited-admin, user, backup-operator, security, none.

```
filesys clean show schedule
```

Display current date and time for the clean schedule. All users may run this command option. Role required: admin, limited-admin, user, backup-operator, security, none.

```
filesys clean show throttle
```

Display throttle setting for cleaning. All users may run this command option. Role required: admin, limited-admin, user, backup-operator, security, none.

```
filesys clean start
```

Start clean process manually. When the process finishes, a message is sent to the system log citing the percentage of available storage space. Role required: admin, limited-admin.

```
filesys clean status
```

Display status of the clean process. The system displays a message if cleaning was aborted because collection replication is initializing. Role required: admin, limited-admin.

```
filesys clean stop
```

Stop the clean process. Stopping the process means all progress is lost. Restarting the process means starting from the beginning. Role required: admin, limited-admin.

If the clean process slows down the system, run the `filesys clean set throttle` command to change the amount of system resources used by the clean process. Changes to system resource usage take effect immediately. Role required: admin, limited-admin.

```
filesys clean watch
```

Monitor the `filesys clean` process. Output of this command continuously updates as the `filesys clean` operation progresses. For example, output of verification phase shows the actual number of files moved to the target. Reporting concludes after the final phase. Role required: admin, limited-admin.

Press Ctrl-C to stop monitoring. Note the `filesys clean` process continues to run. All users may run this command.

Note

Because some files may be dropped during verification, output of the percent completion phase may not reach 100 percent. This is expected behavior.

filesys create

`filesys create`

Create a file system or associated RAID disk group with available and spare storage in the active tier. Change the state from Available to In Use. Role required: admin, limited-admin.

Note

This command fails if the user tries to expand the file system beyond the maximum supported active tier size.

Note

If file system creation fails, and the system is unusable, the system displays the following message: "File system creation failed. Contact Data Domain Support."

filesys destroy

`filesys destroy [and-zero]`

Delete all data in the Data Domain system file system including data configured with Retention Lock Governance, remove Replicator configuration settings, and return file system settings to defaults. When this process is finished, NFS clients connected to the Data Domain system may require a remount. Role required: admin.

Note

This command option is not available on a Retention Lock Compliance system.

By default, this command only marks the file system data as deleted. Disks are not overwritten with zeroes unless you specify the `and-zero` option. File system data marked deleted cannot be recovered, even if the disks have not been overwritten with zeroes. The `and-zero` option adds several hours to the destroy operation.

filesys disable

`filesys disable`

Stops file system operations. If no file system exists, the system displays a message confirming that there is no file system to disable. Role required: admin, limited-admin.

filesys enable

`filesys enable`

Start the file system operations. On systems configured with Retention Lock Compliance, security officer authorization is required if there is time skew in the system clock. See the section on Retention Lock Compliance in the *Data Domain Operating System Administration Guide* for details. Role required: admin, limited-admin.

filesys encryption

```
filesys encryption abort-apply-changes
```

Abort a previously issued `apply-changes` request. This applies to both the active and the retention tiers if DD Extended Retention is enabled. If an `apply-changes` operation is already in progress, the abort request will *not* abort the running operation, which will be allowed to finish. Role required: admin.

Note

If an archive unit is offline and file system is enabled, the system displays a warning that an archive unit is offline and that the `abort-apply-changes` operation will not be applied to this unit.

```
filesys encryption algorithm reset
```

Reset the algorithm to the default (`aes_256_cbc`). After running this command, you must restart the file system with `filesys restart` for the change to take effect. Role required: admin, limited-admin.

```
filesys encryption algorithm set {aes_128_cbc | aes_256_cbc |
aes_128_gcm | aes_256_gcm}
```

Select the encryption algorithm. The `aes_256_gcm` option (AES in the Galois/Counter mode) is the most secure algorithm, but is significantly slower than Cipher Block Chaining (CBC) mode. After running this command, you must restart the file system with `filesys restart` for the change to take effect. Role required: admin, limited-admin.

```
filesys encryption algorithm show
```

Display the encryption algorithm. Output indicates changes are pending, if applicable. Role required: admin, limited-admin, user, backup-operator, security, none.

```
filesys encryption apply-changes
```

Update the file system with the current encryption configuration. Encryption changes are applied to all data in the file system *active* tier during the next *cleaning cycle* and to the file system *retention* tier (if DD Extended Retention is enabled) during the next *space reclamation cycle*. This command does not affect cloud units. Role required: admin, limited-admin.

Note

This process can take a long time to complete depending on the size of the data to be re-encrypted.

Note

If an archive unit is offline and file system is enabled, the system displays a warning that the archive unit is offline and that the `apply-changes` operation will not be applied to this unit.

```
filesys encryption disable [tier active | cloud-unit {unit-name
| all}]
```

Deactivate encryption. Disabling encryption means that new data does not get encrypted. You can then run `apply-changes` to decrypt the existing encrypted data. After running this command, you must restart the file system with `filesys restart` for the change to take effect. Role required: admin, limited-admin.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption embedded-key-manager keys create
```

Create a new key. An alert is raised when the new key is generated. You must run `filesys restart` for the key to be used to encrypt/decrypt any new data that is ingested. Role required: admin, limited-admin.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption embedded-key-manager reset key-rotation-policy
```

Reset key rotation policy of the embedded key manager. The `reset` command resets the key rotation policy to none. The new keys are not created automatically. Role required: admin.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption embedded-key-manager set key-rotation-policy {months | none}
```

Set the key rotation policy of the embedded key manager. The embedded key manager supports a maximum of 254 keys. The argument *months* is an integer between 1 and 12, which is the key rotation period. Each rotation creates a new key, which takes effect after the file system is restarted. If specifying *none*, the results are the same as those of `filesys encryption embedded-key-manager reset key-rotation-policy`. Role required: admin, limited-admin.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption embedded-key-manager show
```

Show the configuration of both the running and the configured embedded key manager. Role required: admin, , limited-admin, user, backup-operator, security, none.

```
filesys encryption enable [tier active | cloud-unit {unit-name | all}]
```

Activate encryption for new data written to the file system. After running this command, you must restart the file system with `filesys restart` for the change to take effect. Role required: admin, limited-admin.

```
filesys encryption key-manager disable
```

Stops the Data Domain system from using the RSA DPM (Data Protection Manager) server for key management. You must restart the file system after running the `filesys encryption key-manager disable` command. (The system will start using the embedded key manager after you run `filesys restart`.) The file system continues to use the latest Activated-RW key for encrypting the data. Role required: admin, limited-admin.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption key-manager enable
```

Enable key management. The RSA DPM (Data Protection Manager) key manager is available for external encryption key management. The local encryption key (which is the embedded key manager) administration method is also available. The RSA DPM key manager enables the use of multiple, rotating keys on a Data Domain system. RSA DPM supports a maximum of 254 keys. See the *Data Domain Operating System Administration Guide* for additional information. Role required: admin, limited-admin.

```
filesys encryption key-manager keys create
```

Creates a new active key in the KeySecure external key manager. The KeySecure key manager enables the use of multiple, rotating keys on a Data Domain system. See the *Data Domain Operating System Administration Guide* for additional information. Role required: security.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption key-manager keys modify {<key-id> | muid <key-muid>}
```

Modify the state of an existing key in the KeySecure key manager to a deactivated state. See the *Data Domain Operating System Administration Guide* for a description of key states. Role required: security.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption key-manager reset
```

Clear the attributes of the key-manager. Role required: admin, limited-admin.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption key-manager reset [key-rotation-policy]
```

Change the existing KeySecure external key manager key rotation policy. See the *Data Domain Operating System Administration Guide* for more information. Role required: security.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption key-manager set {server server-name | port
port-number | fips-mode {enabled | disabled} | key-class key-
class | server server-name port port-number | server-type
{keysecure | rkm} | kmip-user user-id fips-mode {enabled |
disabled} key-class key-class} server-type {keysecure | rkm}
kmip-user user-id
```

Specify the attributes of the key-manager. For more details about configuring and setting up the RSA DPM server, see the *Data Domain Operating System Administration Guide*. Role required: admin, limited-admin.

Note

RSA DPM (Data Protection Manager) supports the Key Class Cipher attributes Key Size, Algorithm, and Mode. Data Domain does not use the RSA DPM attributes Algorithm and Mode. These attributes are configured using `filesys encryption algorithm set {aes_128_cbc | aes_256_cbc | aes_128_gcm | aes_256_gcm}`. For Key Class, RSA DPM attribute “Get Key Behavior” has the choices “New Key Each Time” or “Use Current Key”, however, Data Domain supports only “Use Current Key”. When setting up a Key Class in RSA DPM for use in Data Domain, the Key Size must be 256 bits; otherwise the RSA DPM configuration will fail. An error message is not issued if the Key Class is incorrectly configured to generate a new key each time, but the Data Domain system will not receive the correct key to encrypt the data. For more details about configuring and setting up the RSA DPM server, see the *Data Domain Operating System Administration Guide*.

```
filesys encryption key-manager set key-rotation-policy {every
<n> {weeks | months} | none}
```

Set a key rotation policy on the KeySecure external key manager. 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). See the *Data Domain Operating System Administration Guide* for more information. Role required: security

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption key-manager show
```

Display details about the key manager. See the *Data Domain Operating System Administration Guide* for descriptions of key states. Role required: admin, limited-admin, user, security none.

```
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
```

```
filesys encryption keys delete {key-id | muid key-muid} [tier
{active | archive} | archive-unit unit-name]
```

Delete a specified encryption key from the file system, tier, or retention (archive) unit. Only a Destroyed-Compromised key or a Destroyed key can be deleted. A key can be deleted only if no data is currently encrypted with the key. By default, the key is deleted from entire system. Role required: admin, limited-admin.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

Note

If an archive unit is offline and file system is enabled, the system displays a warning that the archive unit is offline and that the encryption keys delete operation will not be applied to this unit.

```
filesys encryption keys destroy {key-id | muid key-umid} [tier
{active | archive} | archive-unit unit-name]
```

Mark a specified encryption key, from the file system, tier, or retention (archive) unit, to be destroyed. After an encryption key is destroyed, the system creates a flag for a re-encrypt operation and it is carried out the next time `filesys clean` runs. By default, the key is marked for destroy from the entire system. If DD Extended Retention is enabled, SREC (Space Reclamation) is responsible for re-encrypting data on the retention (archive) tier.

The key destroy operation simply flags a key to be destroyed, but it will not take effect right away because there is still data encrypted with it. There is no explicit re-encryption command; that job is scheduled when a key is marked to be in the compromised or destroyed state. Role required: admin, limited-admin.

Note

The re-encryption operation may start in the future and may take a long time depending on how much data needs to be re-encrypted. Use `filesys encryption status` to check the status.

Note

If an archive unit is offline and file system is enabled, the system displays a warning that archive unit is offline and that the encryption keys destroy operation will not be applied to this unit.

```
filesys encryption keys export
```

Export encryption keys. This applies to keys in both the active and the retention tiers if DD Extended Retention is enabled. All encryption keys in the file system are exported to a file that can recover encryption keys in the system if required. The key file is passphrase encrypted, and you will be prompted for a passphrase. To protect the key file, you may enter a new passphrase that differs from the Data Domain system passphrase. Lost or forgotten passphrases cannot be recovered.

If possible:

- Run this command when a new key is created or when a change of state occurs to any of the existing keys.
- Send the exported file via FTP for storage in a secure location, accessible to authorized users only.

Role required: admin, limited-admin.

Note

This command cannot export keys for an offline archive unit.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesystems encryption keys show [tier {active | archive | cloud} |
archive-unit archive-unit-name | cloud-unit cloud-unit-name]
Display information about encryption keys, from the file system, tier, or retention
(archive) unit, including key id, key MUID, key state, and the amount of data
encrypted with each key. Information about all keys in the system is displayed by
default. Role required: admin, limited-admin.
```

Note

This command cannot display key information for an offline unit.

```
filesystems encryption keys show summary
Display summary information for keys on the system. Information is displayed for both
the active and the retention tiers if DD Extended Retention is enabled. Role required:
admin, limited-admin, security.
```

```
# filesystems encryption keys show summary
Key
## MUID Active Tier Retention Tier
----- post-comp size post-comp size
1 164 0
2 7cf9acdbdc28cafe9693b06a7a45876c90663a62df64ef480a929f655030492e 0
3 a9f1571edbd4d6b1129c3267f47b03c46645229cdaf1186bd0fce60d17f3445e 72.00 MiB
-----
* Active Tier post-comp size is based on last cleaning of 2014/07/22 06:00:51.
* Retention Tier post-comp size is based on last space reclamation of 2014/07/22 06:00:51.
```

```
filesystems encryption keys sync
Synchronize the key manager encryption keys. An alert is generated if a new key is
detected. When the file system is restarted, the new key is used for reading and
writing. Role required: admin, limited-admin.
```

Note

A Data Domain system retrieves a key from RSA DPM (Data Protection Manager) by Key Class. Choices for the RSA DPM attribute Get Key Behavior of a Key Class are “New Key Each Time” or “Use Current Key.” Data Domain supports only the Key Class “Use Current Key”. An error message is not issued if the Key Class is incorrectly configured to generate a new key each time; however, the Data Domain system does not receive the correct key to encrypt the data.

```
filesystems encryption lock
```

Note

Before locking the system, you must (1) verify that there are no keys in a compromised state, (2) perform a file system clean (`filesys clean`), and (3) disable the file system (`filesys disable`).

Lock the system by creating a new system passphrase and destroying the cached copy of the current passphrase. This command is useful when preparing a Data Domain system and its external storage devices for shipment. There is only one passphrase for each Data Domain system. After running this command, the system encryption keys are unrecoverable until the system is unlocked with the system passphrase. A new system passphrase is not stored and can be forgotten. It is recommended that you keep a record of the passphrase in a safe location. Data cannot be recovered without the new passphrase. Role required: admin, limited-admin.

Note

This command requires security officer authorization. To enable run-time security officer authorization, login as a security officer, and then enter:

```
# authorization policy set security-officer enabled
```

```
filesys encryption show
```

Check the status of the encryption feature. The system displays the configured key-manager if it is different than the currently running key-manager. Role required: admin, limited-admin, user, backup-operator, security, none.

Example 88 Encryption enabled – no key-manager configured

```
# filesys encryption show
Encryption is enabled
The filesystem is unlocked
Algorithm: aes_256_cbc

Key manager in use:      Embedded Key Manager
Key rotation period:    not-configured
Last key rotation date: N/A
Next key rotation date: N/A
```

Example 89 Key rotation policy set for embedded key manager

```
# filesys encryption show
Encryption is enabled
The filesystem is unlocked
Algorithm: aes_256_cbc

Key manager in use:      Embedded Key Manager
Key rotation period:    2 months
Last key rotation date: N/A
```

Example 90 RSA DPM Key Manager configured and enabled

```
# filesys encryption show
Encryption is enabled
The file system is unlocked
```

Example 90 RSA DPM Key Manager configured and enabled (continued)

```
Algorithm:          aes_256_cbc
Key manager in use: RSA_Key_Manager
Server:            dpmsrv1.mydomain.com
Port:              443
Fips-mode:         enabled
Status:            Online
Key-class:         TestKeyClass1
```

```
filesys encryption status
```

Display status of apply-changes and re-encryption operations. The status is displayed for any enabled tiers. The re-encryption operation is performed when a key is destroyed or marked as compromised, and the data encrypted with such a key needs to be encrypted with the current active key. The operation status can be *none* (no operation is needed), *pending*, *running* (in progress), or *done*. Role required: admin, limited-admin.

Note

- Apply-changes and re-encryption operations are not supported on cloud units.
- This command does not show the status of system-level encryption. It is possible for system-level status to be enabled while active tier and cloud tier encryption are disabled. In this situation, issuing the `filesys encryption enable` command indicates that the encryption feature is already enabled.

Example 91

```
# filesys encryption status
Active Tier:
  Encryption enabled: no
  Apply-changes status: none
  Re-encryption status (for compromised or destroyed keys): none
Cloud Tier:
  Cloud unit: cloud_unit_1
  Encryption enabled: no
```

```
filesys encryption unlock
```

Unlock the file system. The system could be locked for several reasons: it is locked automatically after a headswap or chassis swap, or it could have been locked using `filesys encryption lock`. The system will prompt you for a passphrase. Role required: admin, limited-admin.

Note

This command requires security officer authorization. To log in as a security role, enter:

```
# authorization policy set security-officer enabled
```


Argument Definitions

fips-mode

Indicates whether the imported certificate for key management is FIPS (Federal Information Processing Standards) compliant. The default is `enabled`.

key-class

The key class configured on the RSA DPM (Data Protection Manager) server for the Data Domain system. The Key Class name must be enclosed in single or double quotes if the name contains special characters, such as a comma or a space.

key-id

The identifier for a specific key.

muid

The MUID (manufacturer unique identifier) for a specific key,

port

The port number of the RSA server on which the key manager is listening.

server

The name of the RSA DPM (Data Protection Manager) key manager server or IP address.

tier | archive-unit

For systems enabled with DD Extended Retention, the particular tier [active or retention (archive)] or retention (archive) unit.

filesys expand

```
filesys expand
```

Increase the file system by using all available space in the active tier. Role required: admin, limited-admin.

Note

This command fails if the user tries to expand the file system beyond the maximum supported active tier size.

filesys fastcopy

```
filesys fastcopy [retention-lock] source src destination dest
```

Copy a file, an MTree, or directory tree from a Data Domain system source directory to another destination on the Data Domain system. Role required: admin, limited-admin, backup-operator, security.

Note

Backup-operators cannot copy an MTree using `filesys fastcopy`.

Source names *src* that include spaces or special characters must be entered according to the following conventions.

- Enclose the entire source pathname with double quotation marks:

```
filesys fastcopy source "/data/coll/backup/.snapshot/fast copy"
destination /data/coll/backup/dir
```

OR

- Enter a backslash before the space. Do not add quotation marks:

```
filesys fastcopy source /data/coll/backup/.snapshot/fast\ copy
destination /data/coll/backup/dir2
```

Argument Definitions

retention-lock

Use this argument to propagate the retention lock attributes (worm attributes) for files. Use the retention-lock argument when Retention Lock is required on the destination, if the destination file exists (and cannot be overwritten), and when Retention Lock attributes are set per file. When you use the retention-lock argument, the command output varies in the following situations:

- When both the source and destination MTrees are Retention Lock enabled, Retention Lock attributes are copied.
- When the source MTree is not Retention Lock enabled, the system displays a warning message but allows the file copy.
- When the destination MTree is not Retention Lock enabled, the system displays a warning message but still allows the file to copy, and the Retention Lock attributes are not copied.

DD Retention Lock Governance Edition is supported for on-premises and cloud-based DD VE instances. DD Retention Lock Compliance Edition is not supported for on-premises or cloud-based DD VE instances.

source *src*

The location of the directory or file to copy. The first part of the path must be /data/coll/.

destination *dest*

The destination for the directory or file being copied. The first part of the path must be /data/coll/. If the destination already exists, you will be asked if you want to overwrite it.

filesys option

```
filesys option disable report-replica-as-writable
```

Set the reported read/write status of a replication destination file system to read-only. Use the `filesys disable` command before changing this option and use the `filesys enable` command after changing the option. Role required: admin, limited-admin.

With CIFS, use the `cifs disable` command before changing the option and use the `cifs enable` command after changing the option. Role required: admin, limited-admin.

```
filesys option enable report-replica-as-writable
```

Enable the filesys option. Role required: admin, limited-admin.

Set the reported read/write status of a replication destination file system to read/write. Use the `filesys disable` command before changing this option and use the `filesys enable` command after changing the option.

With CIFS, use the `cifs disable` command before changing the option and use the `cifs enable` command after changing the option.

```
fileSYS option reset {local-compression-type | low-bw-optim |
marker-type | report-replica-as-writable | staging-reserve |
staging-clean | staging-delete-suspend | compute-segfeatures |
app-optimized-compression | warning-space-usage <50-90>|
critical-space-usage <75-98>}
```

Return file system compression to the default settings on the destination Data Domain system. Role required: admin, limited-admin.

Argument Definitions

local-compression-type

Reset the compression algorithm to the default of lz.

low-bw-optim

This option is available only to authorized Data Domain and partner support personnel.

marker-type

Return the marker setting to the default of auto.

report-replica-as-writable

Reset the file system to read-only.

staging-clean

Staging-clean: Controls the automatic start of a cleaning operation after files have been deleted. Specify this as a percentage of the reserve. For example, if the staging reserve is 20% and staging-clean is 80%, then the system will start a cleaning operation when the space to be recovered from deleted files exceeds 16% of the total space. Default 0, range 0-200.

staging-delete-suspend

Intended to prevent runaway deletions. For example, when no more reserve is available to increase available space and the client software keeps deleting files hoping to free up space. Specify as a percentage of the reserve. When the specified amount of space has been freed by deletions, the system allows no further deletions until after a clean is started. Default 0, range 0-400.

compute-segfeatures

This option is available only to authorized Data Domain and partner support personnel.

staging-reserve

Set staging reserve percentage from 0 to 90.

app-optimized-compression

Reset the Oracle Optimized Deduplication to none.

warning-space-usage *50-90*

The system can alert you with a warning message when a percentage (50-90%) of the available space is used. Set the percentage using this argument.

critical-space-usage *75-98*

The system can alert you with a critical message when a percentage (75-98%) of the available space is used. Set the percentage using this argument.

```
filesys option set app-optimized-compression {none | oracle1}
```

When set to "oracle1", the system enables Oracle Optimized Deduplication for Oracle Incrementally Updated backups and RMAN multiplexed backups. Please refer to the *Data Domain and Oracle Incrementally Updated Backup Integration Guide* and the *Data Domain Storage Best Practice Guide: Optimized Oracle Incrementally Updated Backup* documents before changing this setting. Role required: admin.

```
filesys option set critical-space-usage 75-98
```

Set critical space usage percentage. Role required: admin, limited-admin.

Note

It is recommended that you set the critical-space-usage percentage higher than the warning-space-usage percentage.

```
filesys option set local-compression-type {none | lz | gzfast | gz}
```

Set compression type. Role required: admin, limited-admin.

```
filesys option set staging-reserve percent
```

Reserve a percentage of total disk space for disk staging. Range: 0 to 90. Role required: admin, limited-admin.

```
filesys option set warning-space-usage 50-90
```

Set warning space usage percentage. Range: 50 to 90. Role required: admin, limited-admin.

Note

It is recommended that you set the warning-space-usage percentage lower than the critical-space-usage percentage.

```
filesys option show [local-compression-type | low-bw-optim |
marker-type | report-replica-as-writable | staging-reserve |
staging-clean | staging-delete-suspend | compute-segfeatures |
app-optimized-compression | warning-space-usage | critical-
space-usage]
```

Show the file system option settings. By default, all file system options are displayed. To limit the output to a single system option, specify one of the system options. Role required: admin, limited-admin, user, backup-operator, security, none.

Argument Definitions

local-compression-type

Display the current compression algorithm.

marker-type

Display the current marker setting.

low-bw-optim

This option is available only to authorized Data Domain and partner support personnel.

report-replica-as-writable

Display the current reported setting on the destination Data Domain system.

staging-reserve

Set staging reserve percentage from 0 to 90.

staging-clean

This option is available only to authorized Data Domain and partner support personnel.

staging-delete-suspend

This option is available only to authorized Data Domain and partner support personnel.

compute-segfeatures

This option is available only to authorized Data Domain and partner support personnel.

app-optimized-compression

Display the Oracle Optimized Deduplication settings enabled.

warning-space-usage 50-90

The system can alert you with a warning message when a percentage (50-90%) of the available space is used. Set the percentage using this argument.

critical-space-usage 75-98

The system can alert you with a critical message when a percentage (75-98%) of the available space is used. Set the percentage using this argument.

filesys report

```
filesys report generate file-location path {path-name | all}
[output-file filename]
```

Create a report showing the name and location of each file under the specified path. If you specify `output-file filename`, the report is saved in this file under the fixed directory `/ddvar`. If the output file argument is not specified, the report is displayed in standard output. The command returns before the entire report is generated, and a footer indicates that the report is complete. Each line in the report contains a file name and its location. The location is shown as *Active* if the file completely resides in the active tier. If the file resides partially or completely in the retention tier or cloud tier, the retention unit or cloud unit name is shown for its location. An asterisk is appended to the line if the file contents span the active tier and retention unit or cloud unit. Role required: admin, limited-admin.

Example 92

To report files in a directory:

```
# filesys report generate file-location path /backup/dir1 output-file
report.txt
```

or

```
# filesys report generate file-location path /data/col1/mtree-2/dir3
output-file report.txt
```

To report files in an MTree:

Example 92 (continued)

```
# filesys report generate file-location path /data/col1/mtree-2
output-file report.txt
```

To report files in the entire namespace:

```
# filesys report generate file-location path all output-file
report.txt
```

filesys restart

```
filesys restart
```

Disable and enable the file system in a single operation. The system displays a message that the file system will be restarted and that applications may experience interruptions. Role required: admin, limited-admin, user, backup-operator.

filesys show

```
filesys show compression [filename] [recursive] [last n {hours
| days}] [no-sync]
filesys show compression [tier {active | archive | cloud}]
summary | daily | daily-detailed {[last n {hours | days | weeks
| months}] | start date [end date]}
```

These command options display the space used by, and compression achieved for, files and directories in the file system. Information is also shown for the tiers supported by the system. Values are reported in Gigabytes (GiB). See the *Data Domain Operating System Administration Guide* for details. Role required: admin, user, backup-operator, security, none.

In general, the more often a backup procedure is run on a file or file system, the higher the compression. The output does not include global and local compression factors for the Currently Used table, but uses a dash instead. Output for a busy system may not return for several hours, depending on the number of files. Other factors may influence the output display.

Running the command without arguments generates default output that shows a summary of compression statistics for all files and directories in the file system for the last 7 days and the last 24 hours. Output includes details on active and retention tiers for systems with Extended Retention only.

Argument Definitions

recursive (Optional)

Display all files in all subdirectories as well as compression information for each file.

filename (Optional)

Synchronize all modified files to disk and then display compression statistics for the specified file or directory only. To display compression statistics for a specific file or directory without first synchronizing all modified files to disk, include the no-sync option.

Depending on the number of files in the file system, specifying a file name could cause this command to process for several hours before completing.

no-sync (Optional)

Use to not sync the file system prior to getting compression information.

tier {active | archive | cloud} (Optional)

Display results for the specified tier.

last *n* {hours | days | weeks | months} (Optional)

In the summary portion of the output, display file system compression statistics for the specified time frame instead of the past 7 days. The statistics for the last 24 hours remain in the summary output. If you specify a file or directory name, you cannot use this option with the `weeks` keyword or the `months` keyword.

summary (Optional)

Display all compression statistics, summarized in the following categories:

- Storage currently used.
- Data written in the last 7 days. By including the `last n` option or the `start date` option, you can display statistics for a different time frame.
- Data written in the last 24 hours.

daily (Optional)

In addition to the summary output, display the following information for each day, over the previous four full weeks, plus the current partial week. This option is not available if you specify a file or directory name.

daily-detailed (Optional)

Display the daily output, but also include the following information for each day. This option is not available if you specify a file or directory name.

start *date* (Optional)

In the summary portion of the output, display file system compression statistics for the time frame that begins on the specified day instead of the past 7 days. The statistics for the last 24 hours remain in the summary output. If you specify a time frame less than the previous 4 weeks, plus the current full week, the daily or daily-detailed output (if specified) is truncated to the shorter time frame.

Specify *date* using the format `yyyy-mm-dd`. By default, the last day of the time frame specified with this argument is the most recent, full day elapsed.

end *date* (Optional)

Valid only if the `start` option is used. In the summary portion of the output, display file system compression statistics for the time frame that ends on the specified day. In general, the more often a backup is done for a particular file or file system, the higher the compression. On a busy system, this process may not complete for several hours, depending on the number of files. Other factors may also affect results.

On a standard Data Domain system, output includes information on active tier only.

Output Definitions**Pre-Comp**

Data written before compression.

Post-Comp

Storage used after compression.

Global-Comp Factor

Ratio of Pre-Comp / (size after global compression). Not applicable to the storage currently used.

Local-Comp Factor

Ratio of (size after global compression)/Post-Comp. Not applicable to the storage currently used.

Total-Comp Factor

Ratio of Pre-Comp / Post-Comp.

Reduction %

Percentage value (Pre-Comp - Post-Comp) / Pre-Comp * 100. This is the default output format.

Example 93

```
fileSYS show compression filename [recursive]
```

Displays all files in all subdirectories and prints compression information for each file as well as the summary for *filename*.

```
fileSYS show file-info filename
```

Display detailed information about the specified file. Specify the fully qualified path to the file. Role required: admin, limited-admin.

```
fileSYS show space [tier {active | archive | cloud | total} |
archive-unit {all | unit-name} | cloud-unit {all | unit-name}]
```

Displays the space available to and used by file system resources, including per-unit space usage statistics for sealed and cleaning archive units. Values are reported in gigabytes (GiB). Role required: admin, limited-admin, user, backup-operator, security, none. Output includes:

- If the tier option is specified, the system shows a summary for the entire tier.
- If archive-unit option is specified, the system shows space usage for each unit.
- If none is specified, the system shows summary tables for the active tier, archive, and total.
- The total option is valid only for systems with Extended Retention and, if total is used, the system displays a summary of both active and archive tiers.
- A line displays space information on '/ddvar/core' if a separate partition is mounted there.
- For DD Cloud Tier storage, post-comp size is based on the CLOUDTIER-CAPACITY license and might not match what is reported by the cloud storage provider.

Note

The tier and archive-unit keywords are not supported for Data Domain Virtual Edition (DDVE), so `fileSYS show space` does not accept these keywords.

Note

Keywords tier and archive-unit are mutually exclusive. The user can only specify one or the other but not both.

Output Definitions**Size GiB**

Total storage capacity of a file system resource.

Used GiB

Amount of data stored on a file system resource.

Avail GiB

Amount of free space on a file system resource.

Use%

Ratio of data stored to total capacity, multiplied by 100.

Cleanable GiB

Estimated amount of recoverable free space. Command output displays space availability and usage information for the following file system components:

/data: pre-comp

Amount of virtual data stored on the Data Domain system. Virtual data is the amount of data sent to the Data Domain system from backup servers.

/data: post-comp

Amount of total physical disk space available for data, actual physical space used for compressed data, and physical space still available for data storage. For DD Cloud Tier storage, post-comp size is based on the CLOUDTIER-CAPACITY license and might not match what is reported by the cloud storage provider. Warning messages go to the system log and an email alert is generated when the Use% figure reaches 90%, 95%, and 100%. At 100%, the Data Domain system stops accepting data from backup servers.

/ddvar

Approximate amount of space used by and available to the log and core files. Use this directory to free space in this area, remove old logs and core files. You can also delete core files from the /ddvar/core directory or the /ddvar/ext directory if it exists.

The total amount of space available for data storage 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.

If Use% is always high, use the command option `filesys clean show-schedule` to see how often the cleaning operation is scheduled to run automatically. Use `filesys clean schedule` to run the operation more often.

`filesys show uptime`

Display the amount of time passed since the file system was last enabled. The display is in days, hours, and minutes. Role required: admin, limited-admin, user, backup-operator, security, none.

filesystems status

```
filesystems status
```

Display the state of the filesystem process. If the filesystem was shut down with a Data Domain system command, such as `filesystems disable`, the output display includes the command name in square brackets. Role required: admin, limited-admin, user, backup-operator, tenant-admin, tenant-user, security, none.

filesystems sync

```
filesystems sync
```

Synchronize modified files to disk. Role required: admin, limited-admin, backup-operator, security.

CHAPTER 20

ha

High availability (HA) is a licensed feature that allows one Data Domain controller to failover to a second controller connected to it, and to the same sets of disks if the primary controller experiences a failure. The `ha` command creates, manages, modifies, and removes the HA configuration. See the *Data Domain Operating System Administration Guide* for details.

This chapter contains the following topics:

- [ha change history](#)204
- [ha guidelines and restrictions](#) 204
- [ha create](#)205
- [ha destroy](#)206
- [ha failover](#) 206
- [ha offline](#)206
- [ha online](#)206
- [ha status](#) 207

ha change history

There have been no changes to this command in this release.

ha guidelines and restrictions

- HA is supported on the following Data Domain systems:
 - DD6800
 - DD9300
 - DD9500
 - DD9800
- Both nodes in the HA pair must have the same memory configuration, I/O module configuration, and software version.
- Use fixed IP addresses for node management.
- Use floating IP addresses for data access.
- Clean up unused fixed interfaces on standby node, to allow floating address configuration on the same interface as on the active node.
- When configuring a floating IP address, the links on both nodes must be up and running, or the HA pair go into a degraded state and be unable to fail over. If the switch port connected to the floating IP port is disabled, or if the floating IP link has no carrier, the HA pair remains in a degraded state. The HA pair comes back online automatically after the floating IP link is recovered.
- The Data Domain file system (DDFS) only exists on the node that is originally designated as the primary node.
- The system times on each node must be within 10 seconds of each other.
- The HA configuration must be in the `highly available` state to failover.
- After it is initiated, the failover process may take up to 10 minutes to complete.
- After a failover, the following protocols require a manual restart of any jobs that were in progress at the time of the failover:
 - CIFS
 - NDMP
 - VTL
- After a failover, jobs that were in progress at the time of the failover using the following will resume automatically after the failover:
 - DD Boost over FC
 - DD Boost over IP
 - NFS
- Run the `ha offline` command on the standby node when performing maintenance or rebooting to avoid disruption of FC traffic on the active node. Once the operation for the passive node is complete, run the `ha online` command .
- When the standby node reboots, FC I/O on the active node can be disrupted for up to 10 seconds if the `ha offline` and `ha online` commands are not run on

the standby node. Active VTL backup and restore operations may fail and need to be restarted. DFC operations are expected to recover without user intervention.

- The active node reboots if the release resource cannot be delivered or processed within 10 seconds in situations where `scsitgtd` is in the middle of configuration changes. Performing multiple failover or failback endpoints, or `vport disable` operations hang because of pending I/O on the VHBA queue.
- When removing an HA configuration in FC environments:
 1. Disable all ports and endpoints before running the `ha destroy` command.
 2. After the former standby node reboots, run the `scsitarget endpoint modify all wwpn auto` command to change the WWPNS on the node so they are not the same as the WWPNS on the former active node.
 3. Zone the newly generated WWPNS to the FC fabric.

ha create

Create the HA relationship between two Data Domain systems.

```
ha create peer {<ipaddr> | <hostname>} [ha-name <hostname>]
```

Create an HA relationship between the local system, and the specified peer system. Optionally specify a hostname to use as the top-level HA system name. The local system becomes the primary node, and the specified peer becomes the standby node. This command prompts the user for the `sysadmin` password of the peer system. Both nodes automatically reboot after the command completes successfully. Role required: `admin`, `limited-admin`.

If the `ha-name <hostname>` parameter is not specified, the hostname of the local system becomes the HA system name, and the hostnames of the two nodes are assigned as follows:

- Local node: `<ha-system-name>-p0`
- Peer node: `<ha-system-name>-p1`

The HA name, whether specified with the `ha-name <hostname>` parameter, or generated from the hostname of the local system, must be associated with a valid floating IP address to provide access to the system over the network.

To change the HA name, manually update the new HA name for all backup, recovery, and replication operations that identify the system by the HA name. Complete the following sequence to change the HA name:

1. Run the `net set hostname <host>` command to set the hostname on the local system to the desired HA name.
2. Run the `net set hostname ha-system` command to promote the new local system hostname to become the HA name.

Argument definitions

ha-name

The top-level hostname for the HA pair. This hostname is used to access the HA pair, and is not tied to a specific physical node.

ha destroy

Remove the HA relationship between two Data Domain systems to use each system independently.

```
ha destroy
```

Remove the HA relationship between the primary and secondary nodes to use each system independently. Run this command on the same system where the HA pair was first created. After the destroy operation, the HA pair is broken down into two single-node systems. The file system is preserved on the node where it resides.

After this command is successful, complete the following sequence to use both nodes as independent systems:

1. Disconnect the node without the file system from the storage and the HA interconnect.
2. Disable and reenable the file system on the node where it resides to resume activity on the file system.
3. Connect the node without the file system to new storage.
4. Run the GUI or CLI Configuration Wizard on the node without the file system to configure it with its own storage and file system.

Role required: admin, limited-admin.

ha failover

Manually initiate a failover from the current active node to the standby node.

```
ha failover [go-offline]
```

Manually initiate a failover from the current active node to the standby node. The node being switched from active to standby reboots after the command completes successfully. This command can be run from either node. The `go-offline` option can be run on the current active node to take the node offline after failing over to the standby node. Role required: admin, limited-admin.

ha offline

Take the standby node offline.

```
ha offline
```

Take the standby node offline. This command takes the system out of the `highly available` state, therefore a failover cannot occur if the primary node suffers a failure. This command can only be run on the standby node. Role required: admin, limited-admin.

ha online

Bring an offline standby node back online.

```
ha online
```

Bring the standby node back online. The standby node reboots after this command completes successfully. This command returns the system to the `highly available` state, allowing for failover if the primary node suffers a failure. This command can only be run on the standby node. Role required: admin, limited-admin.

ha status

View details of the HA configuration.

```
ha status [detailed]
```

View the details of the HA configuration.

```
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
-----
```

The detailed option provides the following additional information:

- **Heartbeat:** A protocol between the two nodes to provide realtime status of the HA state and individual node health status.
- **Mirroring:** The process of copying all configuration information to the standby node to make sure it is ready for a failover.
- **Node health:** Summary of the health of the ports and I/O modules on the nodes.

```
HA System name:      apollo-ha2a.datadomain.com
HA System Status:   highly available
Interconnect Status: ok
Primary Heartbeat Status: ok
External LAN Heartbeat Status: not ok
Hardware compatibility check: ok
Software Version Check: ok
Highly Availability Ratio: 98.5%
```

```
Node apollo-ha2a-p0.datadomain.com:
  Role:      active
  HA State:  online
  Node Health: ok
```

```
Node apollo-ha2a-p1.datadomain.com:
  Role:      standby
  HA State:  online
  Node Health: ok
```

```
Mirroring Status:
Component Name  Status
-----
nvram           ok
registry        ok
sms             ok
ddboost         ok
cifs            ok
-----
```

Note

The Mirroring Status information only displays when the `ha status detailed` command is run on the active node.

Role required: admin, limited-admin.

ha

CHAPTER 21

help

The Command Line Interface (CLI) displays two types of help, syntax-only help and command-description help that includes the command syntax.

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.

help

CHAPTER 22

ifgroup

The `ifgroup` command configures and displays information about dynamic interface groups. Command options create interface groups, add and delete interfaces and clients, enable and disable interface groups, assign and unassign replication Mtrees and remote hosts, and display configuration and connection information.

This chapter contains the following topics:

- [ifgroup change history](#) 212
- [ifgroup add](#) 212
- [ifgroup create](#) 213
- [ifgroup del](#) 213
- [ifgroup destroy](#) 213
- [ifgroup disable](#) 213
- [ifgroup enable](#) 213
- [ifgroup option](#) 214
- [ifgroup rename](#) 214
- [ifgroup replication assign](#) 215
- [ifgroup replication unassign](#) 215
- [ifgroup reset](#) 215
- [ifgroup show config](#) 216
- [ifgroup show connections](#) 216

ifgroup change history

There have been no changes to this command in this release.

ifgroup add

```
ifgroup add group_name {interface {ipaddr | ipv6addr} | client
host}
```

Add an interface, client, or both to *group-name* or to the default group. Prior to adding an interface you must create the *group_name* unless the group name is the default group. Role required: admin, limited-admin.

This command provides full ifgroup 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. The default group behaves in the same manner as any other group.

- The group-name “default” is created during an upgrade of a fresh install and is always used if *group_name* is not specified.
- You can 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*.

- 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`, wild cards such as `*.datadomain.com` or `*`, a short name such as `ddbboost`, or IP range of the client (`xx.xx.xx.0/24` for IPv4 or `xxxx::0/112` for IPv6, for example). When a client's source IP address is evaluated for access to the ifgroup, the order of precedence is:
 1. IP address of the connected Data Domain system
 2. Connected client IP range. This host-range check is useful for separate VLANs with many clients where there isn't a unique partial hostname (domain).
 - For IPv4, you can select five different range masks, based on network.
 - For IPv6, fixed masks /64, /112, and /128 are available.
 3. Client Name: `abc-11.d1.com`
 4. Client Domain Name: `*.d1.com`
 5. All Clients: `*`

If none of these checks find a match, ifgroup interfaces are not used for this client.

For detailed information about this order of precedence, see the *Data Domain Boost for Partner Integration Administration Guide*.

- By default, the maximum number of groups is eight. It is possible to increase this number by editing the system registry and rebooting.

Additionally, the IP address must be configured on the Data Domain system and its interface must be enabled. You can add public or private IP addresses for data transfer connections. After adding an IP address as an interface, you can enable advanced load balancing and link failover.

See the *Data Domain Boost for Partner Integration Administration Guide* or the *Data Domain Boost for OpenStorage Administration Guide*, and the *Data Domain Operating System Administration Guide* for more information on interface groups.

ifgroup create

```
ifgroup create group-name
```

Create a group with the name *group-name* for the interface. Group names may contain alphanumeric characters, hyphens, and underscores. System hostnames, fully qualified hostnames, and wildcard hostnames indicated by an asterisk may also be used. Reserved group names that cannot be used are `default`, `all`, or `none`. Role required: admin, limited-admin.

ifgroup del

```
ifgroup del group_name {interface {ipaddr | ipv6addr} | client host}
```

Remove an interface, client, or both from *group_name* or default group. Deleting the last IP address interface disables the ifgroup. If this is the case, you have the option of terminating this command option. Role required: admin, limited-admin.

ifgroup destroy

```
ifgroup destroy group-name
```

Destroy the group name. Only empty groups can be destroyed. Interfaces or clients cannot be destroyed but may be removed sequentially or by running the command option `ddboost ifgroup reset group-name`. Role required: admin, limited-admin.

Note

The group-name “default” cannot be destroyed.

ifgroup disable

```
ifgroup disable group-name
```

Disable a specific group by entering the *group-name*. If *group-name* is not specified, the command applies to the default group. Role required: admin, limited-admin.

ifgroup enable

```
ifgroup enable group-name
```

Enable a specific group by entering the *group-name*. If *group-name* is not specified, the command applies to the default group. Role required: admin, limited-admin.

ifgroup option

```
ifgroup option reset {disable-file-replication | enforce-client-interface}
```

Reset replication permissions for ifgroups and interface enforcement settings to their default settings. Role required: admin, limited-admin.

Changed settings impact all interface groups, but they do not impact in-progress jobs. Changed settings take effect during the ifgroup query at the start of a job.

Example 94

```
# ifgroup option reset disable-file-replication
File replication is allowed on ifgroup.
```

Example 95

```
# ifgroup option reset enforce-client-interface
Client may use any interface.
```

```
ifgroup option set {disable-file-replication | enforce-client-interface}
```

Set replication permissions for ifgroups and interface enforcement settings. By default, ifgroup is enabled for file replication, and interface enforcement is disabled. Role required: admin, limited-admin.

Changed settings impact all interface groups, but they do not impact in-progress jobs. Changed settings take effect during the ifgroup query at the start of a job.

Example 96

```
# ifgroup option set disable-file-replication
File replication is not allowed on ifgroup.
```

Example 97

```
# ifgroup option set enforce-client-interface
Client must use interfaces configured in ifgroup.
```

ifgroup rename

```
ifgroup rename group-name new-group-name
```

Rename the ifgroup *group-name* to *new-group-name*. This command option does not require disabling the group. The default group cannot be renamed. Role required: admin, limited-admin.

ifgroup replication assign

```
ifgroup replication assign group_name {mtree mtree-path |
remote hostname | mtree mtree-path remote hostname}
```

Assign a replication MTree and remote host to *group-name*. The full MTree path is required. Role required: admin, limited-admin.

Note

The *hostname* configuration is case-sensitive; however, this command automatically converts input entered as uppercase to lowercase. At upgrade, all previously configured hostnames are automatically converted to lowercase.

Example 98

```
# ifgroup replication assign 10GLab mtree /data/coll/REPLX remote ddp-880-1.datadomain.com
Assigned replication mtree "/data/coll/REPLX" with remote "ddp-880-1.datadomain.com" to
ifgroup "10GLab".
```

ifgroup replication unassign

```
ifgroup replication unassign group_name {mtree mtree-path |
remote hostname | mtree mtree-path remote hostname}
```

Unassign a replication MTree and remote host from *group-name*. Role required: admin, limited-admin.

Example 99

```
# ifgroup replication unassign 10GLab mtree /data/coll/REPLX remote ddp-880-1.datadomain.com
10GLab
Unassigned replication mtree "/data/coll/REPLX" with remote "ddp-880-1.datadomain.com" from
ifgroup "10GLab".
```

ifgroup reset

```
ifgroup reset [group_name] {all | interfaces | clients |
replication}
```

Reset all, interfaces, clients, or replication for *group-name*. If *group-name* is not specified, the command applies to the default group. Role required: admin, limited-admin.

Example 100

```
# ifgroup reset 10GLab replication
ifgroup "10GLab" is enabled with 3 replication assignments.
This command will remove all replication assignments from the group.
Are you sure? (yes|no|?) [no]: yes

ok, proceeding.

Reset ifgroup "10GLab".
```

ifgroup show config

```
ifgroup show config [group_name] {all | summary | interfaces |
clients | replication}
```

Display the configuration of interfaces, clients, or replication for *group-name*. If *group-name* is not specified, information for all groups is shown. Select the all argument to view all configuration options for the selected group. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 101

```
# ifgroup show config summary
Group-name      Status      Interface    Clients      Replication
-----
default         enabled     1            1            1
v6default       enabled     1            0            1
10GLab-192      enabled     2            2            0
10GLab-172      enabled     4            1            1
10GLab-192-REPL disabled    2            0            2
10GV6-2000      enabled     4            1            0
10GV6-3000      enabled     2            4            0
10GLab-172-REPL enabled     3            0            1
-----
File replication is allowed on ifgroup.
Client must use interfaces configured in ifgroup.
```

ifgroup show connections

```
ifgroup show connections
```

Show connections activity for interface groups. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 102

```
# ifgroup show connections
Group-name      Status      Port          Interface    Client Write  Client Read  Repl-out
Repl-in        Total
-----
(null)          disable    eth0a         10.6.109.41  0             0
0              0
default         enabled    eth0a         10.6.109.40  0             0
0              0
10GLab          enabled    eth4a:1       192.168.1.230 0             0
0              0
10GLab          enabled    eth4b:1       192.168.1.231 0             0
0              0
-----
# ifgroup show connections
Group-name      Status      Port          Interface    Client Write  Client Read  Repl-out  Repl-
in  Total
-----
(null)          disable    eth0a         2620::eaf4   0             0           0
0              0
default         disable    eth0b         2620::eaf5   0             0           0
0              0
10GLab          disable    eth4a         3000::230    0             0           0
0              0
```


Example 102 (continued)

```
10GLab      disable  eth4b      3000::231      0           0           0
0           0
-----
-----
-----
```

ifgroup

CHAPTER 23

ipmi

The `ipmi` command monitors and manages a Data Domain system deployed remotely. Command options enable administrators to monitor remote systems and to power the systems on or off as required. The Serial-Over-LAN (SOL) feature is used to view the serial output of a remote system boot sequence. For more information, including the list of supported models, see the *Data Domain Operating System Offline Diagnostics Suite User's Guide*.

This chapter contains the following topics:

- [ipmi change history](#) 220
- [ipmi guidelines and restrictions](#) 220
- [ipmi config](#) 220
- [ipmi disable](#) 220
- [ipmi enable](#) 220
- [ipmi remote](#) 220
- [ipmi reset](#) 221
- [ipmi show](#) 221
- [ipmi user](#) 221

ipmi change history

There have been no changes to this command in this release.

ipmi guidelines and restrictions

- Users cannot log in to IPMI via SSH. See the *Data Domain Operating System Administration Guide* for instructions on managing remote systems.
- Users cannot log in to BMC instead of IPMI.
- IPMI (on/off/cycle/status) and SOL are not supported on models DD140, DD610, and DD630.

ipmi config

```
ipmi config port {dhcp | ipaddress ipaddr netmask mask gateway ipaddr}
```

Configure an IPMI port to get its IPv4 configuration from DHCP, or configure static IP address information. If configuring a static IP address, you must provide the BMC IP address, netmask, and gateway address. To display a list of IPMI ports, enter `ipmi show hardware` or `ipmi show config`. See the *Data Domain Operating System Administration Guide* for details. Role required: admin, limited-admin.

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.

Note

The BMP port and IPMI implementation do not support IPv6 in this release.

ipmi disable

```
ipmi disable {port | all}
```

Disable IPMI remote access through one or all IPMI ports. To display a list of IPMI ports, enter `ipmi show hardware` or `ipmi show config`. Role required: admin, limited-admin.

ipmi enable

```
ipmi enable {port | all}
```

Enable IPMI remote access through one or all IPMI capable ports. To display a list of IPMI capable ports, enter `ipmi show hardware`. Role required: admin, limited-admin.

ipmi remote

```
ipmi remote console ipmi-target {ipaddr | hostname} user user
```

Activates the Serial-Over-Lan (SOL) feature, which enables viewing text-based serial output of a remote Data Domain system without a serial server. SOL is used in combination with the remote power cycle command to view the remote system's boot sequence.

Specify the IP address or hostname of the remote system, and specify an IPMI username that is configured on the remote system. For more information, see the *Data Domain Operating System Administration Guide*. Role required: admin, limited-admin.

```
ipmi remote power {on | off | cycle | status} ipmi-target
{ipaddr | hostname} user user
```

Power on, power off, or power cycle a remote target system from an initiator system. Specify the IP address or hostname of the remote system. Role required: admin, limited-admin.

ipmi reset

```
ipmi reset
```

Resets the LAN configuration for all IPMI ports, and clears the SOL configuration. Role required: admin, limited-admin.

ipmi show

```
ipmi show config
```

View the configuration of local IPMI interfaces. Output includes the dynamic or static IP address, gateway, netmask, and MAC address. Role required: admin, limited-admin.

```
ipmi show hardware
```

View the port names and firmware version of the local BMC. Output also includes the IPMI version, manufacturer, MAC addresses. The Link Status column shows if the LAN cable is connected to the LAN-IPMI shared port.

Link status cannot be determined on the following Data Domain systems: DD640, DD2200, and DD2500. Role required: admin, limited-admin.

ipmi user

```
ipmi user add user [password password]
```

Add a new local IPMI user. The specified username and password are used by remote systems to access the local system. Role required: admin, limited-admin.

Note

User root is not supported for IPMI connections on DD160 systems.

```
ipmi user change user [password password]
```

Change the password of a locally defined IPMI user. Role required: admin, limited-admin.

```
ipmi user del user
```

Delete a locally defined IPMI user. Role required: admin, limited-admin.

```
ipmi user list
```

View a list of locally-defined IPMI users, including names, IDs, and permissions. Role required: admin, limited-admin.

```
ipmi user reset
```

Clear all locally-defined IPMI users. If you are enabling IPMI for the first time, we recommend running this command to clear IPMI users who may be out of synch between two ports, and to disable default users. Role required: admin, limited-admin.

CHAPTER 24

license

The `license` command adds, deletes, and resets keys for licensed features and storage capacity.

This chapter contains the following topics:

- [license change history](#).....224
- [license guidelines and restrictions](#)..... 224
- [license add](#)..... 224
- [license delete](#)..... 224
- [license reset](#)..... 224
- [license show](#).....225

license change history

There have been no changes to this command in this release.

license guidelines and restrictions

- License codes are case-insensitive. Include the hyphens when entering codes.
- These commands are not supported on Data Domain Virtual Edition.
- The following software options require separate licenses. See the Online Support site for details.
 - DD Boost
 - Extended Retention (formerly “Archiver”)
 - Encryption
 - Expanded Storage
 - I/OS
 - Replication
 - Retention Lock Compliance
 - Retention Lock Governance
 - Shelf Capacity
 - Storage Migration
 - Virtual Tape Library (VTL)
 - High Availability

license add

```
license add license-code [license-code ...]
```

Add one or more licenses for features and storage capacity. Enter the license code exactly as provided by Data Domain, including the dashes. Role required: admin, limited-admin.

license delete

```
license del license-feature [license-feature ...] | license-code [license-code ...]
```

Delete one or more licenses for features or storage capacity. To display the license codes and license feature names, enter `license show`. Role required: admin, limited-admin. Security officer authorization is required to delete Retention Lock Compliance licenses.

license reset

```
license reset
```

Remove all licenses. Requires confirmation before deletion. Role required: admin, limited-admin. Security officer authorization is required to delete Retention Lock Compliance licenses.

Example 103

```
#license reset
This will delete all added licenses.
  Do you want to continue? (yes|no) [no]: yes

All licenses deleted.
```

license show

```
license show [scheme]
```

View license codes, which are also called license keys. Feature licenses also display a feature name, which you can use instead of the code when deleting a feature license. The `scheme` argument displays `unknown`, `DD licensing`, or `elicensing`. Role required: `admin`, `limited-admin`, `security`, `backup-operator`, `user`.

license

CHAPTER 25

log

The `log` command manages and displays the Data Domain system log file. Messages from the alerts feature, the autosupport reports, and general system messages are sent to the `log` directory (`/ddvar/log`). A log entry appears for each Data Domain system command given on the system.

Data Domain systems can send network log messages to other systems enabled to listen. The Data Domain system sends the log in the standard syslog format. When remote logging is enabled, all messages in the `messages` and `kern.info` files are exported.

Message selectors include:

***.notice**

Send all messages at the notice priority and higher.

***.alert**

Send all messages at the alert priority and higher (alerts are included in `*.notice`).

kern.*

Send all kernel messages (`kern.info` log files).

This chapter contains the following topics:

• log change history	228
• log host	228
• log list	228
• log view	228
• log watch	230

log change history

There have been no changes to this command in this release.

log host

```
log host add host
```

Add a remote system hostname to the list of hosts to which system log messages are sent. Role required: admin, limited-admin.

Note

If using three or more remote log hosts, they must be added by entering the IP address in the *host* argument instead of the host name.

```
log host del host
```

Remove a hostname from the list of systems that receive system log messages. Role required: admin, limited-admin.

```
log host disable
```

Disable sending log messages to other systems. Role required: admin, limited-admin.

```
log host enable
```

Enable sending log messages to other systems. Role required: admin, limited-admin.

```
log host reset
```

Disable log sending and clear the list of destination hostnames. Role required: admin, limited-admin.

```
log host show
```

Display whether logging is enabled or disabled and the list of destination hostnames. Role required: admin, limited-admin, security, user, backup-operator, or none.

log list

```
log list
```

List the files in the log directory with the date each file was last modified and the size of each file. For information on the log files, see the *Data Domain Operating System Administration Guide*. Role required: admin, limited-admin, security, user, backup-operator, or none.

log view

```
log view [filename]
```

Display the specified log file. To display the available log files, enter `log list`. If a filename is not specified, the command displays the current messages file.

When viewing the log, use the up and down arrows to scroll through the file. Use the q key to quit. Enter a forward slash to search forward or a question mark to search backward for a pattern such as a date. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
log view access-info [authentication-failures {all | known-
users | unknown-users} | access-history {all | logins |
logouts} | user-management] [user <user-name>] [last <n> {hours
```

```
| days | weeks | months} | start <MMDDhhmm[[CC]YY]> [end
<MMDDhhmm[[CC]YY]>]]
```

Displays a history of user logins and logouts on the system, including both successful and unsuccessful attempts to log in.

```
log view audit-info [authorization-errors | all-errors] [user
<user-name>] | user-role {admin | security | user | backup-
operator | none}} [tenant-unit <tenant-unit>] [host <host>]
[application {CLI | REST | GUI | VDISK}] [string <str>] [last
<n> {hours | days | weeks | months} | start <MMDDhhmm[[CC]YY]>
[end <MMDDhhmm[[CC]YY]>]]
```

Displays a list of all system management configuration changes, and provides the following the following details:

- Username of the user who initiated the configuration change
- Timestamp
- Requested operation
- Operation outcome

log watch

```
log watch [filename]
```

View new log entries for the specified log file as they occur. To display the available log files, enter `log list`. If a filename is not specified, the command displays the messages file entries.

Use Ctrl-C to stop the display. Role required: admin, limited-admin, security, user, backup-operator, or none.

CHAPTER 26

migration

The `migration` command copies all data from one DD system to another. Use this command when upgrading to a larger capacity DD system. Migration is typically performed in a LAN environment.

Migration may also be used to copy replication configurations, known as “contexts.” See the *Data Domain Operating System Administration Guide* for instructions.

This chapter contains the following topics:

• migration change history	232
• migration abort	232
• migration commit	232
• migration receive	233
• migration send	234
• migration show stats	236
• migration status	236
• migration watch	236

migration change history

There have been no changes to this command in this release.

migration abort

```
migration abort
```

Stop a migration process and return the DD system to its previous state. If the migration source is part of a replication pair, you must run `migration abort` on the source, and replication will be restarted. You cannot run this command on the migration destination. After you run `migration abort` on the destination, you must also run `filesys destroy` on the destination before the file system can be reenabled. After running `migration abort`, the password on the destination will be the same as the password on the source. Role required: admin, limited-admin.

migration commit

```
migration commit
```

Limit migration to data received by the source at the time the command is entered. You can use this command anytime after entering `migration send`. After `migration commit`, all data on the source, including new data for contexts migrated to the destination, is sent only to the destination. Write access to the source is blocked after you enter `migration commit` and during the time required to complete migration. After the migration process is finished, the source is opened for write access, but new data is not migrated to the destination. Role required: admin, limited-admin.

Example 104

To migrate data from source hostA to destination hostB (no replication):

1. On hostB (destination), enter:

```
# filesys disable
# filesys destroy
# filesys create
# migration receive source-host hostA
```

2. On either host, enter:

```
# migration send /backup destination-host hostB
```

3. At the appropriate time for your site, create a migration end point. The three migration phases may take many hours. During that time, new data sent to the source is also marked for migration.
4. After the three migration phases are finished, enter the following command on hostA first, and then on destination hostB:

```
# migration commit
```

Example 105

To migrate data and a context from source hostA to destination hostC, when hostA is also a directory replication source for hostB:

Example 105 (continued)

1. On hostC (migration destination), enter:

```
# fileys disable
# fileys destroy
# fileys create
# migration receive source-host hostA
```

2. On hostA (migration and replication source), enter:

```
# migration send dir://hostB/backup/dir2 destination-host hostC
# migration watch
```

3. First on hostA and then on hostC, enter (this command also disables the file system):

```
# migration commit
```

4. On hostB (replication destination), enter the following to change the replication source to hostC:

```
# fileys disable
# replication modify dir://hostB/backup/dir2
  source-host hostC
# fileys enable
```

migration receive

```
migration receive source-host src-hostname
```

Prepare a DD system to be a migration destination. This migration destination:

- Must have an empty file system
- Must have equal or larger capacity than the used space on the migration source (with the exception of collection replication)
- Must have a replication and/or encryption license if the source is licensed for those software options

This command should be run:

- Only on the migration destination
- After running `fileys destroy` and `fileys create` on the migration destination
- Before entering `migration send` on the migration source

Role required: admin, limited-admin.

Example 106

To prepare a destination for migration from the source hostA:

```
# fileys destroy
# fileys create
# migration receive source-host hostA
```

Argument Definitions

src-hostname

The migration source host, which can be a simple host name, an IP address, a partially qualified domain name, or a fully qualified domain name.

migration send

```
migration send {obj-spec-list | all} destination-host dst-hostname
```

Start migration, which will continue until you run `migration commit`.

This command should be run:

- Only on the migration source
- Only when no backup data is being sent to the migration source
- After running `migration receive` on the migration destination

New data written to the source is marked for migration until you run `migration commit` (which should be run first on the source, then the destination). New data written to the source after `migration commit` is not migrated. Write access to the source is blocked from the time you run `migration commit` until the migration process concludes.

Any setting of the system's replication throttle also applies to migration. If the migration source has throttle settings, use `replication throttle set override` to set the throttle to the maximum (unlimited) before starting migration.

With the exception of licenses and key-manager settings, all data on the migration source is always migrated, even when a single directory replication context is specified. Role required: admin, limited-admin.

Note

After you run `migration send`, the migration source remains in read-only mode until all replication contexts are synchronized. To avoid excessive time in this mode, it is recommended that you first synchronize these contexts by running `replication sync` and then run `migration send` immediately after synchronization concludes.

Example 107

To start migration of data only (excluding replication contexts, even if replication contexts are configured) to a migration destination hostC:

```
# migration send /backup destination-host hostC
```

Example 108

To start a migration that includes a collection replication context (replication destination string) of col://hostB:

```
# migration send col://hostB destination-host hostC
```

Example 109

Example 109 (continued)

To start migration with a directory replication context of `dir://hostB/backup/`
`dir2`:

```
# migration send dir://hostB/backup/dir2 destination-host hostC
```

Example 110

To start migration with two replication contexts using context numbers 2 and 3:

```
# migration send rctx://2 rctx://3 destination-host hostC
```

Example 111

To migrate all replication contexts:

```
# migration send all destination-host hostC
```

Example 112

If a migration source has encryption enabled, you must do the following on the destination before starting the migration process.

1. Add the encryption license.

```
# license add license-code
```

2. Enable encryption. This command prompts you for a passphrase. Use the same passphrase as the migration source.

```
# filesys encryption enable
```

3. Restart the file system.

```
# filesys restart
```

4. If the migration source has DPM key manager configured and enabled, clear the DPM attributes on the destination.

```
# filesys disable
# filesys encryption key-manager reset
# filesys restart
```

5. After migration concludes, configure the DPM attributes on the destination to be the same as the DPM attributes on the source, and then enable the DPM key manager.

Argument Definitions***dst-hostname***

The migration destination, which can be a simple host name, an IP address, a partially qualified domain name, or a fully qualified domain name.

obj-spec-list

The specified replication contexts or paths, which can be one of the following:

- For systems that do not have a replication license:

```
/backup
```

- For systems with replication, this argument represents one or more contexts from the migration source. After you migrate a context, all data from the context remains on the source, but the context configuration is moved to the sh destination. Thus, this argument can be:
 - The destination string, as defined when setting up replication, for example:


```
dir://hostB/backup/dir2col://hostBpool://hostB/pool2
```
 - The context number, such as `rctx://2`, as shown in the output from `replication status`
 - The keyword `all`, which migrates all contexts from the migration source to the destination

migration show stats

```
migration show stats
```

Display migration statistics during the migration process. Role required: admin, limited-admin.

Output Definitions

Bytes Received

The total number of bytes received at the destination. On the destination, this value includes data, overhead, and network overhead. On the source, this value includes overhead and network overhead. Use this value (and the **Bytes Sent** value) to estimate network traffic generated by migration.

Bytes Remaining

The total number of bytes remaining to be sent. This information is shown only on the migration source.

Bytes Sent

The total number of bytes sent from the migration source. This value includes backup data, overhead, and network overhead. On the destination, this value includes overhead and network overhead. Use this value (and the **Bytes Received** value) to estimate network traffic generated by migration.

Sync'ed-as-of Time

The last time stamp for which data has been synchronized between the two systems.

migration status

```
migration status
```

Display the status of migration at the time the command is run. Role required: admin, limited-admin.

migration watch

```
migration watch
```

Track the initial phase of migration (when write access is blocked). The command output shows the percentage of the migration process that has been completed. Role required: admin, limited-admin.

migration

CHAPTER 27

mtree

The `mtree` command enables operations on a single “managed tree” (MTree) of a filesystem. An MTree is a logical partition of the namespace in the file system that can group together a set of files for management purposes; for example, snapshot schedules, replication, or retention locking.

This chapter contains the following topics:

- [mtree change history](#) 240
- [mtree create](#) 240
- [mtree delete](#) 241
- [mtree list](#) 241
- [mtree modify](#) 242
- [mtree option](#) 243
- [mtree rename](#) 243
- [mtree retention-lock](#) 244
- [mtree show](#) 246
- [mtree undelete](#) 248

mtree change history

Modified behavior in DD OS 6.1.2

`mtree delete mtree-path`

This command will fail to delete the specified MTree if there is a data movement policy configured on the MTree, or a data movement operation is in progress when the delete command is issued.

mtree create

```
mtree create mtree-path [tenant-unit tenant-unit] [quota-soft-limit n {MiB|GiB|TiB|PiB}] [quota-hard-limit n {MiB|GiB|TiB|PiB}]
```

Create an MTree under the specified path. The format of the *mtree-path* is `/data/coll/mtree-name`. An error message notifies you to enter a different name if another MTree with the same name exists. Role required: admin, limited-admin.

Naming conventions for creating MTrees include uppercase and lowercase letters (A-Z, a-z), numbers 0-9, single, non-leading embedded space, exclamation point, hash, dollar sign, ampersand, caret, tilde, left and right parentheses, left and right brackets, left and right braces.

If no quota option is specified, the default is unlimited for both soft and hard limits, meaning there are no quota limits.

When setting quota limits, a warning appears if the new limit is lower than the current space usage of the MTree. The command does not fail, but subsequent writes to the MTree are rejected. An error message appears if you are setting a soft limit that is greater than or equal to the hard limit. When the hard limit is reached for an MTree quota, write operations stop and no more data can be written to the MTree. Data can be deleted.

Argument Definitions

mtree-path

Displays MTrees under a specified path only.

tenant-unit (Optional)

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system.

Example 113

To create MTree `/data/coll/backup1` with no quota limits:

```
# mtree create /data/coll/backup1
```

Example 114

To set a soft limit quota of 10 GiB on MTree `/data/coll/backup1`:

```
# mtree create /data/coll/backup1 quota-soft-limit 10 GiB
```


Example 115

To set a hard limit quota of 10 TiB on MTree `/data/col1/backup2`:

```
# mtree create /data/col1/backup2 quota-hard-limit 10 TiB
```

Example 116

To set a tenant-unit on `/data/col1/backup3`:

```
# mtree create /data/col1/backup3 tenant-unit tenant1
```

mtree delete

```
mtree delete mtree-path
```

Delete the specified MTree (denoted by the pathname). MTrees marked for deletion remain in the file system until the `filesys clean` command is run. This command option is not allowed on Retention Lock Governance or Retention Lock Compliance MTrees unless they are empty. You can revert the marked-for-deletion state of that MTree by running the `mtree undelete` command. See the *Data Domain Operating System Administration Guide* for details on Retention Lock Compliance and Governance. If the MTree is a storage unit, the system returns an error. Role required: admin.

This command will fail to delete the specified MTree if there is a data movement policy configured on the MTree, or a data movement operation is in progress when the delete command is issued.

Note

For systems that use the DD Boost protocol, you can use the `ddbboost storage-unit delete` command to delete a storage unit.

Effects of deleting an MTree include:

- The MTree appears in the output of the `mtree list` command option and is marked with the status value D.
- File service to a deleted MTree is rejected. Deleted MTrees are not visible through NFS or CIFS clients.
- When an MTree is removed from the file system, snapshots associated with that MTree are also deleted from the `/data/col1/mtree-name/.snapshot/` directory.

mtree list

```
mtree list [mtree-path] [tenant-unit tenant-unit]
```

Display the list of MTrees. When Secure Multi-tenancy (SMT) is not enabled, the system displays three columns: Name, Pre-Comp (GiB), and Status. When SMT is enabled, the system also displays Tenant Unit. Role required: admin, limited-admin, user, backup-operator, tenant-admin, tenant-user, security, none.

Argument Definitions

mtree-path (Optional)

Display MTrees under the specified path only. This command supports the asterisk (*) wildcard character in the MTree pathname. Values include:

- /data/col1/mtree1
- /data/col1/mtree*
- *mtree*

tenant-unit (Optional)

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system.

Output Definitions

When SMT is enabled, tenant-unit will be displayed if it is configured. If it is not configured, the system will display "-". Output includes the MTree pathname, pre-compression, and status. Status is based on pre-defined values:

D

Marked for deletion. MTree will be removed from the file system by the `filesys clean` command. Can be unmarked for deletion by using the `mtree undelete` command only if the `filesys clean` command has not been run.

Q

Quota defined.

RO

Read-only access.

RW

Read/write access.

RD

Replication destination.

RLCE

Retention Lock Compliance enabled.

RLGE

Retention Lock Governance enabled.

RLGD

Retention Lock Governance disabled.

mtree modify

```
mtree modify mtree-path tenant-unit tenant-unit name
```

Assign an MTree to a tenant-unit. If the MTree is a storage unit, the system returns an error. Role required: admin, limited-admin.

Note

For systems that use the DD Boost protocol, you can use the `ddboost storage-unit modify` command to modify a storage unit.

Argument Definitions***mtree-path***

Display MTrees under the specified path only.

tenant-unit (Optional)

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system.

tenant-unit name

The name of the tenant-unit you want to associate with the MTree.

mtree option

```
mtree option reset app-optimized-compression mtree mtree_path
```

Reset the Oracle Optimized Deduplication setting on the specified MTree to its default value, none. Role required: admin, limited-admin.

```
mtree option set app-optimized-compression {none | global | oracle1} mtree mtree_path
```

Set Oracle Optimized Deduplication on the specified MTree.

Argument Definitions**none**

Oracle Optimized Deduplication is disabled.

global

The MTree uses the system-level app-optimized-compression value (none or oracle).

oracle1

Oracle Optimized Deduplication is enabled.

```
mtree option show [mtree mtree_path]
```

Display option values for the MTree found at the *mtree_path*. If no MTree is specified, the system displays option values for all MTrees. Role required: admin, limited-admin.

Argument Definitions***mtree_path***

The full path of the MTree in the file system.

mtree rename

```
mtree rename mtree-path new-mtree-path
```

Rename the specified MTree. Note that `/data/coll/backup` cannot be renamed. Retention Lock Governance or Retention Lock Compliance MTrees can only be renamed if they are empty. Role required: admin, limited-admin.

This command option requires security officer authorization if Retention Lock Compliance is enabled on the specified MTree.

mtree retention-lock

```
mtree retention-lock disable mtree mtree-path
```

Disable Retention Lock for the specified MTree. This command option is allowed on Retention Lock Governance MTrees only. It is not allowed on Retention Lock Compliance MTrees. See the *Data Domain Operating System Administrator's Guide* for details on Retention Lock Compliance and Governance. Role required: admin, limited-admin..

```
mtree retention-lock enable mode {compliance | governance}
mtree mtree-path
```

Enable Retention Lock for the specified MTree. Use the compliance argument to meet the strictest data permanence regulatory standards, such as those of SEC17a-4f. Enabling Retention Lock Compliance requires security officer authorization. Role required: admin, limited-admin.

Use the governance argument to propagate the same protection provided in the previous release of DD OS. The level of security protection is lower than Retention Lock Compliance.

When Retention Lock is enabled on an MTree, any file in the MTree may become locked by setting its *atime* to the future. Additionally, renaming a non-empty directory in the MTree is disabled. See the *Data Domain Operating System Administration Guide* for details on Retention Lock Compliance and Governance, and for instructions on setting retention time.

To enable Retention Lock Compliance on an MTree, enter: # `mtree retention-lock enable mode compliance mtree /data/coll/mtree_name`

Note that `/data/coll/backup` cannot be configured for Retention Lock Compliance.

```
mtree retention-lock reset {min-retention-period | max-
retention-period} mtree mtree-path
```

Reset the minimum or maximum retention period for the specified MTree to its default value. The command option is allowed on MTrees with Retention Lock Governance enabled. Role required: admin, limited-admin.

See the *Data Domain Operating System Administration Guide* for details on Retention Lock Compliance and Governance and for instructions on setting retention time.

```
mtree retention-lock report generate retention-details mtrees
{mtree-list | all} [format {text | tsv | csv}] [output-file
filename]
```

Lists all retention-lock files in one or multiple mtrees, their expiration time, mode of retention, and size. If the output-file *filename* option is specified, then the report will be written to `/ddvar/log/debug/retention-lock-reports/filename`; otherwise, the report will go to standard output. The report includes a timestamp indicating the time it was generated. The default output format is text. If the file already exists, an error is generated. Role required: admin, limited-admin.

```
mtree retention-lock revert path
```

Revert all Retention Lock files in a specified path to non-Retention Lock files. If the path points to an MTree, all files within the MTree will be reverted. Note that directories and files within Retention Lock Compliance MTrees cannot be reverted. Role required: admin, limited-admin.

The base of the path must be `/data/coll/mtree-name/` or `data/coll/backup/`.

Reverting Retention Lock Governance generates a Data Domain system alert (at the `Alert` severity level) and logs the names of the reset files. Data Domain recommends when a recipient receives the alert, he or she confirms the reset operation was intended.

Note

For Retention Lock Governance files (only), you can delete retention locked files using a two step process: First use the `mtree retention-lock revert path` command to revert the retention locked file. Next, delete the file on the client system using the `rm filename` command.

```
mtree retention-lock set {min-retention-period | max-retention-period} period mtree mtree-path
```

Set the minimum or maximum retention period for the specified MTree. This command option requires security officer authorization if Retention Lock Compliance is enabled on the MTree. Role required: admin, limited-admin.

Users cannot set the minimum retention period to fewer than 12 hours. Doing so generates a message notifying the user that the entry was invalid and stating the minimum retention period allowed.

When setting the lock period for Retention Lock Compliance MTrees, users cannot set the period to be less than the current minimum or maximum period allowed. Doing so generates a message notifying the user that the entry was invalid and stating the minimum or maximum retention period allowed.

The retention period is specified in the format [number] [unit]. Possible unit values are:

- min
- hr
- day
- mo
- year

The retention period cannot exceed 70 years. Setting a value greater than 70 years results in an error.

Example 117

To set the min-retention-period to 24 months for mtree1: # `mtree retention-lock set min-retention-period 24mo mtree /data/coll/mtree1`

```
mtree retention-lock show {min-retention-period | max-retention-period} mtree mtree-path
```

Show the minimum and maximum for the specified MTree. Role required: admin, limited-admin, user, backup-operator, security, none.

```
mtree retention-lock status mtree mtree-path
```

Show Retention Lock status for the specified MTree. Possible values are enabled, disabled, previously enabled, and MTree Retention Lock mode: Compliance or Governance. Role required: admin, limited-admin, user, backup-operator, security, none.

mtree show

```
mtree show compression {mtree-path | tenant-unit tenant-unit}
[tier {active | archive | cloud}] [summary | daily | daily-
detailed] [last n {hours | days | weeks | months} | start date
[end date]]
```

Display compression statistics for a specific MTree. Values are reported in Gigabytes (GiB). Running the command without arguments generates default output that displays a summary of compression statistics for all files and directories in the file system for the last 7 days and the last 24 hours. Role required: admin, limited-admin, user, backup-operator, tenant-admin, tenant-user, security, none.

Argument Definitions

mtree-path

The pathname of the MTree for which to display compression statistics.

tenant-unit

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system.

tier {active | archive | cloud} (Optional)

Display results for the specified tier.

summary (Optional)

Display all compression statistics, summarized as:

- Data written in the last 7 days. By including the last *n* option or the start *date* option, you can display statistics for a time frame other than the last 7 days.
- Data written in the last 24 hours.

daily (Optional)

In addition to the summary output, display detailed information for each day, over the previous four full weeks, plus the current partial week. This option is not available if you specify a file or directory name.

daily-detailed (Optional)

Display the daily output and include the following information for each day. This option is not available if you specify a file or directory name.

last *n* {hours | days | weeks | months}(Optional)

In the summary portion of the output, display file system compression statistics for the specified time frame instead of for the past 7 days. The statistics for the last 24 hours remain in the summary output. If you specify a file or directory name, you cannot use this option with the weeks keyword or the months keyword.

start *date* (Optional)

In the summary portion of the output, display file system compression statistics for the time frame that begins on the specified day instead of the past 7 hours. The statistics for the last 24 hours remain in the summary output. If you specify a time frame less than the previous four weeks, plus the current full week, the daily or daily-detailed output (if specified) is truncated to the shorter time frame.

Specify *date* in the format *yyyy-mm-dd* (for example, 2013-04-07). By default, the last day of the time frame specified with this argument is the most recent, full day elapsed.

end *date* (Optional)

Valid only if the start option is used. In the summary portion of the output, display file system compression statistics for the time frame that ends on the specified day.

Example 118 Output

```
# mtree show compression /data/coll/mtree-13
      Pre-Comp   Post-Comp   Global-Comp   Local-Comp   Total-Comp
      (GiB)      (GiB)        Factor        Factor        Factor
      -----   -----   -----   -----   -----
      (Reduction %)
Written:
  Last 7 days      820.3      802.9      1.0x      1.0x      1.0x (2.1)
  Last 24 hrs
-----
```

```
mtree show performance {mtree-path | tenant-unit tenant-unit}
[interval n {mins | hours}] [last n {hours | days | weeks |
months} | start MMDDhhmm[[CC]YY] [end MMDDhhmm[[CC]YY]]]
```

Displays MTree performance statistics. Replicate write data is not included in the output. Role required: admin, limited-admin, user, backup-operator, tenant-admin, tenant-user, security, none.

Example 119 Output

```
Date      Time      Throughput      Streams
YYYY-MM-DD HH:MM  read    write    rs/ws/rr/wr/r+/w+
                MB/s      MB/s      MB/s      #
-----
2015-04-22 21:10  0.00  0.00      0/0/0/0/0/0
2015-04-22 21:20  0.00  0.00      0/0/0/0/0/0
.
.
.
2015-04-23 20:30  0.00  0.00      0/0/0/0/0/0
2015-04-23 20:40  0.00  0.00      0/0/0/0/0/0
-----
```

Where:

- rs: read sequential access streams
- ws: write sequential access streams
- rr: read random access streams
- wr: write random access streams
- r+: reopened read streams in last 30 seconds
- w+: reopened write streams in last 30 seconds

Argument Definitions

mtree-path

The pathname of the MTree for which to display performance statistics.

tenant-unit

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system.

interval *mins* or *hours* (Optional)

The interval is an optional number of minutes or hours.

last *n* {*hours* | *days* | *weeks* | *months*}

In the summary portion of the output, display file system performance statistics for the specified time frame instead of for the past 7 days. The statistics for the last 24 hours remain in the summary output. If you specify a file or directory name, you cannot use this option with the `weeks` keyword or the `months` keyword.

start (Optional)

In the summary portion of the output, display file system performance statistics for the time frame that begins on the specified day instead of the past 7 hours. The statistics for the last 24 hours remain in the summary output. If you specify a time frame less than the previous four weeks, plus the current full week, the daily or daily-detailed output (if specified) is truncated to the shorter time frame. Specify the starting date in the format: `MMDDhhmm[[CC]YY]`

end (Optional)

Valid only if the `start` option is used. In the summary portion of the output, display file system performance statistics for the time frame that ends on the specified day. Specify the ending date in the format: `MMDDhhmm[[CC]YY]`

Example 120

```
# sysadmin@ddr9# mtree show performance /data/coll/55_source
INTERVAL: 10 mins
 "-" indicates that the data is not available for the intervals
```

Date YYYY/MM/DD	Time HH:MM	Throughput		Streams rd/wr/r+/w+ #
		read MB/s	write MB/s	
2014/01/09	15:10	0.00	0.00	0/0/0/0
2014/01/09	15:20	0.00	41.66	0/3/0/0
2014/01/09	15:30	0.00	49.85	0/51/0/0
2014/01/09	15:40	0.00	23.04	0/51/0/0

mtree undelete

```
mtree undelete mtree-path
```

Mark as not deleted the marked-for-deletion MTree at the specified path. This command reverses a previous `mtree delete` command. Role required: `admin`, `limited-admin`.

Note

To undelete an MTree, cleaning must not have run before executing the `undelete` command.

Note

If a user tries to undelete a storage unit, the system displays the following message:
"MTree 'mtree-path' contains a DD Boost storage unit and cannot be undeleted."

Example 121

To reverse a previous `mtree delete` command request that included the MTree at `/data/col1/myMTree`:

```
# mtree undelete /data/col1/myMTree
```

mtree

CHAPTER 28

ndmpd

The `ndmpd` command is the top-level command for the NDMP (Network Data Management Protocol) daemon running on a DD system. The NDMP daemon provides access to VTL-created devices using the NDMP version 4 protocol. Use of this command requires a VTL license. A VTL used by the NDMP tapeserver must be in the TapeServer access group.

This chapter contains the following topics:

- [ndmpd change history](#).....252
- [ndmpd disable](#)..... 252
- [ndmpd enable](#).....252
- [ndmpd option](#)..... 252
- [ndmpd show](#).....252
- [ndmpd status](#)..... 253
- [ndmpd stop](#)..... 253
- [ndmpd user](#)..... 253

ndmpd change history

There have been no changes to this command in this release.

ndmpd disable

```
ndmpd disable
```

Disable the NDMP (Network Data Management Protocol) daemon. Role required: admin, limited-admin.

ndmpd enable

```
ndmpd enable
```

Enable the NDMP (Network Data Management Protocol) daemon. Role required: admin, limited-admin.

ndmpd option

```
ndmpd option reset option-name | all
```

Reset all NDMP (Network Data Management Protocol) daemon options or just a specific option. Role required: admin, limited-admin.

```
ndmpd option set option-name value
```

Set a specific NDMP daemon option. Role required: admin, limited-admin.

```
ndmpd option show option-name | all
```

Show the values for all NDMP daemon options or just for a specific option. Role required: admin, limited-admin.

Argument Definitions

option-name

The NDMP daemon option, which can be authentication, debug, port, or preferred-ip.

value

The value for the particular NDMP daemon option.

ndmpd show

```
ndmpd show devicenames
```

View the device name, VTL virtual name, SCSI vendor and product code, and the serial numbers of devices controlled by the NDMP (Network Data Management Protocol) daemon. Typically, this information is displayed during device discovery and configuration. However, you can use this command to verify the VTL TapeServer group configuration and perform a manual configuration, if required.

If there is no output in the NDMP Device column, either the VTL service is not running or there are no devices registered with the VTL TapeServer. A series of hyphens in the NDMP Device column means the VTL service is running on the system, but has not registered the devices. Restart the VTL service to correct this behavior. If this problem persists, go to the Online Support website for assistance. Role required: admin, limited-admin.

```
ndmpd show sessions
```

View active sessions. Role required: admin, limited-admin.

```
ndmpd show stats session-id | all
```

View statistics of a single session or all sessions. Session numbers are displayed by `ndmpd show sessions`. Role required: admin, limited-admin.

ndmpd status

```
ndmpd status
```

Display the NDMP (Network Data Management Protocol) daemon status. Role required: admin, limited-admin.

ndmpd stop

```
ndmpd stop session session-id | all
```

Stop all NDMP (Network Data Management Protocol) daemon sessions or stop a single session. Role required: admin, limited-admin.

ndmpd user

```
ndmpd user add user-name
```

Add (only) one user name and password for NDMP (Network Data Management Protocol) daemon MD5 authentication. Role required: admin, limited-admin.

```
ndmpd user del user-name
```

Delete the configured NDMP daemon MD5 user name and password. Role required: admin, limited-admin.

```
ndmpd user modify user-name
```

Set the password for the NDMP daemon MD5 user name. Role required: admin, limited-admin.

```
ndmpd user show
```

Show the NDMP daemon MD5 user name. Role required: admin, limited-admin.

ndmpd

CHAPTER 29

net

The `net` command manages the use of all IP network features and displays network information and status.

Federal certification requirements state that the DD OS must be IPv6-capable and that interoperability with IPv4 be maintained in an heterogeneous environment. As a result, several `net` command options include arguments for both versions of Internet Protocol. 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.

If you do not specify an IP version, the default is IPv4 to maintain compatibility with DD OS versions prior to 5.2. The exception is `show` commands. If the version is not specified in the `show` command option (as in `route show table`), both address versions are displayed. To view the IPv4 routes only, you must specify the `IPv4` argument.

For some commands, you must include the IPv6 command argument if the host is to be accessed using its IPv6 address. This is required when a hostname is specified and the host name format resembles an IPv4 address.

This chapter contains the following topics:

• net change history	256
• net guidelines and restrictions	256
• net aggregate	256
• net config	259
• net congestion-check	263
• net create	267
• net ddns	267
• net destroy	268
• net disable	269
• net enable	269
• net failover	269
• net filter	271
• net hosts	274
• net iperf	275
• net lookup	278
• net modify	278
• net option	278
• net ping	278
• net reset	279
• net route	279
• net set	285
• net show	286
• net tcpdump	292
• net troubleshooting	292

net change history

There have been no changes to this command in this release.

net guidelines and restrictions

- Changes made by the net command to disabled Ethernet interfaces flush the routing table. If possible, make interface changes only during scheduled downtime. After making changes to disabled interfaces, you must reconfigure all routing rules and gateways.
- IPv4 is the default IP version.
- The output of the net show settings command displays !! for a failed slave, but the net aggregate show command will not display the failed slave under the bonded interface. The !! disappears after the failed slave is replaced.
- When creating network filters, some system rules supersede any user-created filters:
 - There is no blocking for outbound network traffic. When the Data Domain systems initiates outbound network traffic, the responses to it are considered established and are allowed, even if the response comes from a blocked address or interface.
 - If a network address of interface is blocked after a connection between it and the Data Domain system is already established, traffic will continue over that connection. After that connection is dropped, new inbound connections from that address or interface are blocked.
 - SSH and HTTPS connections to the admin interface from a blocked address are allowed.

Note

The default port numbers for SSH (port 22) and HTTPS (port 443) cannot be changed.

net aggregate

```
net aggregate add virtual-ifname interfaces physical-ifname-list [mode {roundrobin | balanced hash {xor-L2 | xor-L3L4 | xor-L2L3} | lacp hash {xor-L2 | xor-L3L4 | xor-L2L3} [rate {fast | slow}]] [up {time | default}] [down {time | default}]
```

Add physical interfaces to an aggregate virtual interface. Setting the mode is required on initial configuration and when there is no default aggregate mode, but optional when adding interfaces to an existing aggregate interface. Choose the mode compatible with the specifications of the system to which the ports are attached. Balanced and LACP modes require a hash selection.

Note

The *rate* argument can only be used with LACP mode.

Role required: admin, limited-admin.

Argument Definitions

interfaces *physical-ifname-list*

Specifies the physical interfaces to be added to the aggregate virtual interface. To display the physical interfaces on the system, enter `net show hardware`. The interface names appear in the Port column. For information on supported interfaces, refer to the *Data Domain Operating System Administration Guide*.

mode {roundrobin | balanced hash {xor-L2 | xor-L3L4 | xor-L2L3} | lacp hash {xor-L2 | xor-L3L4 | xor-L2L3}}

Specifies how traffic is routed over the aggregate interfaces.

balanced hash {xor-L2 | xor-L3L4 | xor-L2L3}

Data is sent over interfaces as determined by the hash method selected. Balanced mode requires a hash configuration.

lacp hash {xor-L2 | xor-L3L4 | xor-L2L3}

LACP is a link aggregation mode based on the Link Aggregation Control Protocol (LACP, IEEE 802.3ad). From a switch perspective, this configuration is always an active LACP configuration; it cannot be set to passive. When this mode is selected, both ends must be configured with LACP. LACP mode requires a hash configuration.

For successful communication, an interface must be able to communicate with its directly attached partner, and carrier must be up. The switch LACP ports must reside on a single switch except for special cases of virtual switch ports. To fail across switches, failover bonding must be used.

roundrobin

Packets are transmitted sequentially, beginning with the first available link and ending with the last link in the aggregated group.

xor-L2

Transmission of packets from a specific slave interface is based on static balanced mode or LACP mode aggregation with an XOR based on a hash policy. An XOR of source and destination MAC addresses is used to generate the hash.

xor-L2L3

Transmission of packets from a specific slave interface is based on static balanced and LACP mode aggregation with an XOR based on a hash policy. An XOR of source and destination's upper layers (L2 and L3) protocol information is used to generate the hash.

xor-L3L4

Transmission of packets from a specific slave interface is based on static balanced and LACP mode aggregation with an XOR based on a hash policy. An XOR of source and destination's upper layers (L3 and L4) protocol information is used to generate the hash.

rate {fast | slow}

Specifies how often an LACP message is sent to the switch or system that is connected to the Data Domain system. The message identifies the aggregated interface and serves as a type of heartbeat. The rate determines how fast LACP recognizes when an interface can and cannot be used.

`Slow` is the default setting, which sends the message once every 30 seconds. `Fast` sends the message every second. The `Fast` setting generates more traffic comprised of small packets (100 bytes or less) across all aggregated LACP interfaces, but it can detect data transfer failures faster and might be better for faster 10 Gb interfaces.

up { *time* | default}, down { *time* | default}

The length of delay allowed before the link is considered up or down. When interface carrier is present for the interval configured in `up time`, the interface is considered *up*. When interface carrier is absent for the interval configured in `down time`, the interface is considered *down* and not available. The up and down times are rounded down to a multiple of 900 milliseconds. For example if 10,000 milliseconds is configured, 9,900 milliseconds is used. The default up and down times are 29,700 milliseconds.

When the link is down:

- Data is no longer sent to the interface.
- For aggregation bonding, aggregation is recalculated.
- For failover bonding, if the affected interface is the active interface, then the active interface is switched to another interface that is up.

When the link is up:

- Data can be sent over it.
- For aggregation bonding, aggregation is recalculated to include the up link.

virtual-ifname

Specifies a virtual interface to create or modify. 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. To display a list of aggregate virtual interfaces, enter `net aggregate show`.

Example 122

The following command enables link aggregation on virtual interface `veth1` to physical interfaces `eth1a` and `eth2a` in mode `lacp hash xor-L2`.

```
# net aggregate add veth1 interfaces eth1a eth2a mode lacp hash xor-L2
```

```
net aggregate del virtual-ifname interfaces {physical-ifname-list | all}
```

Delete one or more physical interfaces from the specified aggregate virtual interface. To display information on the aggregate virtual interfaces, enter `net aggregate show`. Role required: `admin`, `limited-admin`.

Example 123

To delete physical interfaces `eth2a` and `eth3a` from the aggregate virtual interface `veth1`:

```
# net aggregate del veth1 interfaces eth2a,eth3a
```

```
net aggregate modify virtual-ifname [mode {roundrobin | balanced hash {xor-L2 | xor-L3L4 | xor-L2L3} | lacp hash {xor-
```

```
L2| xor-L3L4 | xor-L2L3} [rate {fast | slow}]]] [up {time |
default}] [down {time | default}]
```

Change the configuration of an existing aggregate virtual interface. Choose the mode compatible with the specifications of the system to which the ports are attached. Balanced and LACP modes require a hash selection. The argument definitions are the same as for `net aggregate add`. Role required: admin, limited-admin.

Example 124

Use the following command to change link aggregation on virtual interface veth1 to mode lACP hash xor-L2. Stating the previous configuration is not required.

```
# net aggregate modify veth1 mode lACP hash xor-L2
```

```
net aggregate show
```

Display basic information on the aggregate setup. If there are no slave interfaces in the up state, the output displays No interface in the aggregate mode. Role required: admin, limited-admin, security, user, backup-operator, or none.

Note

With the exception of `net aggregate show`, `net aggregate` commands control link aggregation. The recommended and supported maximum is four ports, but there are no restrictions on the Data Domain system for having more aggregate slaves.

net config

```
net config addresses type {fixed | floating}
```

Bulk convert IP addresses on the Data Domain system to the specified type. The system prompts for each IP address individually. This command only works on HA systems. Role required: admin.

```
net config ifname {[ipaddr [netmask mask]] | [ipv6addr/prefix]
| [type {fixed | floating}] | [dhcp {yes [ipversion {ipv4 |
ipv6}] | no}]} {[autoneg] | [duplex {full | half} speed {10|
100|1000|10000}] [up | down] [mtu {size | default}] [txqueuelen
size]
```

Display the physical interface configuration or configure a base interface or an alias interface. A base interface is a physical interface to which an IP address is assigned. An alias interface is used to add an additional IP address to a base interface, and you can create multiple alias interfaces to add multiple IP addresses to a base interface.

Note

An alias interface does not operate as an independent interface. DD OS generates statistics and supports additional configuration settings only for a base interface. The only function of an alias interface is to add an additional IP address to the base interface.

To create an alias interface, enter the base interface and alias name in the following format: *base_interface.alias_name* and specify an IPv4 or IPv6 address. The following are some sample alias names.

- eth5a:35—The base interface is physical interface eth5a, and the alias name is 35.
- veth4:26—The base interface is virtual interface veth4, and the alias name is 26.

- eth5a.82:162—The base interface is VLAN interface eth5a.82, and the alias name is 162.

To delete an alias interface, assign the 0 value to the IP address as follows: `net config eth0a:200 0`

Role required: admin, limited-admin.

Argument Definitions

autoneg

Specify this option to configure the interface to autonegotiate the duplex and speed settings with the remote interface.

dhcp {yes [ipversion {ipv4 | ipv6}] | no}

Set the `dhcp` option to `yes` to configure the interface to receive the IP address configuration from a DHCP server, and set this option to `no` when you want to manually configure the IP address. The default option requests an IPv4 address from DHCP, but you can select either IPv4 or IPv6 when you enable DHCP. When you use DHCP, the IP address delivered by DHCP replaces any static IP address previously configured for the base interface.

Note

DHCP over IPv6 does not supply a host name. If you set an interface to use DHCP over IPv6, complete the configuration, run `net show hostname`, and verify that the hostname is correct. If there is no host name or if it is no longer correct, configure the hostname using `net set hostname` or with DD System Manager at **Hardware > Network > Settings**.

duplex {full | half} speed {10|100|1000|10000}

Specify this option when you want to manually configure the duplex setting or speed. The speed settings are 10, 100, 1,000, or 10,000 Mbps. This option automatically disables autonegotiation on the interface.

ifname

Specify the interface to configure and one or more arguments to change the configuration. If you omit the interface name, the command lists the configuration for all the interfaces. If you specify an interface without any additional arguments, the command lists the configuration for the interface.

To create an alias interface, enter the alias in the following format: `base_interface.alias_name`. The alias name must be a number in the range of 1 to 9999.

type { fixed | floating}

HA systems use two types of IP addresses. Use the `fixed` IP option for node-specific configuration/management, which can be static or DHCP, IPv6 SLAAC, or IPv6 Link Local.

Note

The IPv6 SLAAC and IPv6 Link Local addresses cannot be configured. They are automatically configured when the interface is brought up.

The SLAAC addresses are generated based on the response from the router and is based on the mac address. The Link Local is also based on the mac address but is generated whenever the physical, bonded, or VLAN is brought up to the running state.

Use the `floating` IP option for file system access and most configuration. The `floating` IP is static.

Note

`Floating` IP addresses only exist on an HA system and should be configured on the active node. When upgrading from a single node, the `fixed` IP will need to be manually converted to a `floating` IP address and requires the `type floating` argument. During failover, the IP will "float" to the new active node. When the HA configuration is destroyed, all floating IPs convert to fixed IPs.

ipaddr [netmask mask]

Specify an IPv4 address for the interface. The `dhcp` option must be set to `no` to support manual IP address configuration.

Use the `netmask` option to specify a network mask that is different from the default `netmask`. The `netmask` can only be specified when an IPv4 address is specified.

ipv6addr/prefix

Specify an IPv6 address for the interface. The `dhcp` option must be set to `no` to support manual IP address configuration. The `dhcp` option is automatically set to `no` if a static address is set.

If an IPv6 address is specified, there is no associated `netmask`. Instead, a prefix length is used to determine the subnet. The default prefix length is 64. To use a prefix length different from 64, it must be specified with the address by adding a forward slash followed by a number. For example, if the prefix length is 52, the notation is: `2026:3456:cafe::f00d:1/52`.

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.

mtu {size | default}

The range for the MTU size is 350 - 9000 for IPv4 and 1280 - 9000 for IPv6. To ensure backward compatibility DD OS accepts an MTU size of 9014, but sets it to 9000 if the MTU requested is greater than 9000 and less than or equal to 9014.

txqueuelen size

Specify the transmit queue length. The range is 500 to 10,000 packet pointers, and the default value is 1000.

up | down

Use the `up` argument to bring up an interface with or without an IP address. (Using `net enable` fails if no IP address is configured on the interface.) Use the `down` argument to bring down an interface.

Note

If no address is given, the `up` option might fail because there is no registry entry for an IP address. This typically occurs after a fresh install. If this occurs, specify an address of 0 to allow a registry address location to be created.

Example 125

The following example shows an excerpt from the `net config` display when no arguments are entered.

```
eth1d  Link encap:Ethernet HWaddr 00:1B:21:5F:E2:4D
      inet6 addr: 2100:bad:dead:f00d::e4b:100/64 Scope:Global
      inet6 addr: 2100:bad:dead:f00d::e4b:210/64 Scope:Global
      inet6 addr: fe80::21b:21ff:fe5f:e24d/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:0 errors:0 dropped:0 overruns:0 frame:0
      TX packets:37274 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:0 (0.0 b) TX bytes:2431901 (2.3 MiB)

eth1d:10 Link encap:Ethernet HWaddr 00:1B:21:5F:E2:4D
      inet addr:192.168.141.20 Bcast:192.168.141.255 Mask:255.255.255.0
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1

eth1d:100 IPv6 alias address, 2100:bad:dead:f00d::e4b:100/64, is on the interface eth1d
when up

eth1d:200 Link encap:Ethernet HWaddr 00:1B:21:5F:E2:4D
      inet addr:192.168.141.200 Bcast:192.168.141.255 Mask:255.255.255.0
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1

eth1d:210 IPv6 alias address, 2100:bad:dead:f00d::e4b:210/64, is on the interface eth1d
when up
```

Interface `eth1d` represents the physical interface. Interfaces `eth1d:10` and `eth1d:200` are alias interfaces that each add an IPv4 address to the base interface, and `eth1d:100` and `eth1d:210` are alias interfaces that add IPv6 addresses to the same base interface. The IPv6 alias addresses are available when the base interface is in the running state and the alias interface state is *up*. Notice that the IPv6 alias addresses in the example above are displayed with the alias interfaces and the base interface.

Example 126

The following example adds an alias named `200` to the `eth0a` interface and assigns an IPv6 address to it.

```
# net config eth0a:200 2620:0:170:1a04:28c:faff:fe05:6c91/64
Creating interface ...
Done.
Configuring interface...
done.
```

Example 127

The following example deletes alias 200 from the eth0a interface.

```
# net config eth0a:200 0
Alias is destroyed.
```

net congestion-check

```
net congestion-check modify [sample-interval secs] [capture-
window secs] [every mins] [detailed {on | off}] [logfile
filename] [logfilev6 filename] [iperf-client {none | iperf-
server-host | iperf-server-ipaddr} [nodelay {on | off}]] [port
{port | default} ] [window-size bytes] [connections count]
[data {random|default}]]]
```

Congestion data is collected during a period of time defined by the *capture window* argument. Within the capture window, data is captured at intervals defined by the *sample-interval* argument. If the *every* argument is non-zero, a new capture window starts at intervals determined by the *every* argument. For example, if the capture window is 60 seconds, the sample interval is 5 seconds, and the *every* argument is set to 60, data is collected every 60 minutes for a period of 60 seconds at 5 second intervals. The output displays as one line per remote IP address.

This command modifies options for the congestion monitor whether or not the monitor program is activated. (The congestion monitor is activated by the `net congestion-check start` command.) The settings configured with this command are stored in the registry and replace the default values used by `net congestion-check start` command. If the congestion monitor is scheduled, the new registry values are used when it runs. If the monitor is not scheduled, the values are used as defaults when it is started. Typically this command option is used after the monitor is scheduled to run and the user does not want to stop and restart the monitor.

Output values for rates and error numbers are added together. Values that may increase or decrease, such as the capture-window, are averaged over time. Role required: admin, limited-admin.

Argument Definitions

The argument definitions are the same as described for the `net congestion-check start` command.

```
net congestion-check run [sample-interval secs] [capture-window
secs] [every mins] [detailed {on | off}] [logfile filename]
[{{iperf-client {none | {iperf-server-host | iperf-server-
ipaddr} [nodelay {on | off}]] [port {port | default}}] [window-
size bytes] [connections count] [data {random | default}]]]
```

Run the congestion check program with the `run` option to display the results as screen output when the *capture-window* time is complete. When the command option is entered without arguments, defaults are used. When the command option includes arguments, the arguments override the defaults during the procedure but return to the configured defaults after the procedure concludes. Default values for the `run` command are always the same and are not affected by the `net congestion-check modify` or the `net congestion-check start` commands.

Argument Definitions

The following argument definitions are unique to this command. The rest of the argument definitions are the same as described for the `net congestion-check start` command.

port {port | default}

The TCP port number for the target iperf server. The default is 5002, which is one more than the iperf default, 5001.

```
net congestion-check start [sample-interval secs] [capture-
window secs] [every mins] [detailed {on | off}] [logfile
filename] [{iperf-client {none | {iperf-server-host | iperf-
server-ipaddr} [nodelay {on | off}] [port {port | default}]
>window-size bytes] [connections count] [data {random|
default}}}]
```

Start the congestion monitor and schedule when it is to be run using the time arguments: `sample-interval`, `capture-window`, and `every`. Command output is stored in the `congestion.log` and `congestion6.log` files, unless the names are changed from the command line. When the command option is run with arguments, the arguments override the defaults and become the new default values. The remaining arguments of `net congestion-check start` command are used to configure in detail how the monitor is run.

Note

After entering the command, there is a slight delay during which the process actually starts the monitor. After the monitor is started, the specified time arguments take over. To get information immediately, use the `net congestion-check run` command instead.

Output is one line per external destination. All connections to and from an external address are merged into a single line of data.

Value types from the output vary. Amounts of data or packets increase. These amounts are added together across all connections to a specific IP address to give the total value to or from the external location. Rates are relatively constant but are also added together to give the total flow rate to the pipe at the remote location. Other values are relatively static across all connections, such as the `mss`, `rtt`, `window scale factor`, or `congestion window`. These are given as an average with the minimum and maximum. Errors and losses are treated the same as rates and are added across all interfaces. Role required: `admin`, `limited-admin`.

Argument Definitions

capture-window secs

Specify the period during which data is captured. The initial value is 60 seconds; the range is 10 to 3600 seconds. The configured value must be less than the `every` argument and greater than the `sample-interval` argument.

connections count

The `connections` argument determines how many parallel TCP connections to establish between the iperf client and server. The default value is 1, which is typically satisfactory if the window size is set appropriately. Larger values are supported between DD OS iperf clients and servers, but are not supported by all iperf servers. Increasing the connection count can improve performance, but too many connections will negatively impact network performance. This argument is equivalent to the Linux argument for parallel tests: `-p number`.

data {random | default}

The `data default` argument allows iperf to send "normal" data and uses fewer system resources than the `data random` argument. If you suspect that WAN accelerators are contributing to artificially high performance statistics, you can use the `data random` argument to have iperf use random data that is difficult for WAN accelerators to accelerate. This argument is equivalent to the Linux argument: `-R`

detailed

By default, detailed information is saved, but setting the argument to `off` saves basic information. The basic setting is mainly for replication on the source system and focuses on congestion conditions between the source and destination. The `detailed on` argument adds receive information and other entries useful for the general network environments of the Data Domain system. The initial value is `off`.

every mins

Specify the period between the start of each capture window. The initial period is 60 minutes. The range is 10 to 60 minutes. The configured period must be greater than that for the `capture-window` argument. Because this command configures an ongoing monitor, value 0 is not supported.

iperf-client {none | iperf-server-host | iperf-server-ipaddr} [nodelay {on | off}] [port {port | default}] [window-size bytes] [connections count] [data {random | default}]}

The `iperf-client` argument can be used to generate network traffic for throughput testing. This argument is typically used when there is insufficient normal traffic for a capacity test. This argument should be used with caution on production networks because iperf is disruptive to the network. Consider using iperf for brief periods, especially if there is other traffic using the network.

The iperf client requires an iperf server to communicate with. You can use `net iperf server` to start an iperf server on a remote system.

When the congestion monitor is configured to use the iperf client, iperf starts at the beginning of a capture window. If you want to control iperf operation manually, you can start an iperf client with the `net iperf client` command before the congestion check is performed. The advantage of letting the congestion-check manage iperf operation is that the iperf client runs only when the congestion-check requires it. Otherwise iperf does not run.

Iperf is an open-source utility. For more information on iperf, search for iperf on the World Wide Web.

iperf-server-host | iperf-server-ipaddr

The `iperf-server-host` and `iperf-server-ipaddr` arguments enable the iperf client to run during congestion checks and specify a target iperf server hostname or IP address.

logfile

Set the log file name used to save the IPv4 data collected. The initial default is `/ddvar/log/default/congestion.log`.

Do not change the file name unless absolutely required. The default file name is on a rotation system where the file size cannot exceed 10 MB, and up to 10 files are saved for a maximum of 100 MB of disk space. Changing the file name voids the space restrictions, meaning there is no limit to the space the files may consume.

logfilev6 filename

Set the log file name used to save the IPv6 data collected. The initial default is `/ddvar/log/default/congestion.log`.

Do not change the filename unless absolutely required. The default file name is on a rotation system where the file size cannot exceed 10 MB, and up to 10 files are saved for a maximum of 100 MB of disk space. Changing the file name voids the space restrictions, meaning there is no limit to the space the files may consume.

nodelay

The `nodelay on` argument eliminates the wait time between sends. The `nodelay off` argument requires iperf to wait for an ACK message after each send. This argument is equivalent to the Linux argument: `-N`.

none

The `none` argument prevents the iperf client from running during a network congestion check. If iperf was previously enabled and is no longer needed, use the `none` argument to disable iperf use by the congestion monitor. The initial value is `none`.

port {port | default}

The TCP port number for the target iperf server. The initial default number is 5002. If a different port number is specified with this command, that port number becomes the new default value. If the port number is changed with `net congestion-check modify`, the new port number is used the next time a congestion check is scheduled to run.

sample-interval secs

Specify the sample period within the capture window. The initial value is 4 seconds; the range is 2 to 3600 seconds. The value of the `sample-interval` argument must be less than the value of the `capture-window` argument.

window-size bytes

The size of the socket buffer to use. The default is 32,000; the range is 8,000 to 10,000,000. For long latencies, this size may be too small. Consider setting the size to 250,000 or 10,000,000.

```
net congestion-check status
```

Display the state of the congestion monitor. The congestion monitor is started when the `net congestion-check start` command is issued. The `status` argument displays the configured timings, the level of logging, the log file, the monitored connections, if the monitor is actually running or scheduled to run, and if iperf is specified to run. It also shows if iperf is currently running and which connections are being monitored. Role required: `admin`, `limited-admin`, `security`, `user`, `backup-operator`, or `none`.

```
net congestion-check stop
```

If the congestion monitor is running, this command stops the monitor. If the congestion monitor is scheduled to run, this command option un schedules the monitor. A message notifies you if the congestion monitor is not scheduled and no action is taken. Role required: `admin`, `limited-admin`.

net create

```
net create interface {physical-ifname | virtual-ifname} {vlan
vlan-id | alias alias-id}
```

Create a VLAN interface or alias on the specified physical or virtual interface. A VLAN is created immediately in the kernel, and the number given must be between 1 and 4094 inclusive. An alias is not created in the kernel until an IP address is specified. Role required: admin, limited-admin.

Note

The alias argument is deprecated and will be removed in a future release. You can use the alias option in this release, and if you do, the system displays a message regarding the deprecation and the preferred way to enter an alias.

```
net create virtual vethid
```

Create a virtual interface. The virtual interface name *vethid* must begin with veth. The remainder of the name is a decimal number. Interface names must be unique.

There are no restrictions except for the size and the number. The maximum size for an interface name is 15 characters, which includes VLAN, alias names, and the associated dot and colon. The virtual interface name must be kept at a minimum. If possible, use a number in the range of 0 to 99. However, the maximum value is 9999.

The number of virtual interfaces cannot exceed the number of physical interfaces. For example, if there are 10 physical interfaces there can be no more than 10 virtual interfaces. Role required: admin, limited-admin.

net ddns

```
net ddns add {ifname-list | all | ifname interface-hostname
hostname}
```

Add interfaces to the Dynamic DNS (DDNS) registration list. Role required: admin, limited-admin.

Note

When DDNS is configured for UNIX mode, this feature supports physical interfaces and aliases for physical interfaces. In this release, VLAN and virtual interfaces (and any aliases for those interfaces) are not supported in DDNS UNIX mode.

Argument Definitions

all

When DDNS is enabled for the Windows environment, this option is enabled and specifies that host names be registered for all interfaces.

ifname-list

When DDNS is enabled for the Windows environment, this option is enabled and specifies that host names be registered for the specified interfaces.

ifname *interface-hostname*

When DDNS is enabled for the UNIX environment, this option is enabled and specifies an interface and hostname to be registered with DDNS.

```
net ddns del {ifname-list | all}
```

Remove one or all interfaces from the DDNS registration list. To display the list entries, enter `net ddns show`. Role required: admin, limited-admin.

```
net ddns disable
```

Disable DDNS updates. Role required: admin, limited-admin.

```
net ddns enable [windows | unix [TSIG-key key]]
```

Enable DDNS updates for either Windows or UNIX environments. The transaction signature key option (TSIG-key) is for secure connections to the server where the key and a secret are defined. If you enter a TSIG key, the system prompts you to enter the corresponding *secret*. Role required: admin, limited-admin.

Note

If DDNS is already enabled, you must disable DDNS before selecting a different mode.

```
net ddns register
```

Register configured interfaces with DNS. Role required: admin, limited-admin.

```
net ddns reset
```

Clear the DDNS interface list and disable registration. In Windows mode, the registration list is set to auto. In UNIX mode, the TSIG key is also deleted. Role required: admin, limited-admin.

```
net ddns reset TSIG-key
```

Clear the TSIG key and secret. Role required: admin, limited-admin.

```
net ddns set TSIG-key key
```

Set the TSIG key and secret. The system will prompt you for the secret. Role required: admin, limited-admin.

```
net ddns show
```

In Windows mode, display the enabled interfaces. In UNIX mode, display the UNIX mode status and the enabled interfaces. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
net ddns status
```

Display only the DDNS status, which can be enabled in Windows mode, enabled in UNIX mode, or disabled. Role required: admin, limited-admin, security, user, backup-operator, or none.

net destroy

```
net destroy {virtual-ifname | vlan-ifname | ipalias-ifname}
```

Remove a VLAN, IP alias, or virtual interface. If VLANs and aliases are associated with a virtual interface, or if aliases are associated with a VLAN, these entities are also destroyed when the virtual interface or VLAN interface is destroyed. Role required: admin, limited-admin.

Note

Setting the address to zero for an alias will also cause it to be destroyed.

Example 128

The following commands remove a VLAN named eth1a.35, an alias on a virtual interface named veth23:2, and the alias interface eth1b:57.

Example 128 (continued)

```
# net destroy eth1a.35
# net destroy veth23:2
# net config eth1b:57 0
```

net disable

```
net disable ifname
```

Disable an Ethernet interface on the Data Domain system. Role required: admin, limited-admin.

net enable

```
net enable ifname
```

Enable or reenables an Ethernet interface on the Data Domain system, where *ifname* is the name of an interface. An IP address must be assigned to the interface. When the interface is configured properly, this command brings up the interface to the RUNNING state. If the interface does not go into the RUNNING state, the command fails, the interface is set to the DOWN state, and then set to disabled. Role required: admin, limited-admin.

net failover

```
net failover add virtual-ifname interfaces ifname-list [primary ifname] [up {time | default}] [down {time | default}]
```

Add interfaces to a failover virtual interface. Note that you can add an aggregated interface to a failover interface. Role required: admin, limited-admin.

Argument Definitions

interfaces *ifname-list*

Specifies one or more slave interfaces to be added to the failover virtual interface. The slave interfaces must be in a down (disabled) state when added to the virtual interface. (Use `net show settings` to view the link state of all interfaces.)

virtual-ifname

Specifies a virtual interface to modify. To display a list of failover virtual interfaces, enter `net failover show`.

primary *ifname*

Specifies an interface as the primary failover slave interface.

up {*time* | default}, down {*time* | default}

The length of delay allowed before the link is considered up or down. When interface carrier is present for the interval configured in `up time`, the interface is considered *up*. When interface carrier is absent for the interval configured in `down time`, the interface is considered *down* and not available. The up and down times are rounded down to a multiple of 900 milliseconds. For example if 10,000 milliseconds is configured, 9,900 milliseconds is used. The default up and down times are 29,700 milliseconds.

When the link is down:

- Data is no longer sent to the interface.
- For failover bonding, if the affected interface is the active interface, then the active interface is switched to another interface that is up.

When the link is up:

- Data can be sent over it.
- If the interface is the primary interface or the sole slave interface, it becomes the active interface and traffic is diverted to it. Any other slave interface is added into the failover interface pool.

Example 129

The following command example associates a failover virtual interface named `veth1` with the physical interfaces `eth2a` and `eth3a` and designates `eth2a` as the primary interface.

```
# net failover add veth1 interfaces eth2a eth3a primary eth2a
```

```
net failover del virtual-ifname interfaces {ifname-list | all}
```

Delete slave interfaces from a failover interface. The freed interface remains disabled after being removed from the virtual interface. Use commas, spaces, or both to separate list entries, or specify `all` to delete all slave interfaces. To delete a primary interface, use `net failover modify` to specify another interface as primary or set the primary to `none`. The argument definitions are the same as for `net failover add`. Role required: admin, limited-admin.

Example 130

The following command removes `eth2a` from the virtual interface `veth1`, for which `eth2a` and `eth3a` are slaves and `eth3a` is the primary interface.

```
# net failover del veth1 interfaces eth2a
```

```
net failover modify virtual-ifname [primary {ifname | none}]
[up {time | default}] [down {time | default}]
```

Modify the primary network interface, the up /down times for a failover interface, or both. A down interface must transition and stay up for the amount of `time` to be designated up. An up interface must transition and stay down for the amount of `time` to be designated down.

The up and down time is given in milliseconds and is adjusted internally to the largest multiple of 900, less than or equal to the specified value. For example, if the time you want is 10 seconds and 10000 is specified, the actual value is 9900. The default value is 30 seconds but the actual resulting value is 29.7 seconds.

A primary interface cannot be removed from failover. To remove a primary use `primary ifname none` first. The argument definitions are the same as for `net failover add`. Role required: admin, limited-admin.

Example 131

```
# net failover modify veth1 up 5000 down 10000
```

Example 131 (continued)

The up time value used is 4500 (4.5 seconds) and the down time value is 9900 (9.9 seconds).

```
net failover show
```

Display the full configuration details for each of the failover interfaces for which at least one of the physical interfaces are up. The displayed information includes the MAC address, the list of configured interfaces, the primary interface (if any), and the up and down delays. Role required: admin, limited-admin, security, user, backup-operator, or none.

Note

If a physical interface is down, none of the associated failover interfaces appear in this list. To see all failover interfaces, regardless of the states of the physical interfaces, use `net show settings`.

net filter

```
net filter add [seq-id n] operation {allow | block} [clients
{host-list | ipaddr-list}] [except-clients {host-list | ipaddr-
list}] [interfaces {ifname-list | ipaddr-list}] [except-
interfaces {ifname-list | ipaddr-list}] [ipversion {ipv4 |
ipv6}]
```

Add a set of rules to the iptables; service names are restricted to what is supported. Role required: admin, limited-admin.

Argument definitions**seq-id n**

The sequence number of where to add the function into the current filter functions. If it is not specified it will be appended to the end of the user generated filter functions but before the default functions.

operation {allow | block}

This option determines if the packets with the specified information will be allowed to be further process the packet or to discard the packet if it contains the specified information.

clients {host-list | ipaddr-list}

Either a list of client host names or client addresses (is the source address in the packet) on which the operation is to be performed. No more than 25 IPs, host names, or interfaces are allowed.

except-clients {host-list | ipaddr-list}

Either a list of client host names or client addresses (is the source address in the packet) on which the reverse of the operation is to be performed. No more than 25 IPs, host names, or interfaces are allowed.

interface {interface-list | ipaddr-list}

Either a list of client local interfaces or local addresses (is the destination of the packet) on which the operation is to be performed. No more than 25 IPs, host names, or interfaces are allowed.

except-interface {interface-list | ipaddr-list}

Either a list of client local interfaces or local addresses (is the destination of the packet) on which the opposite of the operation is to be performed. No more than 25 IPs, host names, or interfaces are allowed.

ipversion {ipv4 | ipv6}

This option indicates whether the address is to be applied to IPv4 or IPv6 filter functions. The default is IPv4 if none is given.

```
net filter auto-list add ports {all}
```

Configures the net filter to allow connections to the system only from ports with a listen thread. When the auto-list function is on, the default net filter functions are disabled except for functions specifically enabled by the user. There is no way to configure this functionality for individual ports, it is either enabled or disabled for all ports. Role required: admin, limited-admin.

```
net filter auto-list delete ports {all}
```

Configures the net filter to allow connections to the system from all supported ports. When the auto-list function is off, the default net filter functions are enabled. There is no way to configure this functionality for individual ports, it is either enabled or disabled for all ports. Role required: admin, limited-admin.

```
net filter clear stats
```

Clear all iptables rules statistics.

```
net filter config reset [admin-interface] [ipversion {ipv4 | ipv6}]
```

Set the net filter configuration option to the default value; reset all if no options specified. Role required: admin, limited-admin.

```
net filter config set [admin-interface ifname] [ipversion {ipv4 | ipv6}]
```

Set a net filter option. Role required: admin, limited-admin.

Argument definitions**admin-interface**

Set the admin interface to the specified name if it is available and in a running state. If an alias is given, the immediate base interface will be used. For example, if eth0a:55 is given then eth0a is used. If eth4b.67:22 is used eth4b.67 is used. The alias cannot be specified as an admin-interface.

```
net filter config show [admin-interface] [ipversion {ipv4 | ipv6}]
```

Displays the net filter configuration option value; displays all if no options specified. Role required: admin, limited-admin.

Argument definitions**admin-interface**

Shows what the admin interface is set to. If not specified, all configurations are shown.

```
net filter delete {seq-id | all} [ipversion {ipv4 | ipv6}]
```

Delete one net filter command or all of them. A default function cannot be deleted. Role required: admin, limited-admin.

Argument definitions**seq-id**

The sequence number of functions that are being deleted. The default functions can not be deleted.

all

Delete all functions. The default functions can not be deleted.

```
net filter disable {seq-id | all} [ipversion {ipv4 | ipv6}]
Disable one net filter command or all of them. Role required: admin, limited-admin.
```

Argument definitions**seq-id**

The sequence number of functions that are being disabled except the default functions. Only one default function can be disabled per command.

all

Disable all functions. Only one default function can be disabled per command.

```
net filter enable {seq-id | all} [ipversion {ipv4 | ipv6}]
Enable one net filter command or all of them. Role required: admin, limited-admin.
```

Argument definitions**seq-id**

The sequence number of functions you want to enable.

all

Enable all functions.

```
net filter log start
Start writing to the net filter log. Role required: admin, limited-admin.

net filter log stop
Stop writing to the net filter log. Role required: admin, limited-admin.

net filter move seq-id new-seq-id [ipversion {ipv4 | ipv6}]
Move command at seq-id to new-seq-id. Role required: admin, limited-admin.
```

Argument definitions**seq-id**

The sequence number of functions you want to move.

new-seq-id

The sequence number indicating where the function is to be moved.

```
net filter show kernel [ipversion {ipv4 | ipv6}] [chain {chain-
name | all}]
Show iptable rules configured. If a chain is given only the rules for that chain are
displayed. If all is given then all the rules for all of the chains are displayed. The format
is the same format used by iptables to display the rules. Role required: none.
```

Argument definitions**ipversion {ipv4 | ipv6}**

Both are displayed unless one of the protocols is specified.

chain {chain-name | all}

Displays only the specified chain. If none are specified, all are shown.

```
net filter show map {ids | all} [ipversion {ipv4 | ipv6}]
[chain {chain-name | all}]
```

Show the net filter functions configured and the iptable rules associated with each function. One *id* or a list of *ids* can be given. Role required: none.

Argument definitions**ids**

The sequence numbers of the functions to display. If none are specified, all are displayed.

ipversion {ipv4 | ipv6}

If one of the protocols is not specified, both are displayed.

chain {chain-name | all}

If one of the chains is not specified, all are displayed.

```
net filter show seq-id-list {ids | all} [ipversion {ipv4 |
ipv6}] [chain {chain-name | all}]
```

Show the net filter commands configured. One *id* or a list of *ids* can be given. Role required: none.

Argument definitions**ipversion {ipv4 | ipv6}**

If one of the protocols is not specified, both are displayed.

chain {chain-name | all}

If one of the chains is not specified, all are displayed.

```
net filter start
```

Load the iptables module and add required default commands. Role required: none.

net hosts

```
net hosts add {ipaddr | ipv6addr} host-list
```

Add a host list entry. Associate an IP address with a hostname. The address can be an IPv4 or an IPv6. The hostname is a fully qualified domain name, a hostname, or an alias. The entry is added to the `/etc/hosts` file. Entries in the list can be separated by commas, spaces, or both. Role required: admin, limited-admin.

Example 132

To associate the fully qualified domain name `bkup20.yourcompany.com` and the hostname of `bkup20` with an IP address of `192.168.3.3`, enter the following command.

```
# net hosts add 192.168.3.3 bkup20.yourcompany.com bkup20
```

```
net hosts del {ipaddr | ipv6addr}
```

Delete a host list entry from the `//etc/hosts` file. Role required: admin, limited-admin.

```
net hosts reset
```

Clear the hosts list from the `/etc/hosts` file. Role required: admin, limited-admin.

```
net hosts show
```

Display hostnames and IP addresses from the `/etc/hosts` file. Role required: admin, limited-admin, security, user, backup-operator, or none.

net iperf

```
net iperf client {ipaddr | ipv6addr | hostname [ipversion {ipv4 | ipv6}]} [port port] [window-size bytes] [data {random | default}] [interval secs] [{transmit-size bytes | duration secs}] [connections count] [nodelay]
```

This command starts iperf client software, which can be used to generate network traffic and display throughput test results. This command should be used with caution on production networks because iperf is disruptive to the network. Consider using iperf for brief periods, especially if there is other traffic using the network.

The iperf client requires an iperf server to communicate with. You can use `net iperf server` to start an iperf server on a remote system. Role required: admin, limited-admin.

Argument Definitions

connections *count*

The `connections` argument determines how many parallel TCP connections to establish between the iperf client and server. The default value is 1, which is typically satisfactory if the window size is set appropriately. Larger values are supported between DD OS iperf clients and servers, but are not supported by all iperf servers. Increasing the connection count can improve performance, but too many connections will negatively impact network performance. This argument is equivalent to the Linux argument for parallel tests: `-p number`.

data {random | default}

The `data default` argument performs a less-stringent test and uses fewer system resources than the `data random` argument. The `data random` argument performs a more stringent test for traffic optimized and accelerated networks.

duration *secs*

The `duration` argument indicates how many seconds iperf transmits packets. This argument is equivalent to the Linux argument: `-t secs`.

interval *secs*

This argument indicates the time between reports. If this is not given, one is reported at the end. If this is given, a progress report is displayed every "secs" (seconds). Equivalent to the Linux argument: `-i secs`.

***ipaddr | ipv6addr | hostname* [ipversion {ipv4 | ipv6}]**

Identifies the iperf server host. If a hostname is given and the hostname translates to an IPv6 address, the `ipversion` argument must also be specified. The default is an IPv4 address.

nodelay

The `nodelay on` argument eliminates the wait time between sends. If `nodelay` argument is not specified, iperf will wait for an ACK message after each send. This argument is equivalent to the Linux argument: `-N`.

port *port*

The `port` argument can be used to specify the TCP port number for the target iperf server. The default port number is 5001, which is the default value for iperf. The port number used by the iperf client must match the port number used by the iperf server. Typically, you might change the port number to bypass network filters or test specific ports. This argument is equivalent to the Linux argument: `-p port`.

transmit-size *bytes*

The `transmit-size` argument defines how much data iperf will send before closing. This is equivalent to the Linux argument: `-n num`.

window-size *bytes*

The `window-size` argument increases the amount of data sent at one time (socket buffer size). This is equivalent to the Linux argument: `-w iperf-bytes`.

```
net iperf server [run] [ipversion {ipv4 | ipv6}] [port {port |
congestion-check-port}] [window-size bytes]
```

Starts iperf in server mode. Role required: admin, limited-admin.

Argument Definitions**ipversion {ipv4 | ipv6}**

Specifies the type of addressing.

port {*port* | congestion-check-port}

The `port` argument specifies the TCP port number to use instead of the iperf default, which is 5001. Use this argument to specify a port number or the keyword `congestion-check-port`. The `congestion-check-port` keyword selects port 5002. This argument is equivalent to the Linux argument: `-p port`.

window-size *bytes*

The `window-size` argument specifies the amount of data sent at one time. This is equivalent to the Linux argument: `-w iperf-bytes`.

```
net iperf server start [port {port | congestion-check-port}]
[ipversion {ipv4 | ipv6}] [window-size bytes]
```

Runs iperf in the background in server (-s) mode until stopped with `net iperf server stop`. This command enables the terminal to be used for other operations, such as a network congestion check, while iperf is running. Do not use this command except when running in conjunction with `net congestion-check start iperf-client`. Role required: admin, limited-admin.

Argument Definitions**ipversion {ipv4 | ipv6}**

Specifies the type of addressing.

port {*port* | congestion-check-port}

The `port` argument specifies a port number to use instead of the default. The initial default port number is 5002. If the port number is changed, the default becomes the last port number specified. Use this argument to specify a port number or the keyword `congestion-check-port`. The `congestion-check-port` keyword selects port 5002. This argument is equivalent to the Linux argument: `-p port`.

window-size *bytes*

The `window-size` argument specifies the amount of data sent at one time (socket buffer size). This is equivalent to the Linux argument: `-w iperf-bytes`.

```
net iperf server status
```

When the iperf server is running in the background (as invoked by `net_server start`), this command option displays the iperf server status and what connections the server is using. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
net iperf server stop
```

When the iperf server is running in the background (as invoked by `net_server start`), this command option stops iperf. Role required: admin, limited-admin.

net lookup

```
net lookup {ipaddr | ipv6addr | hostname}
```

Search DNS entries. This command may be used with IPv4 or IPv6 addresses. Role required: admin, limited-admin, security, user, backup-operator, or none.

net modify

```
net modify virtual-ifname bonding {aggregate | failover}
```

Change the behavior of the specified virtual interface from aggregate to failover or from failover to aggregate. Role required: admin, limited-admin.

The result is the default's value target function with the same slaves. The default for failover is no primary and up-and-down delays of 29,700 milliseconds. The default for link aggregation is LACP with hash of L3L4, and a rate of slow and up/down times of 29,700 milliseconds.

net option

```
net option show
```

Display settings for network options. Role required: admin, limited-admin, security, user, backup-operator, or none.

net ping

```
net ping {ipaddr | ipv6addr | hostname [ipversion {ipv4 | ipv6}]} [broadcast] [count n] [interface ifname] [packet-size bytes] [path-mtu {do | dont | want}] [pattern pattern] [numeric] [verbose]
```

Verify the Data Domain system can communicate with a remote host. Role required: admin, limited-admin, security, user, backup-operator, or none.

Argument Definitions

broadcast

Enable pinging a broadcast address (available for IPv4 only).

count *n*

Number of pings to issue.

interface *ifname*

Name of interface from which to send the ping. You can ping from physical, virtual, and VLAN interfaces.

***ipaddr | ipv6addr | hostname* [ipversion {ipv4 | ipv6}]**

Identifies the host to ping. Specify an IPv4 or IPv6 address or a hostname. If a host name is entered without specifying an IP version, IPv4 is used. To ping an IPv6 host using the hostname, you must specify `ipversion ipv6` after the hostname.

numeric

Ping the IP address, not the hostname.

packet-size *bytes*

Set packet size.

path-mtu {do | dont | want}

Define the MTU discovery and packet fragmentation strategy.

- Select `do` when you do want to drop packets that are too large (no fragmentation).
- Select `dont` when you don't want to drop oversize packets. Some packets may be dropped during path MTU discovery, but once the path MTU is determined, packets are fragmented locally for the entire path. Fragmentation is not supported after the packet leaves the local system.
- Select `want` when you want the packets delivered and not dropped. Fragmentation can take place locally or at any device along the path.

pattern *pattern*

Send packets with the specified pattern.

verbose

Display expanded output.

net reset

```
net reset {domainname | searchdomains}
```

Reset Data Domain system DNS servers or domain names to the default settings. This usually clears any static settings and ensures DNS addresses provided by DHCP are used. If DHCP is not being used or DHCP does not supply any DNS servers or domain names, then no DNS addresses are used. Role required: admin, limited-admin.

```
net reset dns
```

Reset DNS list to default values. This usually clears any static settings and ensures DNS addresses provided by DHCP are used. If DHCP is not being used or DHCP does not supply any DNS servers or domain names, then no DNS addresses are used. Role required: admin, limited-admin.

```
net reset hostname
```

Reset the hostname to the default value. This usually clears any static settings and ensures DNS addresses provided by DHCP are used. If DHCP is not being used or DHCP does not supply any DNS servers or domain names, then no DNS addresses are used. Role required: admin, limited-admin.

net route

The `net route` command manages routing between Data Domain systems and backup hosts. An additional routing rule in the Kernel IP routing table and in the Data Domain system Net Route Config list shows a list of static routes reapplied at each system boot. Each interface is assigned a route based on its assigned address.

In addition, depending on the default gateway subnet and the gateway owner, a route is added to an interface automatically if the interface is in the subnet of a default gateway address. If the address is an IPv4 type, a routing table is created for the interface and default routes for that address are set up in that table.

net route guidelines and restrictions

Changes to Ethernet interfaces made with `net route` command options flush the routing table. All routing information is lost and data movement using routing is cut off immediately. You should make interface changes only during a scheduled downtime. You must also reconfigure routing rules and gateways after making interface changes.

```
net route add [ipversion {ipv4 | ipv6}] route spec
```

Add an IPv4 or IPv6 static route for a network or network host. Role required: admin, limited-admin.

Arguments and definitions**ipv4address**

The default gateway's IP address. It can only be an IPv4 address type.

interface name

The name of the interface for adding the default gateway. This makes this default gateway a "targeted" default gateway, which means this default will be used to route traffic for the IP address on this interface as long as the address is in the same subnet as this default gateway.

```
net route add [ipversion {ipv4 | ipv6}] [type {fixed | floating}] route spec
```

Add a fixed or floating static route in a high-availability (HA) system. Role required: admin, limited-admin.

Arguments and definitions**fixed**

Specifies that the static route is fixed.

floating

Specifies that the static route is floating.

gw gateway

Specifies the IP address of the gateway to use to reach the destination network or host. If no gateway is specified, the route uses the default gateway.

ipv4address

Specifies the IPv4 address for the destination network or host. If no gateway is specified, the command fails if the destination host is not found on the local network or through the default gateway.

ipv6address

Specifies that the route is for IPv6 routing. This argument is not required when an IPv6 address is specified.

-netmask

Specifies the network mask that applies to the destination network or network host.

type

Specifies the type of static route is either a fixed or floating IP.

Note

Except for the cases where specific gateways are used to get to specific addresses or set of addresses, it is recommended to use the gateway as a default gateway instead of specified in a static route.

Example 133

The following example shows an IPv4 route added to network 192.168.1.0 with netmask 255.255.255.0 using the `srvr12` gateway.

```
# net route add 192.168.1.0 netmask 255.255.255.0 gw srvr12
```

Example 134

The following example shows an IPv4 route added to network 192.168.1.0 with netmask 255.255.255.0.

```
# net route add 192.168.1.0 netmask 255.255.255.0 gw srvr12 table teth5a
```

Example 135

The following example shows an IPv4 route added to a host named `user24` through the `srvr12` gateway.

```
# route add user24 gw srvr12
```

```
net route add gateway ipv4address [interface name]
```

Add a gateway address to the list of gateway addresses on the Data Domain system. Optionally specify a specific interface to associate with the gateway address. If the gateway is unreachable, the system displays a warning, but still adds the gateway. Role required: admin, limited-admin.

If the same route is needed from multiple NICs, then consider adding static gateways:

```
net route add gateway gateway IP interface NIC name
```

Example 136

The following example shows the addition of default gateways on specific NICs.

```
# net route add gateway 192.168.1.2 interface eth0b
# net route add gateway 192.168.1.2 interface eth1b
```

Arguments and definitions**ipv4address**

The default gateway's IP address. It can only be an IPv4 address type.

interface name

The name of the interface for adding the default gateway. This makes this default gateway a "targeted" default gateway, which means this default will be used to route traffic for the IP address on this interface only as long as the address is in the same subnet as this default gateway.

```
net route del [ipversion {ipv4 | ipv6}] route spec
```

Delete an IPv4 or IPv6 static route for a network or network host. Role required: admin, limited-admin.

The IPv4 route spec syntax is: `ipv4address [netmask] gw gateway`

The IPv6 route spec syntax is: `ipv6address gw gateway`

```
net route del gateway ipv4address [interface name]
```

Deletes the specified gateway or routing table along with associated route entries and route rules. If the gateway is "targeted" (associated with a specific interface), the interface must be also given. Role required: admin, limited-admin.

Arguments and definitions

ipv4address

The default gateway's IP address. It can only be an IPv4 address type.

interface name

If this is a targeted default gateway, the associated interface also needs to be given for the targeted gateway to be deleted. Otherwise, the interface will be deleted from the list of the added default gateways.

```
net route reset [ipversion {ipv4 | ipv6}]
```

Delete the configured routing gateway for the specified protocol. If no protocol is specified, the IPv4 gateway configuration is deleted. Role required: admin, limited-admin.

```
net route set gateway {ipaddr | ipv6addr}
```

Configure the IP address for the IPv4 or IPv6 default gateway. When the default gateway is added or changed, the Data Domain operating system automatically adds a route to default gateway for each interface with the same subnet. Role required: admin, limited-admin.

Note

When configuring an IPv6 address, a command failure might not produce an error message in the CLI. If the new gateway is not visible using the `route show gateway` and `route show table` commands, check the messages log file for information on why the command failed.

Example 137

The following example shows the device configured at 192.168.10.1 the default IPv4 gateway.

```
# net route set gateway 192.168.10.1
```

```
net route show config [routing-table-name name]
```

Display the configured static routes. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
net route show gateways [detailed] [ipversion {ipv4 | ipv6}]
[<ipv4addr>]
```

Displays the configured or DHCP-supplied IPv4 and IPv6 gateways as specified. The `detailed` option displays the network interface or type, associated routing tables, interface addresses, and owners if applicable. If no IP version is specified, both

gateways are displayed. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
net route show tables [table-name-list | ipversion {ipv4 | ipv6}]
```

Displays the IPv4 and IPv6 routing tables as specified. If no IP version is specified, both table versions are displayed. For IPv4, a list of routing tables is displayed; the routing within each table is displayed and the rules associated with the table is displayed. The main IPv4 routing table is displayed in the same form as was done previously, but the other tables are displayed in the same form as the output from the `ip route list` command. The output for IPv6 versions is the same as it was previously. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
# net route show tables ipversion ipv4

IP Routing Tables and IDs in the Kernel
254   main
1     teth0b
2     teth0a
3     teth1a.555

    Table: main
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref    Use Iface
0.0.0.0          10.25.128.1     0.0.0.0          UG    0      0      0 eth0a
10.25.128.0      0.0.0.0         255.255.240.0    U      0      0      0 eth0a
10.25.160.0      0.0.0.0         255.255.240.0    U      0      0      0 eth0b
127.0.0.0        0.0.0.0         255.0.0.0        U      0      0      0 lo
172.16.32.0      0.0.0.0         255.255.240.0    U      0      0      0 eth1b.100
172.16.144.0     0.0.0.0         255.255.240.0    U      0      0      0 eth1b
172.16.208.0     0.0.0.0         255.255.240.0    U      0      0      0 eth1a.200
172.16.240.0     0.0.0.0         255.255.240.0    U      0      0      0 eth1a.100
172.17.48.0      0.0.0.0         255.255.240.0    U      0      0      0 eth1a.555
192.168.112.0    0.0.0.0         255.255.255.0    U      0      0      0 eth1d

    Routing rules:
32766:    from all lookup main

    Table: teth0a
Kernel IP routing table:
default via 10.25.128.1 dev eth0a
10.25.128.0/20 dev eth0a scope link src 10.25.142.166
    Routing rules:
32764:    from all oif eth0a lookup teth0a
32765:    from 10.25.142.166 lookup teth0a

    Table: teth0b
Kernel IP routing table:
default via 10.25.160.1 dev eth0b
10.25.160.0/20 dev eth0b scope link src 10.25.167.241
    Routing rules:
32762:    from all oif eth0b lookup teth0b
32763:    from 10.25.167.241 lookup teth0b

    Table: teth1a.555
Kernel IP routing table:
default via 172.17.55.1 dev eth1a.555
172.17.48.0/20 dev eth1a.555 scope link src 172.17.55.55
    Routing rules:
32760:    from all oif eth1a.555 lookup teth1a.555
32761:    from 172.17.55.55 lookup teth1a.555
```

```
# net route show tables ipversion ipv4

IP Routing Tables and IDs in the Kernel
```

net

```
254 main
1 teth0b
2 teth0a
3 teth1a.555
```

```
Table: main
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
0.0.0.0 10.25.128.1 0.0.0.0 UG 0 0 0 eth0a
10.25.128.0 0.0.0.0 255.255.240.0 U 0 0 0 eth0a
10.25.160.0 0.0.0.0 255.255.240.0 U 0 0 0 eth0b
127.0.0.0 0.0.0.0 255.0.0.0 U 0 0 0 lo
172.16.32.0 0.0.0.0 255.255.240.0 U 0 0 0 eth1b.100
172.16.144.0 0.0.0.0 255.255.240.0 U 0 0 0 eth1b
172.16.208.0 0.0.0.0 255.255.240.0 U 0 0 0 eth1a.200
172.16.240.0 0.0.0.0 255.255.240.0 U 0 0 0 eth1a.100
172.17.48.0 0.0.0.0 255.255.240.0 U 0 0 0 eth1a.555
192.168.112.0 0.0.0.0 255.255.255.0 U 0 0 0 eth1d
```

```
Routing rules:
32766: from all lookup main
```

```
Table: teth0a
Kernel IP routing table:
default via 10.25.128.1 dev eth0a
10.25.128.0/20 dev eth0a scope link src 10.25.142.166
```

```
Routing rules:
32764: from all oif eth0a lookup teth0a
32765: from 10.25.142.166 lookup teth0a
```

```
Table: teth0b
Kernel IP routing table:
default via 10.25.160.1 dev eth0b
10.25.160.0/20 dev eth0b scope link src 10.25.167.241
```

```
Routing rules:
32762: from all oif eth0b lookup teth0b
32763: from 10.25.167.241 lookup teth0b
```

```
Table: teth1a.555
Kernel IP routing table:
default via 172.17.55.1 dev eth1a.555
172.17.48.0/20 dev eth1a.555 scope link src 172.17.55.55
```

```
Routing rules:
32760: from all oif eth1a.555 lookup teth1a.555
32761: from 172.17.55.55 lookup teth1a.555
```

```
net route show [ipversion {ipv4 | ipv6}] [type {fixed |
floating}]
```

Displays the type of IP address as specified.

```
net route trace ipv4addr | ipv6addr | {hostname [ipversion
{ipv4 | ipv6}]}
```

Displays a route used by a Data Domain system to connect with the specified destination. Role required: admin, limited-admin, security, user, backup-operator, or none.

To trace the route to `svr24`:

```
# net route trace svr24
Traceroute to svr24.yourcompany.com (192.168.1.6), 30 hops max, 38
byte packets
1 svr24 (192.168.1.6) 0.163 ms 0.178 ms 0.147 ms
```

net set

```
net set {domainname local-domain-name | searchdomains search-domain-list}
```

Set the domain name or search domains used by the Data Domain system. The default for `domainname` is the return from DHCP, or domain portion of the hostname configured with `net set hostname`. The default for `searchdomains` is the domain name configured with `net set domainname`. The configured domain name is always included in the list of search domains. Role required: admin, limited-admin.

Example 138

```
# net set domainname yourcompany-ny.com
# net set searchdomains yourcompany2.com, yourcompany3.com
```

The `searchdomains` list is `yourcompany-ny.com`, `yourcompany2.com` and `yourcompany3.com`.

If the domain names provided cannot be resolved, a warning appears.

```
net set dns ipv4-ipv6-addr-list
```

Set the DNS server list using addresses for IP version 4, IP version 6, or both. Separate the IP addresses with a comma or a space. This command overwrites the current list of DNS servers. Only servers included in the most recently issued command are available to a Data Domain system. Role required: admin, limited-admin.

Example 139

```
# net set dns 10.0.0.1, 10.0.0.2, 10.0.0.3
The Name (DNS) server list is:
    10.0.0.1, 10.0.0.2, 10.0.0.3
```

Example 140

```
# net set dns 2100:bad:cafe:f00d::1:101 10.24.255.146
10.24.255.150
The Name (DNS) server list is:
    2100:bad:cafe:f00d::1:101, 10.24.255.146, 10.24.255.150
```

```
net set hostname host
```

Set the hostname of the Data Domain system. If you do not statically set the hostname, the system uses a DHCP hostname from one of the system interfaces. If multiple interfaces have DHCP hostnames, then during some DD OS upgrades, the system hostname might change to a hostname from a different interface. The best practice is to use this command to statically set the system hostname.

Note that some browsers may prevent logins to the host if the hostname contains an underscore. Data Domain recommends using hostnames without underscores to ensure the GUI can recognize and manage the host. Role required: admin, limited-admin.

Note

If the Data Domain system is using CIFS with Active Directory authentication, changing the hostname causes the Data Domain system to drop out of the domain. Use the `cifs set authentication` command option to rejoin the Active Directory domain.

Note

This command accepts domain names and validates that the domain name is made up of valid characters separated by periods. Although an IPv4 address passes the validation for a domain name, this command does not recognize the IP address as such and does not validate the IP address. This is not an issue for IPv6 addresses because they contain colon characters, which are invalid in host names.

```
net set hostname ha-system
```

Promote the hostname of the system to be the HA system name for the HA pair.

```
net set portnaming {slot-based | legacy}
```

Change the port naming scheme. Role required: admin, limited-admin.

net show

```
net show {domainname | searchdomains}
```

Display the domain name or search domains used for email sent by a Data Domain system. Role required: admin, limited-admin, security, user, backup-operator, or none.

Example 141

```
# net show domainname
The Domainname is: emc.com
```

```
# net show searchdomains
#   Searchdomains
-   -----
1   emc.com (local domain)
-   -----
```

```
net show all
```

Display all networking information produced by the other `net show` commands and a Network Stats table. Role required: admin, limited-admin, security, user, backup-operator, or none.

Output Definitions

Most of the command output is described for other `net show` commands. The following definitions are for the columns in the Network Stats table.

Foreign Address

The connection IP address and port used on a destination device or application.

Local Address

The connection IP address and port used on the local system or application.

Proto

Protocol in use for the listed connection. This is always TCP because UDP is a connectionless protocol.

Recv-Q

A count of the protocol packets in the receive queue.

Send-Q

A count of the protocol packets in the send queue.

State

The state of the connection as signaled by the TCP protocol. The state for active sessions is ESTABLISHED. The TIME_WAIT state appears when the local connection is closing and reserving the port number for a wait period in case any additional packets arrive. The SYN_RCVD, SYN_SENT, and CLOSE_WAIT states are usually so brief that they do not appear, but if one of these states do appear in several consecutive command displays, it might indicate a connection problem.

```
net show config [ifname]
```

Display the configuration for a specific Ethernet interface. Exclude the keyword *ifname* to view the configuration for all Ethernet interfaces. This command also shows auto-generated IPv6 addresses, which are automatically generated and assigned to the base interface. Role required: admin, limited-admin, security, user, backup-operator, or none.

Example 142

```
# net show config
eth0a    Link encap:Ethernet  HWaddr 00:8C:FA:08:92:19
         inet  addr:10.110.141.187  Bcast:10.110.143.255  Mask:255.255.248.0
         inet6 addr: 2620:0:170:1a04:28c:faff:fe08:9219/64 Scope:Global
         inet6 addr: fe80::28c:faff:fe08:9219/64 Scope:Link
         UP BROADCAST NOTRAILERS RUNNING MULTICAST  MTU:1500  Metric:1
         RX packets:6891161 errors:0 dropped:0 overruns:0 frame:0
         TX packets:339319 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:726690965 (693.0 MiB)  TX bytes:102448918 (97.7 MiB)
```

Output Definitions**Bcast**

IPv4 network broadcast address.

Collisions

Network collisions.

HWaddr

MAC address.

inet addr

IPv4 network address.

inet6 addr

IPv6 network address. An interface can have multiple IPv6 IP addresses.

Link encap

Link encapsulation used, typically Ethernet.

Mask

IPv4 network mask.

MTU

Maximum transfer unit.

RX bytes

Bytes of data received.

RX packets

Network packets received.

TX bytes

Bytes of data transmitted.

TX packets

Network packets transmitted.

txqueuelen

Transmit queue length.

```
net show dns
```

Display a list of DNS servers used by the Data Domain system. The final line in the output shows if the servers were configured manually or by DHCP. Role required: admin, limited-admin, security, user, backup-operator, or none.

Example 143

```
# net show dns
#   Server
-   -----
1   10.24.255.146
2   10.24.255.150
3   10.110.188.5
-   -----
Showing DNS servers configured manually.
```

```
net show hardware
```

Display Ethernet port hardware information from the kernel. Role required: admin, limited-admin, security, user, backup-operator, or none.

Figure 1 Output: net show hardware

```
# net show hardware
Port      Speed      Duplex      Supp Speeds  Hardware Address  Physical  Link Status  State
-----
eth0a     1000Mb/s   full        10/100/1000  00:8c:fa:05:6c:91  Copper    yes          running
eth0b     unknown    unknown     10/100/1000  00:8c:fa:05:6c:90  Copper    unknown     down
```

Output Definitions**Duplex**

Full, half, or unknown duplex protocol. Unknown means the interface is not available.

Hardware Address

The MAC address.

Link Status

The status is `yes` if the link is receiving carrier from the remote system and `no` if no carrier is present. Carrier must be present for the link to support data transfer. The status is `unknown` when the link is administratively down and the link state cannot be determined.

Physical

Copper, DA Copper, or Fibre.

Port

The Ethernet interfaces on the system.

Speed

The actual speed at which the port processes data.

State

The port state indicates whether the port is administratively `up` or `down` and whether the link is ready for traffic. A port in the `running` state is enabled and receiving carrier from the remote system, so it is ready to send and receive data.

Supp Speeds

Lists speeds the port is capable of using.

```
net show hostname
```

Display the hostname of the Data Domain system. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
net show settings
```

Display the registry settings for the network interfaces available on the system. Settings include the name for the physical, virtual, VLAN, and alias interfaces. The settings also indicate if an interface is enabled and if the interface information is from DHCP using IPv4 or IPv6. A type field for interconnect or floating, and additional settings such as bonding information are also included. Role required: admin, limited-admin, security, user, backup-operator, or none.

Figure 2 Output: net show settings

```
sysadmin@dd860-81# sysadmin@dd860-81# net show settings
port      enabled  state   DHCP  IP address          netmask      type  additional setting
-----  -
eth0a     yes      running ipv4   10.110.143.200*    255.255.248.0*  n/a
          /64
          fe80::28c:faff:fe05:6c91**
eth0a:20  yes      running no     10.110.154.201    255.255.255.0  n/a
eth0b     no       down    ipv4   n/a                n/a            n/a
-----  -
* Value from DHCP
** auto_generated IPv6 address
```

net

Figure 3 Output: net show settings for HA (active)

```
SE@apollo-ha7a-p1(active:1)## net show settings
port      enabled  state  DHCP  IP address                               netmask  type  additional setting
-----  -
ethMa     yes     running  ipv4  10.25.132.160*                          255.255.240.0*  n/a
          /64
          2620:0:170:1106:260:16ff:fe5c:8f70**
          fe80::260:16ff:fe5c:8f70**
ethMb     no      down    ipv4  n/a                                       n/a           n/a
ethMc     no      down    ipv4  n/a                                       n/a           n/a
ethMd     no      down    ipv4  n/a                                       n/a           n/a
eth1a     yes     running  no    1.1.1.1                                  255.255.255.0  floating
          /64
          2620:0:170:1106:260:16ff:fe52:14a4**
          fe80::260:16ff:fe52:14a4**
eth1a.1   yes     running  no    2.2.2.1                                  255.255.255.0  floating
          /64
          fe80::260:16ff:fe52:14a4**
eth1a.1:1 yes     running  no    3.3.3.1                                  255.255.255.0  floating
eth1b     no      down    no    n/a                                       n/a           n/a
eth1c     no      down    no    n/a                                       n/a           n/a
eth1d     no      down    no    n/a                                       n/a           n/a
eth4a     no      down    no    n/a                                       n/a           n/a
eth4b     no      down    no    n/a                                       n/a           n/a
eth4c     no      down    no    n/a                                       n/a           n/a
eth4d     no      down    no    n/a                                       n/a           n/a
eth11a    yes     running  n/a   n/a                                       n/a           interconnect  bonded to veth99
eth11b    yes     running  n/a   n/a                                       n/a           interconnect  bonded to veth99
eth11c    yes     running  n/a   n/a                                       n/a           interconnect  bonded to veth99
eth11d    yes     running  n/a   n/a                                       n/a           interconnect  bonded to veth99
veth99    yes     running  no    d:d:d:d:0060:1651:efe0                    / (not specified)  interconnect  lacp hash xor-L3L4:
          fe80::260:16ff:fe51:efe0**
          /64
```

Figure 4 Output: net show settings for HA (standby)

```
SE@apollo-ha7a-p0(standby:0)## net show set
port      enabled  state  DHCP  IP address                               netmask  type  additional setting
-----  -
ethMa     yes     running  ipv4  10.25.132.148*                          255.255.240.0*  n/a
          /64
          2620:0:170:1106:260:16ff:fe5c:8f7c**
          fe80::260:16ff:fe5c:8f7c**
ethMb     no      down    ipv4  n/a                                       n/a           n/a
ethMc     no      down    ipv4  n/a                                       n/a           n/a
ethMd     no      down    ipv4  n/a                                       n/a           n/a
eth1a     yes     running  no    n/a                                       n/a           floating
          /64
          2620:0:170:1106:260:16ff:fe52:1780**
          fe80::260:16ff:fe52:1780**
eth1a.1   yes     running  no    n/a                                       n/a           floating
          /64
          fe80::260:16ff:fe52:1780**
eth1a.1:1 yes     running  no    n/a                                       n/a           floating
eth1b     no      down    no    n/a                                       n/a           n/a
eth1c     no      down    no    n/a                                       n/a           n/a
eth1d     no      down    no    n/a                                       n/a           n/a
eth11a    yes     running  n/a   n/a                                       n/a           interconnect  bonded to veth99
eth11b    yes     running  n/a   n/a                                       n/a           interconnect  bonded to veth99
eth11c    yes     running  n/a   n/a                                       n/a           interconnect  bonded to veth99
eth11d    yes     running  n/a   n/a                                       n/a           interconnect  bonded to veth99
veth99    yes     running  no    d:d:d:d:0060:1651:fda0                    / (not specified)  interconnect  lacp hash xor-L3L4:
          fe80::260:16ff:fe51:fda0**
          /64
```

Output Definitions

DHCP

The DHCP configuration for the interface, which is ipv4 (enabled for IPv4), ipv6 (enabled for IPv6), disabled (no), or not applicable (n/a).

Enabled

The target state of the interface, which is yes (enabled) or no (disabled).

Floating

The floating keyword means that the IP address is a floating type. If the column does not indicate a floating type, then it is a fixed IP address.

IP address

The IPv4 and IPv6 addresses that are assigned to the interface. The auto-generated IPv6 addresses are followed by one asterisk (*).

Interconnect

Used for internal communication between HA nodes.

Netmask/prefix length

The IPv4 network mask or IPv6 addresses prefix assigned to the interface.

Port

The Ethernet interfaces on the system. Interface eth1d represents the physical interface. Interface eth1d:10 is an alias interface that adds an IPv4 address to the base interface, and eth1d:100 is an alias interface that adds an IPv6 address to the same base interface.

State

The port state indicates whether the port is administratively `up` or `down` and whether the link is ready for traffic. A port in the `running` state is enabled and receiving carrier from the remote system, so it is ready to send and receive data.

Type

The label assigned to the interface with the `net config ifname type` command.

```
net show stats [[ipversion {ipv4 | ipv6}] [all | listening]
[detailed] |[ipversion {ipv4 | ipv6}] route | interfaces |
[ipversion {ipv4 | ipv6}] statistics]
```

Display network statistics. Role required: admin, limited-admin, security, user, backup-operator, or none.

Argument Definitions**all**

Lists local client connections for the TCP and UDP protocols. Also lists client and server connections for the UNIX protocol.

detailed

Adds the associated processes for each connection.

interfaces

Displays a table of the driver statistics for each interface that is UP.

ipversion {ipv4 | ipv6}

Limits the display output to IPv4 or IPv6 statistics only. When this option is omitted, the system shows all statistics.

ipversion {ipv4 | ipv6}] route

Displays the route table (default is the IPv4 main table only).

listening

Lists local server TCP connections.

statistics

Displays the statistics for IP, IP extended, ICMP, TCP, TCP extended, UDP, and UDP Lite.

net tcpdump

```
net tcpdump capture filename [interface iface] [{host host
[ipversion {ipv4 | ipv6}] | net {ipaddr [mask mask] |
ipv6addr[/prefixlength]}]}] [port port] [snaplen bytes]
```

Capture data, and then copy the output file to another system for analysis. This command converts the options from the command line to equivalent `tcpdump` options. Output files are placed in `/ddvar/traces` where you can upload them to autosupport. Values for *bytes* may be followed by the K, M, or G to scale the value accordingly. A maximum of 10 output files may be retained on the system. If this limit is reached, you are prompted to delete some of the files. Role required: admin.

Argument Definitions

filename

Specifies the output filename. Equivalent Linux argument: `-w /ddvar/traces/tcpdump_filename -C 100 -W 5`.

host *host* [ipversion {ipv4 | ipv6}]

Equivalent Linux argument: `host host`.

interface *iface*

Equivalent Linux argument: `-i iface`.

ipv6addr/ *prefixlength*

IPv6 address.

net {*ipaddr* [mask *mask*]}

Equivalent Linux arguments:

- net *net*
- mask *mask*

port *port*

Equivalent Linux argument: `port port`

snaplen *bytes*

Equivalent Linux argument: `-s bytes`

```
net tcpdump del {filename | all}
```

Delete output files created by the `net tcpdump capture` command. Specify a *filename* to delete files matching the pattern `/ddvar/traces/tcpdump_filename*`. Specify `all` to remove all `net tcpdump` output files. Role required: admin.

net troubleshooting

```
net troubleshooting duplicate-ip
```

Detect duplicate IP addresses in the local network. Role required: admin, security, user, backup-operator, or none.

CHAPTER 30

nfs

The `nfs` command enables you to add NFS clients and manage access to a Data Domain system. It also enables you to display status information, such as verifying that the NFS system is active, and the time required for specific NFS operations.

This chapter contains the following topics:

- [nfs change history](#)294
- [nfs add](#) 294
- [nfs del](#) 297
- [nfs disable](#) 297
- [nfs enable](#)297
- [nfs export add](#)298
- [nfs export create](#) 300
- [nfs export del](#)301
- [nfs export destroy](#)302
- [nfs export modify](#)302
- [nfs export rename](#)303
- [nfs export show](#) 303
- [nfs option](#)305
- [nfs reset](#)306
- [nfs show](#)307
- [nfs status](#)308

nfs change history

Modified argument and modified behavior in DD OS 6.1.2

```
nfs export create [export-name] path path [clients client-list
[options option-list]] [referral referral-name remote-servers
address-list [remote-path path]]
```

The command supports adding clients to the NFS export at the same time the export is created.

nfs add

```
nfs add path client-list [(option-list)]
```

Add NFS clients that can access the Data Domain system. A client can be a fully qualified domain hostname, class-C IP address, IP address with netmask or length, IPv6 address, NIS netgroup name with the prefix @, or an asterisk wildcard for the domain name, such as *.yourcompany.com. Role required: admin, limited-admin.

An asterisk by itself means no restrictions. A client added to a subdirectory under /data/coll/backup has access only to that subdirectory.

The *options-list* is comma or space separated, enclosed by parentheses. If no option is specified, the default options are `rw`, `root_squash`, `no_all_squash`, and `secure`.

The following wildcards are supported:

- ? means any single character.
- * represents any number of characters.

NFS Options

ro

Enable read-only permission.

rw

Enable read and write permissions (default value).

root_squash

Map requests from uid or gid 0 to the anonymous uid/gid (default value).

no_root_squash

Turn off root squashing.

all_squash

Map all user requests to the anonymous uid/gid.

no_all_squash

Turn off the mapping of all user requests to the anonymous uid/gid (default value).

default_root_squash

force_minimum_root_squash

secure

Require that requests originate on an Internet port that is less than IPPORT_RESERVED (1024) (default value).

insecure

Turn off the secure option.

anonuid=*id*

Set an explicit user ID for the anonymous account. The ID is an integer bounded from 0 to 65535.

anongid=*id*

Set an explicit group ID for the anonymous account. The ID is an integer bounded from 0 to 65535.

sec

Set sec equal to one or more of the following options to activate different types of authentication security options. The default for sec is sys.

sys: Allow unauthenticated connections. Select to not use authentication. This is the default.

krb5: Allow Kerberos-5 NFS authenticated connections.

krb5i: (krb5 integrity) Allow connections that checksum NFS arguments and results.

krb5p: (krb5 privacy) Allow connections that encrypt NFS arguments and results.

Note

You can use any combination of the sec options. Security options are colon separated.

⚠ CAUTION

If authentication options (sec options) on the DDR are selected and a client tries to connect to the DDR without setting the respective setting(s) on the client, the client will be denied with an authentication failure. If multiple authentication options are present for an export, the clients will be able to mount the export using any one of the specified authentication options.

Example 144

```
nfs add path client-list (sec=sys:krb5:krb5i:krb5p)
```

Add NFS clients for a path using all of the security options.

Example 145

```
nfs add / backup * (sec=krb5:krb5i:krb5p:sys)
```

Export backup to all users so that any client can access the mount point, and all of the security options will be activated.

Example 146

Example 146 (continued)

To add an NFS client with an IP address of 192.168.1.02 and read/write access to /backup with the secure option, enter:

```
# nfs add /data/col1/backup 192.168.1.02
```

Example 147

To add a subnet client using its IP address followed by a length and a netmask, enter:

```
# nfs add /data/col1/test-mtree 192.168.1.02/24
# nfs add /data/col1/test-mtree 192.168.1.02/255.255.255.0
```

Example 148

To add an NFS client with an IPv6 address of 2620:0:170:1a01:250:56ff:fe8d:c6ae and read/write access to /data/col1/test-mtree with the secure option, enter:

```
# nfs add /data/col1/test-mtree 2620:0:170:1a01:250:56ff:fe8d:c6ae
```

Example 149

To add a client that uses an IPv6 address followed by a prefix length, enter:

```
# nfs add /data/col1/test-mtree 2620:0:170:1a01:250:56ff:fe8d:c6ae/64
```

Note

IPv6 addresses do not use subnet masks.

Example 150

To add an NFS export for /data/col1/test_su

```
# nfs add /data/col1/test_su *
```

```
nfs show clients
```

path	client	options
/data/col1/test_su	*	(rw,no_root_squash,no_all_squash,secure)

To modify the /data/col1/test_su export by changing rw to ro, and secure to insecure:

```
nfs add /data/col1/test_su * (ro,insecure)
```


Example 150 (continued)**Note**

You must include a space between the client (in this case "*") and the opening parenthesis.

```
nfs show clients
```

path	client	options
/data/coll/test_su	*	(ro,no_root_squash,no_all_squash,insecure)

nfs del

```
nfs del path client-list
```

Delete an export for a given set of client(s). The *client-list* can contain IPv4 and IPv6 addresses, hostnames, or an asterisk that represents all clients. Role required: admin, limited-admin.

Example 151

To delete an NFS client with an IP address of 192.168.1.01 from the `/ddvar` directory, enter:

```
# nfs del /ddvar 192.168.1.01
```

Example 152

To delete an NFS client with an IPv6 address of 2620:0:170:1a01:250:56ff:fe8d:c6ae from the `/ddvar` directory, enter:

```
# nfs del /ddvar 2620:0:170:1a01:250:56ff:fe8d:c6ae
```

nfs disable

```
nfs disable
```

Disable the NFS server, effectively disabling access from the clients. Role required: admin, limited-admin.

nfs enable

```
nfs enable
```

Allow all NFS-defined clients to access the Data Domain system. Role required: admin, limited-admin.

nfs export add

```
nfs export add {<export-spec> | all} clients <client-list>
[options <option-list>]
```

Add a client or list of clients to one or more exports. A client can be a fully qualified domain hostname, a class-C IP address, an IP address with netmask or length, an IPV6 address, an NIS netgroup name with the prefix @, or an asterisk wildcard for the domain name such as *.yourcompany.com. Role required: admin, limited-admin.

An asterisk by itself means no restrictions.

The *options-list* is comma separated and enclosed by quotes if more than one option is provided. If no option is specified, the default options are `rw`, `no_root_squash`, `no_all_squash`, and `secure`.

Note

NFSv4 has the same export options as those that exist for NFSv3.

NFS options

version

Select the appropriate version or versions of NFS, which can be 3, 4, 3:4, or all.

ro

Enable read-only permission.

rw

Enable read and write permissions (default value).

root_squash

Map requests from uid or gid 0 to the anonymous uid/gid.

no_root_squash

Turn off root squashing.

Note

`no_root_squash` is the default value.

all_squash

Map all user requests to the anonymous uid/gid.

no_all_squash

Turn off the mapping of all user requests to the anonymous uid/gid (default value).

default_root_squash

force_minimum_root_squash

secure

Require that requests originate on an Internet port that is less than IPPORT_RESERVED (1024) (default value).

insecure

Turn off the secure option.

anongid=*id*

Set an explicit user ID for the anonymous account. The ID is an integer bounded from 0 to 65635.

sec

Set sec equal to one or more of the following options to activate different types of authentication security options. The default for sec is sys.

sys: Allow unauthenticated connections. Select to not use authentication. This is the default.

krb5: Allow Kerberos-5 NFS authenticated connections.

krb5i: (krb5 integrity) Allow connections that checksum NFS arguments and results.

krb5p: (krb5 privacy) Allow connections that encrypt NFS arguments and results.

Note

You can use any combination of the sec options. Security options are colon separated.

⚠ CAUTION

If authentication options (sec options) on the DDR are selected and a client tries to connect to the DDR without setting the respective setting(s) on the client, the client will be denied with an authentication failure. If multiple authentication options are present for an export, the clients will be able to mount the export using any one of the specified authentication options.

```
nfs export add {<export-spec> | all} referral <referral-name>
remote-servers <address-list> [ remote-path <path> ]
```

Add a referral location to the export defined in <export-spec>. Role required: admin.

<referral-name> defines the name of the referral. If the name you choose is the same as an existing referral, you will see an error message. Each export can have multiple referrals, each with a unique name.

When adding <referral-name> or <referral-list>, consider the following guidelines:

- Both <referral-name> and <referral-list> can accept embedded spaces; if you use embedded spaces for an item, that item must be contained within double quotation marks.
 - Specify a single referral name as ref1 without quotation marks if the name contains no spaces.
 - Specify a single referral name as “ref1 space” with quotation marks if the name contains spaces.
 - Use quotation marks with lists of items whether or not they have embedded spaces; for example, “ref1,ref2” and “ref1,ref3 space”

remote servers <address-list> defines the remote network address or addresses to be used in the referral. The following must be true for each export:

- Each referral location must refer to only one NFS server, although the server can contain multiple network addresses.

- Each NFS server should be associated with only one referral location.

`remote-path <path>` allows you to specify a remote path name. If you do not specify a path, the current export path is used.

In the following example, you would add a referral to a single remote server for an existing export and use a different path on the referred system:

```
# nfs export add db_backups referral db_backups
  remote-servers db_backups.domain.name
  remote-path /data/coll/db_backups2
```

In the following example, you would create referrals for several exports using referral export locations with two remote addresses for the same server. You would also use the same path on the server that you use on the Data Domain system:

```
# nfs export add dd_backups,dd_locks
  referral backups2
  remote-servers db_backups2a.domain.name,
                db_backups2b.domain.name
```

In the following example, you would create referrals for several exports using two referral locations to indicate two different servers, and using the same path on each server. Compare this with the previous example:

```
# nfs export add "dd_backups,dd_locks" referral backups2
  remote-servers db_backups2.domain.name
NFS referral added.

# nfs export add "export1,export2" referral referral1 remote-servers
test.domain.com remote-path /test1
2 NFS referrals added.
```

If you try to create a referral with a duplicate location, you will see an error message:

```
# nfs export add "dd_backups,dd_locks"
  referral backups2
  remote-servers db_backups2a.domain.name,
                db_backups2b.domain.name
NFS referral entry(s) added.

# nfs export add "dd_backups,dd_locks"
  referral backups2
  remote-servers db_backups2a.domain.name,
                db_backups2b.domain.name
**** Referral 'backups2' already exists.
```

nfs export create

```
nfs export create [export-name] path path [clients client-list
[options option-list]] [referral referral-name remote-servers
address-list [remote-path path]]
```

Create a named export and add a path. If you do not provide an export name, the name simply defaults to the path. Use the `clients` parameter to optionally add a client or list of clients to the export. A client can be a fully qualified domain hostname,

a class-C IP address, an IP address with netmask or length, an IPV6 address, an NIS netgroup name with the prefix @, or an asterisk wildcard for the domain name such as *.yourcompany.com. Role required: admin.

In the following example, a named export is created with one added client:

```
# nfs export create path /data/coll/new_data clients
emc.datadomain.com options version=all
NFS export '/data/coll/new_data' created.
```

nfs export del

```
nfs export del {<export-spec> | all} clients {<client-list> |
all}
```

Removes a client or a list of clients from existing exports. You can remove a single client or a list of clients, with the name of each client separated by a comma. Role required: admin.

Note

If either referral lists or client lists have comma separators in them, the entire list must be enclosed in double quotes.

```
nfs export del {<export-spec> | all} referrals {<referral-list>
| all}
```

Removes specified NFSv4 referrals. You can remove a single referral or a list of referrals, with the name of each referral separated by a comma. If a referral in a referral list does not exist on one or more of the specified exports, you will see an error message and the exports will remain unchanged.

If more than one referral is given (with each separated by a comma), the entire list must be enclosed within double quotation marks. If you specify `all`, all referrals are removed from the indicated exports.

Role required: admin.

Delete all referrals for every export:

```
# nfs export del all referrals all
<count> NFS referral(s) deleted.
```

Delete all referrals for the export `db_backups`:

```
# nfs export del db_backups referrals all
<count> NFS referral(s) deleted.
```

Delete the specific referral `backups1`:

```
# nfs export del db_backups referrals backups1
<count> NFS referral(s) deleted.
```

Delete a referral that does not exist, and the system indicates the nonexistent referral cannot be found:

```
# nfs export del db_backups referrals backups1UUU
**** Referral 'backups1UUU' was not found.
```

nfs export destroy

```
nfs export destroy {<export-spec> | all}
Destroys one or multiple NFS exports. Role required: admin.
```

nfs export modify

```
nfs export modify {<export-spec> | all} clients {<client-list>
| all} options <option-list>
```

Updates an existing client or clients to an export or set of exports identified in {**export-spec** | **all**}. Role required: admin.

```
nfs export modify {<export-spec> | all} referral <referral-
name> [remote-servers <address-list>] [remote-path {<path> |
default}]
```

Updates an existing referral to an export or set of exports identified in {**export-spec** | **all**}.

- If the referral specified in **<referral-name>** does not exist in one or more of the specified exports, you see an error message and no change occurs. Similarly, if the requested specific client does not exist, you see an error message and no change occurs.
- **remote-servers** defines the remote network address or addresses to be used in the referral. It replaces the existing remote network address list, if used.
- If **remote-path** is specified, enter the remote path name in **<path>**; otherwise, the current path is unchanged. If you use the default instead of a specific path, the export path is used.

Role required: admin.

In the following example, a referral is modified to a single remote server for an existing export:

```
# nfs export modify db_backups referral db_backups
remote-servers db_backups3.<domain-name>
<count> NFS referral(s) modified.
```

In the following example, you can see a referral modified to a single remote server for an existing export, but with an invalid referral:

```
# nfs export modify db_backups referral db_backups
remote-servers db_backups3.mycorp.com
****Referral 'db_backups' was not found.
```

nfs export rename

```
nfs export rename <export-name> <new-export-name>
```

Rename a specific export. Role required: admin.

nfs export show

```
nfs export show list [<export-spec>] [path <path-spec>]
[clients <client-list>] [tenant-unit <tenant-unit>]
```

Enables you to view a list of exports. Role required: admin, limited-admin.

Note

The NFS data path security feature filters the Linux 'showmount' output on the client to match the client permissions in the export list. The system displays only the client's activity. Because NFSv4 does not use the mountd daemon, NFSv4 exports are not listed.

The following example output shows a list of all exports:

```
# nfs export show list
Export Path # Client Tenant-Unit
          Entries
-----
/data /data 0 -
/ddvar /ddvar 0 -
finance /data/coll/m1 3 -
-----
(3 exports found)
```

The following example output shows a specific list of exports that share the prefix `hr`:

```
# nfs export show list hr*
Export Path # Client Tenant-Unit
          Entries
-----
hr1 /data/coll/m2 3 -
hr2 /data/coll/m4 0 -
-----
(2 exports found)
```

The following example output shows information related to two specific clients, `c1` and `c4`:

```
# nfs export show list clients "c1,c4"
Export Path # Client Tenant-Unit
          Entries
-----
finance /data/coll/m1 3 -
hr1 /data/coll/m2 3 -
-----
(2 exports found)
```

The following example displays existing referrals for the export `db_backups`:

```
# nfs export show detailed db_backups
NFS Export: db_backups
  Path: /data/col1/backups_15
  Tenant-Unit: -
  NFSv3 Mounts: 200
  Active NFSv3 clients: 300
  NFSv4 Clients instances: 100
  Active NFSv4 clients: 22

Client      Options
-----
oradb.mycorp.com (rw,secure,root_squash,sec=krb5i)
*           (ro,insecure,root_squash)
-----
Total Client Entries: 2

Referrals:
Name Remote Path      Remote Servers
-----
db2 /mnt/data1/db2 ddr1a.myco.com,ddr1b.myco.com
db4 /mnt/data1/db4 ddr17.myco.com
-----
Total Referral locations: 2
```

```
nfs export show detailed [<export-spec>] [path <path-spec>]
[clients <client-list>] [tenant-unit <tenant-unit>]
```

Allows you to apply filters to selectively view certain exports, clients, paths, and tenant-units. Some filters accept limited wildcards. Role required: admin.

The following shows detailed information for an export called `finance`:

```
# nfs export show detailed finance
NFS Export: finance
  Path: /data/col1/m1
  Tenant-Unit: -
  NFSv3 Mounts: 0
  Active NFSv3 clients: 0
  Active NFSv4 clients: 0

Client      Options
-----
1.1.1.1    (sec=sys,rw,root_squash,no_all_squash,secure,version=3)
2.2.2.2    (sec=sys,rw,root_squash,no_all_squash,secure,version=4)
3.3.3.3    (sec=sys,ro,root_squash,no_all_squash,secure,version=3:4)
-----
Total Client Entries: 3

No referrals found.
```

```
nfs export show stats [<export-spec>] [interval <secs>] [count
<count>]
```

Shows NFS export statistics. The interval is an optional number of seconds with a minimum of 1. The count is an optional ordinal value with a minimum of 1. Role required: admin, limited-admin.

```
nfs export show summary [tenant-unit <tenant-unit>]
```

Show summary information for NFS exports. Role required: admin, limited-admin.

nfs option

```
nfs option reset default-server-version
```

Reset the NFS server to NFSv3. Role required: admin, limited-admin.

```
nfs option set default-server-version
```

Set the NFS server to use NFSv4 as the default. Role required: admin, limited-admin.

```
# nfs option set default-server-version 4
NFS option 'default-server-version' set to '4'.
```

```
nfs option show
```

Show the NFS version that is currently used as the default. Role required: admin, limited-admin.

```
# nfs option show
Option                                     Value
-----
default-export-version                    3
default-server-version                    3
nfs4-grace-period                          30
nfs4-lease-interval                       300
mountd-port                               2052
nfs4-port                                 2049
nfs3-port                                 2049
nfs4-domain                               brs.lab.emc.com
nfs4-idmap-out-numeric                     map-first
nfs4-idmap-active-directory                disabled
nfs4-acls                                  disabled
default-root-squash                       enabled
force-minimum-root-squash-default         disabled
```

Output Definitions

default-export-version

The default version or versions for client exports. This takes effect for future client exports only.

Any legal version string is allowed ("3," "3:4", "all").

default-server-version

The NFS server version or versions enabled by default.

Any legal version string is allowed ("3," "3:4", "all").

nfs4-grace-period

The grace period for NFSv4 recovery measured in seconds.

The minimum is 5 seconds; the maximum is 120 seconds.

The default is 30 seconds.

nfs4-lease-interval

The client lease interval measured in seconds.

The minimum is 120 seconds, the maximum is 3600 seconds.

The default is 300 seconds.

mountd-port

The IP port for mountd. Changing this port requires an NFS server restart.

The default port is 2052.

nfs4-port

The IP port for the NFSv4 server. Changing this port requires an NFS server restart.

The default port is 2049.

nfs3-port

The IP port for the NFSv3 server and related protocols. Changing this port requires an NFS server restart.

The default port is 2049.

nfs4-domain

The NFSv4 sever domain. Any valid domain name is permitted.

The Data Domain DNS domain name is the default; "" is the default if the domain name is not set.

nfs4-idmap-out-numeric

Set output mapping of NFSv4 owner/group ids (e.g. fred@emc.com) as numeric values or names in output attributes and ACL ACE entries.

The default is map-first; use numeric ID mapping if normal mapping fails.

Use numeric ID mapping if allowed. Numeric IDs are never sent; if mapping fails, the server sends the ID "nobody".

nfs4-idmap-active-directory

Determine whether NFSv4 should use CIFS active directory (AD) for name resolution and ID mapping.

Disabled is the default setting.

Active-Directory mapping may be used to increase interoperability in a mixed CIFS/NFS environment.

nfs4-acls

Determine whether NFSv4 ACLs (access control lists) are enabled.

Disabled is the default setting.

nfs reset

```
nfs reset clients
```

Removes the existing client/share configuration, resetting the client list to the factory default (empty). In non-interactive mode, for example when the command is run as part of a script, the system will not pause. However, in interactive mode, the command warns the user and asks for confirmation before proceeding. NFS clients can access the Data Domain system when the client list is empty. Role required: admin, limited-admin.

Note

In interactive mode, the system will prompt the user with the following warning message:

```
This command will delete all exports and client configurations.
Do you want to proceed? (yes|no) [no]
```

```
nfs reset stats
```

Clear the NFS statistics. Role required: admin, limited-admin.

nfs show

```
nfs show active [tenant-unit tenant-unit]
```

List clients active in the past 15 minutes and the mount path for each. Optionally, list NFS clients assigned to a tenant-unit. Role required: admin, limited-admin, user, backup-operator, security, tenant-user, tenant-admin.

Note

The NFS data path security feature filters the Linux 'showmount' output on the client to match the client permissions in the export list. The system displays only the client's activity. Because NFSv4 does not use the mountd daemon, NFSv4 exports are not listed.

Argument definitions

tenant-unit (Optional)

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system.

```
nfs show clients [tenant-unit tenant-unit]
```

Lists NFS clients allowed to access the Data Domain system and the mount path as well as NFS options for each. Optionally, list NFS clients assigned to a tenant-unit. A client added using a hostname is displayed using the client's hostname. Security options and the log option are displayed for each mount point. If client is added using a hostname, and both sides support IPv6 and IPv4, then the client can connect using both addresses. Role required: admin, limited-admin, user, backup-operator, security, tenant-user, tenant-admin.

Note

The NFS data path security feature filters the Linux `showmount` output on the client to match the client permissions in the export list; the system does not display output that is not relevant to the client.

When you run the `showmount` command with the tenant-unit hostname, you see only the exports that tenant-unit owns.

```
nfs show detailed-stats
```

Display NFS cache entries and status to facilitate troubleshooting. Role required: admin, limited-admin, user, backup-operator, security.

```
nfs show histogram
```

Display NFS operations in a histogram. Users with user role permissions may run this command. Role required: admin, limited-admin, user, backup-operator, security.

Output Definitions

mean (ms)

The mathematical mean time for completion of the operations.

std-dev

The standard deviation for time to complete operations, derived from the mean time.

max

The maximum time taken for a single operation.

min

The minimum time taken for a single operation.

2ms

The number of operations that took 2 ms or less.

4ms

The number of operations that took between 2ms and 4ms.

6ms

The number of operations that took between 4ms and 6ms.

8ms

The number of operations that took between 6ms and 8ms.

10ms

The number of operations that took between 8ms and 10ms.

100ms

The number of operations that took between 10ms and 100ms.

1s

The number of operations that took between 100ms and 1 second.

10s

The number of operations that took between 1 second and 10 seconds.

>10s

The number of operations that took over 10 seconds.

`nfs show port`

Display NFS port information. Role required: admin, limited-admin, user, backup-operator, security.

`nfs show stats`

Display NFS statistics, including NFS Kerberos (only) related GSSAPI (Generic Security Services API) statistics. Role required: admin, limited-admin, user, backup-operator, security.

nfs status

`nfs status`

Enter this option to determine if the NFS system is operational. When the filesystem is active and running, the output shows the total number of NFS requests since the filesystem started, or since the last time that the NFS statistics were reset.

CHAPTER 31

ntp

The `ntp` command synchronizes a Data Domain system with an NTP time server, manages the NTP service, or turns off the local NTP server.

A Data Domain system can use a time server supplied through the default multicast operation, received from Dynamic Host Configuration Protocol (DHCP), or set manually with the Data Domain system `ntp add` command.

This chapter contains the following topics:

• ntp change history	310
• ntp guidelines and restrictions	310
• ntp add	310
• ntp del	310
• ntp disable	310
• ntp enable	311
• ntp reset	311
• ntp show	311
• ntp status	311

ntp change history

There have been no changes to this command in this release.

ntp guidelines and restrictions

- Default system settings for NTP service are enabled and multicast.
- Time servers set with the `ntp add` command override time servers from DHCP and from multicast operations.
- Time servers from DHCP override time servers from multicast operations.
- The Data Domain system `ntp del` and `ntp reset` commands act only on manually added time servers, not on DHCP-supplied time servers. You cannot delete DHCP time servers or reset to multicast when DHCP time servers are supplied.

ntp add

```
ntp add timeserver server-name
```

Add a remote time server hostname to the NTP timeserver list. Role required: admin, limited-admin. This command option requires security officer authorization for Retention Lock Compliance systems.

Example 153

To add an NTP time server named `svr26.yourcompany.com` to the list, enter:

```
# ntp add timeserver svr26.yourcompany.com
```

ntp del

```
ntp del timeserver server-name
```

Delete a manually added time server hostname from the NTP server list. Role required: admin, limited-admin. This command option requires security officer authorization for Retention Lock Compliance systems.

Example 154

To delete an NTP time server named `svr26.yourcompany.com` from the list, enter:

```
# ntp del timeserver svr26.yourcompany.com
```

ntp disable

```
ntp disable
```

Disable NTP service on a Data Domain system. Role required: admin, limited-admin. This command option requires security officer authorization for Retention Lock Compliance systems.

ntp enable

```
ntp enable
```

Enable NTP service on a Data Domain system. Role required: admin, limited-admin. This command option requires security officer authorization for Retention Lock Compliance systems.

ntp reset

```
ntp reset
```

Reset the NTP configuration to the default settings. Role required: admin. This command option requires security officer authorization for Retention Lock Compliance systems.

```
ntp reset timeservers
```

Reset the time server list from manually entered time servers to DHCP time servers (if supplied) or to the multicast mode (if no DHCP time servers supplied). Role required: admin, limited-admin. This command option requires security officer authorization for Retention Lock Compliance systems.

ntp show

```
ntp show config
```

Display whether NTP is enabled or disabled and show the time server list. Role required: admin, limited-admin, security, user, backup-operator, or none.

ntp status

```
ntp status
```

Display the local NTP service status, time, and synchronization information. Role required: admin, limited-admin, security, user, backup-operator, or none.

ntp

CHAPTER 32

qos

The `qos` command displays, modifies, or resets the value of the Random I/O throttle.

This chapter contains the following topics:

- [qos change history](#)314
- [qos randomio](#) 314

qos change history

There have been no changes to this command in this release.

qos randomio

```
qos randomio throttle reset
```

Reset the Random I/O throttle to its default value of 40 percent. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
qos randomio throttle set percent
```

Set the Random I/O throttle to a percent value from 1 to 100, where 1 allocates the fewest resources for Random I/O workloads and 100 allocates the most resources. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
qos randomio throttle show
```

Display the current value of the Random I/O throttle. Role required: admin, limited-admin, security, user, backup-operator, or none.

CHAPTER 33

quota

The `quota` command lets you modify the amount of storage space for MTrees and for VTL and DD Boost storage units. There are two quota limits: hard and soft. The hard limit prevents writes from exceeding the quota. An error message is issued if the hard limit is exceeded. The soft limit allows writes to exceed the quota. However, an alert is generated if this happens. The soft limit value must be less than the hard limit value. Quota limit values must be specified as integers.

You can set a hard limit, a soft limit, or both, depending on your requirements. For example, an administrator may choose to enforce only a soft limit to prevent overnight backup jobs from failing when the quota limit is reached. Or the administrator may choose to enforce only a hard limit to block a user from writing when the quota limit is reached.

Snapshots capture quota information at a precise point in time. Usage tracking in the active file system does not account for the space of a snapshot, so quota limits are not enforced on snapshots.

This chapter contains the following topics:

• quota change history	316
• quota capacity	316
• quota disable	317
• quota enable	317
• quota reset	317
• quota set	318
• quota show	318
• quota status	318
• quota streams	318

quota change history

There have been no changes to this command in this release.

quota capacity

```
quota capacity disable
```

Disable capacity quota. Also disables MTree quota limits and restores the limits to the default state (unlimited). Role required: admin, limited-admin.

```
quota capacity enable
```

Enable capacity quota. Role required: admin, limited-admin.

```
quota capacity reset { all | mtrees mtree-list | storage-units
storage-unit-list } [soft-limit] [hard-limit]
```

Reset capacity quota limits. Both the *mtree-list* and the *storage-unit-list* are colon-separated lists. If hard or soft limits are not entered, both are reset to the default state (unlimited). Role required: admin, limited-admin.

To reset hard and soft limits for an MTree:

```
# quota capacity reset mtrees /data/coll/backup1
```

To reset only a soft limit for an MTree:

```
# quota capacity reset mtrees /data/coll/backup1 soft-limit
```

To reset only a hard limit for an MTree:

```
# quota capacity reset mtrees /data/coll/backup3 hard-limit
```

To reset hard and soft limits for a storage unit:

```
# quota capacity reset storage-units DDBOOST_STRESS_SU
```

```
quota capacity set {all | mtrees mtree-list | storage-units
storage-unit-list} {soft-limit n {MiB|GiB|TiB|PiB} | hard-limit
n {MiB|GiB|TiB|PiB} | soft-limit n {MiB|GiB|TiB|PiB} hard-limit
n {MiB|GiB|TiB|PiB}}
```

Set capacity quota limits during runtime for multiple MTrees (*mtree-list* is a colon-separated list). When used for storage units (*storage-unit-list* is a colon-separated list), this sets limits only after the storage unit is created. Note that the quota feature must be enabled, because limits are otherwise not enforced. Setting quotas does not require disabling the file system and therefore does not affect system performance. Role required: admin, limited-admin.

To set a soft limit quota of 10 GiB on MTree `/data/coll/backup1` when the quota feature is disabled:

```
# quota capacity set mtrees /data/coll/backup1 soft-limit 10 GiB
```

To set a hard limit quota of 10 TiB on MTree `/data/coll/backup1`:

```
# quota capacity set mtrees /data/coll/backup1 hard-limit 10 GiB
```

To set a soft limit quota of 100 GiB and a hard limit quota of 1 TiB on MTree `/data/coll/backup1`:

```
# quota capacity set mtrees /data/coll/backup1 soft-limit 10 GiB hard-limit 10 TiB
```

To set a soft limit quota of 100 GiB and a hard limit quota of 1 TiB on storage-unit `DDBOOST_STRESS_SU`:

```
# quota capacity set storage-units DDBOOST_STRESS_SU soft-limit 100 GiB hard-limit 1 TiB
```

```
quota capacity show {all | mtrees mtree-list | storage-units storage-unit-list | tenant-unit tenant-unit}
```

List capacity quotas and usage of a particular MTree (*mtree-list* is a colon-separated list) or storage unit (*storage-unit-list* is a colon-separated list), all mtrees or storage units, or all of both. The unit of display for usage and limits is MiB. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

```
quota capacity status
```

Display status of capacity quota enforcement: enabled or disabled. If output includes a note stating that status is disabled, capacity quota limits are not being enforced and are therefore unlimited. Role required: admin, limited-admin, security, user, backup-operator, none.

quota disable

```
quota disable - deprecated
```

This command is deprecated. Use `quota capacity disable` instead. Role required: admin, limited-admin.

quota enable

```
quota enable - deprecated
```

This command is deprecated. Use `quota capacity enable` instead. Role required: admin, limited-admin.

quota reset

```
quota reset - deprecated
```

This command is deprecated. Use `quota capacity reset` instead. Role required: admin, limited-admin.

quota set

quota set - deprecated

This command is deprecated. Use `quota capacity set` instead. Role required: admin, limited-admin.

quota show

quota show - deprecated

This command is deprecated. Use `quota capacity show` instead. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-admin, tenant-user.

quota status

quota status - deprecated

This command is deprecated. Use `quota capacity status` instead. Role required: admin, limited-admin, security, user, backup-operator, none.

quota streams

```
quota streams reset storage-units storage-unit-list [write-
stream-soft-limit] [read-stream-soft-limit] [repl-stream-soft-
limit] [combined-stream-soft-limit] [hard-stream-limit n]
```

Reset streams quota soft limits. The *storage-unit-list* is a colon-separated list. Note that this command controls the same stream limits as `ddboost storage-unit modify`. Role required: admin, limited-admin.

Example 155

```
# quota streams reset storage-units sul write-stream-soft-limit read-
stream-soft-limit repl-stream-soft-limit combined-stream-soft-limit

sul: Stream soft limits: write=none, read=none, repl=none,
combined=none
```

```
quota streams set storage-units storage-unit-list [write-
stream-soft-limit n ] [read-stream-soft-limit n ] [repl-stream-
soft-limit n ] [combined-stream-soft-limit n ] [hard-stream-
limit n]
```

Set streams quota soft limits. The *storage-unit-list* is a colon-separated list. Note that this command controls the same stream limits as `ddboost storage-unit modify`. Role required: admin, limited-admin.

Example 156

```
# quota streams set storage-units sul write-stream-soft-limit 10 read-
stream-soft-limit 3 repl-stream-soft-limit 10 combined-stream-soft-
limit 10

sul: Stream soft limits: write=10, read=3, repl=10, combined=10
```

```
quota streams show {all | storage-unit storage-unit | tenant-
unit tenant-unit}
```

List streams quotas for all storage units or tenant units. Or list streams quotas for a specific storage or tenant unit. Role required: admin, limited-admin, security, user, backup-operator, tenant-user, tenant-admin, none.

Example 157

When SMT is enabled and tenant-units exist, this example displays filtering by tenant-unit tu1.

```
# quota streams show tenant-unit tu1

Tenant-unit: tu1
Storage Unit  Write Streams  Read Streams  Repl Streams  Combined Streams
                Soft-Limit    Soft-Limit    Soft-Limit    Soft-Limit
-----
su1           none          none         none         none
su2           none          none         none         none
-----
DD System Stream Limits: write=16 read=4 repl-in=20 repl-out=20 combined=16
```

Example 158

When SMT is enabled and tenant-units exist, but not all storage-units are contained in tenant-units, this example displays streams quotas for all storage-units, where some storage-units are in tenant-units, and some storage-units are not in any tenant-unit.

```
# quota streams show all

Storage Unit  Write Streams  Read Streams  Repl Streams  Combined Streams
                Soft-Limit    Soft-Limit    Soft-Limit    Soft-Limit
-----
su4           none          none         none         none
su5           none          none         none         none
su6           none          none         none         none
-----
Tenant-unit: tu1
Storage Unit  Write Streams  Read Streams  Repl Streams  Combined Streams
                Soft-Limit    Soft-Limit    Soft-Limit    Soft-Limit
-----
su1           none          none         none         none
su2           none          none         none         none
-----

Tenant-unit: tu2
Storage Unit  Write Streams  Read Streams  Repl Streams  Combined Streams
                Soft-Limit    Soft-Limit    Soft-Limit    Soft-Limit
-----
su3           none          none         none         none
-----
DD System Stream Limits: write=16 read=4 repl-in=20 repl-out=20 combined=16
```

quota

CHAPTER 34

replication

Data Domain Replicator lets you replicate data (copy and synchronize) between two Data Domain systems: a source and a destination. Source and destination configurations, or pairs, are also known as “contexts.” Depending on your objective, you can replicate entire sites, specific directories, MTrees, or files. Replication is a licensed software option. See the *Data Domain Operating System Administration Guide* for details on replication practices and procedures.

This chapter contains the following topics:

• replication change history	322
• replication abort	322
• replication add	322
• replication break	324
• replication dir-to-mtree	324
• replication disable	325
• replication enable	325
• replication initialize	325
• replication modify	326
• replication option	327
• replication reauth	328
• replication recover	328
• replication resync	328
• replication show	329
• replication status	333
• replication sync	334
• replication throttle	334
• replication watch	336

replication change history

There have been no changes to this command in this release.

replication abort

```
replication abort recover destination
```

Stop a recover process. Run this on the destination Data Domain system only. Then, reconfigure replication on the source Data Domain system and restart the recover process. Role required: admin, limited-admin.

```
replication abort resync destination
```

Stop a resync operation. Run this on the source or destination Data Domain system. In case of a directory replication context, run it both on the source and the destination. Role required: admin, limited-admin.

replication add

```
replication add source source destination destination [low-bw-
optim {enabled | disabled}] [encryption {enabled |
[authentication-mode {one-way | two-way | anonymous |
disabled}] [propagate-retention-lock {enabled | disabled}]
[ipversion {ipv4 | ipv6}] [max-repl-streams n] [destination-
tenant-unit tenant-unit]
```

Create a replication pair, which can be for Collection, MTree, or Directory Replication. If the *destination* exists, you will get an error, and you must either delete it or rename it before proceeding. If a source or destination name does not correspond to a Data Domain network name, run `replication modify connection-host` on the source system. When entering names that include spaces or special characters, enclose the entire pathname with double quotation marks, or enter a backslash before the space, but do not use both. A file or a directory may not be renamed or moved into or out of a source. This includes a “cut” operation followed by a “paste” operation in Windows. After replication is initialized, ownership and permissions of the destination are always identical to those of the source. If the context is configured, the destination is kept in a read-only state and can receive data only from the source. Role required: admin, limited-admin.

Example 159 Collection Replication

- The storage capacity of the destination system must be equal to, or greater than, that of the source system. If the destination capacity is less than that of the source, the available capacity on the source is reduced to that of the destination.
- The destination must have been destroyed and subsequently created, but not enabled.
- Each destination and each source can be in only one context at a time.

In this example, notice the prefix `co1` to the URL signifying Collection Replication. The source hostname is `system-dd1`, and the destination hostname is `system-dd2`.

```
# replication add source col://system-dd1.chaos.local destination
col://system-dd2.chaos.local
```

Example 160 MTree Replication

- 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 Extended Retention systems to non-Extended Retention systems if both are running DD OS 5.5 or later.
- The destination Data Domain system must have available storage capacity of at least the post-compressed size of the expected maximum post-compressed size of the source directory or MTree.
- When replication is initialized, a destination MTree is created automatically.
- A Data Domain system can simultaneously be the source for one context and the destination for another context.

In this example, notice the prefix `mtree` to the URL signifying MTree Replication. The source MTree path is `/data/col1/mtree1`, the destination MTree path is `/data/col1/dstmtree1`, the maximum number of replication streams is 6, and the destination Tenant Unit is `tu1`.

```
# replication add source mtree://system-dd1.chaos.local/data/col1/mtree1 destination mtree://system-dd2.chaos.local/data/col1/dstmtree1 max-repl-streams 6 destination-tenant-unit tu1
```

Example 161 Directory Replication

- The destination Data Domain system must have available storage capacity of at least the post-compressed size of the expected maximum post-compressed size of the source directory or MTree.
- When replication is initialized, a destination directory is created automatically.
- A Data Domain system can simultaneously be the source for one context and the destination for another context.

In this example, notice the prefix `dir` to the URL signifying Directory Replication. The source directory name is `dir1`, and it resides in the `/backup` MTree (the default MTree).

```
# replication add source dir://system-dd1.chaos.local/backup/dir1 destination dir://system-dd2.chaos.local/backup/dir1
```

Argument Definitions

authentication-mode {anonymous | one-way | two-way | disabled}

Lets you choose an authentication-mode. If the mode is not specified, anonymous is the default. One-way indicates that only the destination certificate is certified. Two-way indicates that both the source and destination certificates are verified.

Note

Mutual trust must be established before you can use authentication-mode. The `adminaccess trust` section provides more details about establishing mutual trust.

destination-tenant-unit *tenant-unit*

Lets you specify a Tenant Unit only on the destination.

encryption {enabled | disabled}

Enables or disables *encryption over wire*. Both the source and the destination must enable this feature. Encrypted replication uses the ADH-AES256-SHA cipher suite.

low-bw-optim {enabled | disabled}

Enables or disables *low bandwidth optimization*, which improves data transfer over low bandwidth links by adding increased data compression to optimize network bandwidth. Both the source and the destination must enable this feature.

Low bandwidth optimization is not supported if the DD Extended Retention software option is enabled on either Data Domain system. It is also not supported for Collection Replication.

max-repl-streams *n*

The maximum number of replication streams allowed, which must be between 1 and the maximum streams per context for a given Data Domain system model, and is supported only for MTree Replication.

propagate-retention-lock {enabled | disabled}

Enables or disables the propagation of Retention Lock. This cannot be enabled for Directory Replication.

ipversion {ipv4 | ipv6}

Lets you choose your network preference for the replication pair. 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.

replication break

```
replication break {destination | all}
```

Remove the source or destination Data Domain system from a replication pair, or remove all Replicator configurations from a Data Domain system. The `all` option breaks all the replication contexts on the system without individual confirmation prompts. Role required: `security` for Retention Lock Compliance systems; `admin` and `limited-admin` for all other systems.

replication dir-to-mtree

```
replication dir-to-mtree abort source
```

Abort the directory-to-MTree migration process for the specified context. The command stops the ongoing migration and performs the necessary cleanup. When the

process is complete, the MTree replication context and the associated MTrees on both source and destination system are deleted. Role required: admin, limited-admin.

```
replication dir-to-mtree start from source to destination
```

Perform the directory-to-MTree migration from the directory replication context to the MTree replication context. Role required: admin, limited-admin.

```
replication dir-to-mtree status [source | all]
```

Shows you the status of the directory-to-MTree replication for the specified context or contexts. This command allows you to instantly see the status of fastcopy or replication initialization operations. Role required: admin, limited-admin.

```
replication dir-to-mtree watch destination
```

Displays the progress of the directory-to-MTree migration. You can see the percentage of the initialization that is complete and track the virtual and physical bytes that are transferred. Role required: admin, limited-admin.

replication disable

```
replication disable {destination | all}
```

Disable replication. Run this on the source or destination system to halt data replication temporarily. If run on the source, the operation stops sending data to the destination. If run on the destination, the operation stops serving the active connection from the source. Role required: security for Retention Lock Compliance systems; admin, limited-admin for all other systems.

replication enable

```
replication enable {destination | all}
```

Restart replication. If run on the source, the operation resumes sending data to the destination. If run on the destination, the operation resumes serving the active connection from the source. Role required: admin, limited-admin.

replication initialize

```
replication initialize destination
```

Initialize replication. Run this on the source to start replication between a source and destination and to verify that configuration and connections, including checking for a matching tenant on both sides if a *destination-tenant-unit* was set for this context. Error messages are returned if problems appear. Initialization can take several hours, or days, depending on the amount of data in the source. To reduce initialization time, consider placing both Data Domain systems of the replicator pair in the same location with a direct link. The *destination* variable is required. Key-manager settings on a destination are ignored when users set up and initialize a collection replication pair. The keys are copied to the replica, but key-manager settings are not. If the destination is configured with key-manager settings prior to becoming the replication destination, the settings remain on the system but are not used. If a collection replication breaks, you must reconfigure the destination to use the correct key-manager settings and key class. If possible, reset the key-manager on the destination prior to collection replication, and then configure the destination with the correct key manager-server and key class after a collection replication is broken. Role required: admin, limited-admin.

replication modify

```
replication modify destination {source-host | destination-host}
new-host-name
```

Modify the source or destination host name. In this case, you must modify the replication configuration on both the source and the destination; that is, if the host name that changed was the destination, you must run replication modify on both the destination and the source so both sides will be updated. The *new-host-name* must be the name returned by `net show hostname` on the system receiving the new host. When using replication modify, always run `filesystems disable` or `replication disable` first, and conclude with `filesystems enable` or `replication enable`. Then, run `replication show config` to make sure all changes were done. Role required: admin, limited-admin.

```
replication modify destination connection-host new-host-name
[port port]
```

Modify the destination host name, when it does not resolve for the connection, to a new host name or IP address. You may also specify an optional port number. This action may be required when a connection passes through a firewall. It is definitely required when connecting to an alternate listen-port on the destination. It may also be required after adding a new source and destination pair, or after renaming a source or a destination. Role required: admin, limited-admin.

Example 162

If local destination `ca.company.com` is moved from California to New York, run the following on both the source and the destination:

```
# replication disable
# replication modify dir://ca.company.com/backup/dir2 destination-
host ny.company.com
# replication enable
# replication show config
```

```
replication modify destination crepl-gc-bw-optim {enabled |
disabled}
```

Modify the collection replication bandwidth optimization option. The default value is enabled. Disable this option in a high bandwidth environment to enhance throughput. Role required: admin, limited-admin.

```
replication modify destination destination-tenant-unit tenant-
unit
```

Modify the destination Tenant Unit. Note that after the replication context has been initialized, the Tenant Unit for the replica MTree can be modified only by using `mtree modify`. Role required: admin, limited-admin.

Example 163

```
replication modify mtree://ip2/data/coll/mtr1 destination-tenant-
unit tu2
```

```
replication modify destination encryption {enabled | disabled}
```

Modify the state of encryption over wire for the destination. This feature is active only when enabled on both the source and the destination. Role required: admin, limited-admin.

```
replication modify destination ipversion {ipv4 | ipv6}
```

Modify the network preference for the destination. 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. Role required: admin, limited-admin.

```
replication modify destination low-bw-optim {enabled | disabled}
```

Modify the state of low bandwidth optimization for the destination. This feature is active only when enabled on both the source and the destination. This feature is not supported for collection replication or if DD Extended Retention is enabled on either the source or the destination. Role required: admin, limited-admin.

```
replication modify destination max-repl-streams n
```

Modify the number of maximum replication streams allowed, which must be between 1 and the maximum streams per context for a given Data Domain system model, and is supported only for MTree Replication. Role required: admin, limited-admin.

Example 164

```
replication modify mtree://ip2/data/coll/mtrl max-repl-streams 6
"max-repl-streams" changed to 6 for replication context mtree://ip2/
data/coll/mtrl.
```

replication option

```
replication option reset {bandwidth | delay | listen-port |
default-sync-alert-threshold}
```

Reset system bandwidth, delay, listen port, and sync-alert-threshold to default values. Defaults are bandwidth, unlimited; delay, none; listen-port, 2051. Default for sync-alert-threshold is 24 (hours). When using replication option reset, always run `filesys disable` first, and conclude with `filesys enable`. Role required: admin, limited-admin.

```
replication option set bandwidth rate
```

Set the network bandwidth rate for the Data Domain system. You must set the bandwidth and network delay on each side of the connection. Role required: admin, limited-admin.

```
replication option set default-sync-alert-threshold value
```

Set the sync time to configure when an alert is generated. The sync time is set in hours. The default *value* is 24. Role required: admin, limited-admin.

```
replication option set delay value
```

Set the network delay in milliseconds for the Data Domain system. You must set the bandwidth and network-delay on each side of the connection. Role required: admin, limited-admin.

```
replication option set listen-port value
```

Set the listen port for the Data Domain system. On a destination Data Domain system, set the port from which the destination receives data from replication sources (the default is 2051). A destination can have only one listen port used by all sources. The connection-host port used by a source must match the listen port used by the destination. For DD Boost managed file replication, the listen port is used on the source Data Domain system and on the destination Data Domain system to specify the connection-host port. For directory replication, `replication modify`

`connection-host` is used on the source Data Domain system. Role required: admin, limited-admin.

`replication option show`

Display the current bandwidth, network-delay settings, listen port, and sync-alert-threshold. If these settings are configured using default values, `replication option show` returns a command prompt with no setting information. Role required: admin, limited-admin, security, user, backup-operator, none.

replication reauth

`replication reauth destination`

Reset authentication on the source and destination systems. The *destination* variable is required. Messages similar to `Authentication keys out of sync` or `Key out of sync` indicate a reset is required. Reauthorization is primarily used when replacing a source Data Domain system. Role required: admin, limited-admin.

replication recover

`replication recover destination`

Recover replication. Run this on a new source to move data from a destination system. If configuring collection replication, this must be run on the new source only. The *destination* argument is required. This is not available for MTree replication. When using `replication recover`, always run `filesys disable` first, and conclude with `filesys enable`. If `replication break` was previously run, the destination cannot be used to recover a source. If configuring directory replication, the destination directory on the source must be empty. Role required: admin, limited-admin.

replication resync

`replication resync destination`

Bring back into sync (or recover) the data between a source and destination replication pair after a manual break. The replication pair are resynchronized so both endpoints contain the same data. Resynchronization is available for Directory, MTree or Pool Replication, but not for Collection Replication.

Before running `replication resync`, you must run `replication add` to add the source and destination back on the system.

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.

Note the following about using `replication resync` with DD Retention Lock:

- If the destination MTree or directory contains retention-locked files that do not exist on the source, then resync will fail.
- If the destination directory has retention lock enabled, but the source directory does not have retention lock enabled, then a resync of a directory replication will fail.
- With MTree replication, resync will succeed if the source MTree does not have retention lock enabled while the destination MTree has retention lock enabled or

vice versa, as long as the destination MTree does not contain retention-locked files not present on the source.

Role required: admin, limited-admin.

replication show

```
replication show config [destination | tenant-unit tenant-unit
| all]
```

Show replication configuration. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-user, tenant-admin.

Output Definitions

Connection Host and Port

A source system connects to the destination system using the name returned by the hostname command on the destination. It may also connect using a destination name or IP address and port designated by `replication modify connection-host`. The destination hostname may not resolve to the correct IP address when connecting to an alternate interface on the destination, or when passing through a firewall.

Ipversion

The IP version - either IPv4 or IPv6.

Low-bw-optim

The status of low-bandwidth optimization: enabled, disabled, or configuration mismatch.

Crepl-gc-bw-optim

The status of Collection Replication bandwidth optimization: enabled, disabled. The default value is enabled. Disable this optimization to enhance throughput in a high bandwidth environment.

Encryption

The replication process is enabled and available for encryption (yes) or disabled and not available for encryption (no).

Enabled

The replication process is enabled and available to replicate data (yes) or disabled and not available to replicate data (no).

Propagate-retention-lock

The retention lock process is available (enabled) or not available (disabled).

Max-repl-streams

The maximum number of replication streams allowed.

Tenant-unit (if SMT enabled)

The local Tenant Unit to which the local MTree belongs or (-) if the local MTree does not belong to any Tenant Unit.

```
replication show detailed-history {obj-spec-list | tenant-unit
tenant-unit | all} [duration hr] [interval hr]
```

Show details of replication performance history. The *obj-spec-list* is a space- or comma-separated list. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-user, tenant-admin.

Output Definitions**Pre-Comp (KB) Written**

The number of logical bytes ingested to the source corresponding to the CTX.

Pre-Comp (KB) Remaining

For directory replication only, this is the sum of file sizes remaining to be replicated for the context. Output includes the entire logical size of the current file being replicated. If a large file is being replicated, this number may take a lot of time to change. The number changes only after the current file finishes.

Replicated (KB) Pre-compressed

The amount of pre-compressed data replicated.

Replicated (KB) Post-synthetic-optim

The amount of data replicated after synthetic optimization was applied.

Replicated (KB) Post-filtered

The amount of data replicated after identity filtering (dedup).

Replicated (KB) Post-low-bw-optim

The amount of data replicated after delta compression (low-bandwidth optimization).

Replicated (KB) Post-local-comp

The amount of data replicated after local compression.

Replicated (KB) Network

The amount of data replicated over the wire.

Sync-as-of Time

The time when the most recently replicated data on the destination was generated on the source. A value of unknown appears during replication initialization.

```
replication show detailed-stats [destination | tenant-unit
tenant-unit | all]
```

Display cumulative statistics beginning from when the context was created. Byte-count statistics are provided related to identity filtering, delta compression, and local compression. The ratio of the values Bytes after filtering by destination to Bytes after low bandwidth optimization gives additional compression ratio supplied by delta compression. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-user, tenant-admin.

Output Definitions**Network bytes sent to destination**

The number of physical bytes sent to the destination over the wire.

Pre-compressed bytes written to source

The number of bytes received by the source, including logical bytes associated with the file being replicated.

Pre-compressed bytes sent to destination

The number of bytes sent to the destination, including logical bytes associated with the file being replicated.

Bytes after synthetic optimization

The number of bytes still needed to send/receive after synthetic replication optimization.

Bytes after filtering by destination

The number of bytes sent after identity filtering (dedup).

Bytes after low bandwidth optimization

The number of bytes sent after delta compression (low-bandwidth optimization).

Bytes after local compression

The number of bytes sent after local compression.

Pre-compressed bytes remaining

For directory replication only, this is the sum of file sizes remaining to be replicated for the context. Output includes the entire logical size of the current file being replicated. If a large file is being replicated, this number may take a lot of time to change. The number changes only after the current file finishes.

Compression ratio

The ratio of the value of logical bytes ingested to the source to physical bytes actually sent to the destination over the wire.

Sync'ed-as-of Time

The time when the most recently replicated data on the destination was generated on the source. A value of unknown appears during replication initialization.

```
replication show history {obj-spec-list | tenant-unit tenant-unit | all} [duration hr] [interval hr]
```

Show replication performance history. The *obj-spec-list* is a space- or comma-separated list. Statistics are generated hourly. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-user, tenant-admin.

Output Definitions**Pre-Comp (KB) Written**

The number of logical bytes ingested to the source corresponding to the CTX.

Pre-Comp (KB) Remaining

The amount of pre-compression data not replicated.

Replicated (KB) Pre-Comp

The amount of pre-compressed data replicated.

Replicated (KB) Network

The amount of compressed data sent over the network.

Low-bw-optim

The additional compression ratio supplied by delta compression (low-bandwidth optimization).

Sync-as-of Time

The time when the most recently replicated data on the destination was generated on the source. A value of unknown appears during replication initialization.

```
replication show performance {obj-spec-list | tenant-unit
tenant-unit | all} [interval sec] [count count]
```

Display current replication activity. The *obj-spec-list* is a space- or comma-separated list. Default interval is two seconds. If a single source context is specified, four additional columns are presented. These columns show the relative amounts of time spent working on, or waiting for, replication sender threads for the specified context. Values are calculated by the amount of time spent for the activity, multiplying the time by 100, and dividing the time by the duration of the reporting interval.

Due to the presence of multiple threads working on behalf of the specified replication context, the average values are displayed. This average is calculated by adding all the values from all the threads that worked on behalf of the replication context, and dividing this sum with the number of threads that worked on behalf of the context. When a replication throttle is configured, you may see a large amount of output followed by a period of none while viewing performance or statistics. This behavior is the result of how statistics are calculated by replication, combined with the default Data Domain system replication configuration of using large TCP buffers. When a throttle is in effect, data is buffered before being sent on the network. Users may configure the replication bandwidth and delay arguments in `replication option reset` on the source and destination to use smaller TCP socket buffers. This reduces the total amount of data on the network, increases how often replication writes data to its sockets, and, as a result, increases the frequency of updates for statistics counters. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-user, tenant-admin.

Output Definitions

Pre-comp (KB/s)

The size value before compression is applied. Sometimes referred to as *logical size*.

Network (KB/s)

The amount of compressed data transferred over the network per second.

Streams

An internal system resource associated with reads and writes. One replication context can use multiple streams for better performance.

Busy Reading

The time spent reading file system data from the local file system. This number is typically the second highest number after Network. On a deployment with high network bandwidth, Busy Reading may be the largest column.

Busy Meta

The time spent on miscellaneous bookkeeping activities and replicating file system namespace operations. This value is typically under 50. If this value exceeds 50 on a sustained basis, it may indicate an unusual workload (a large number of file attribute updates, for example).

Waiting Dest

The time spent waiting because the receiver is not providing the sender enough information on what data to send. Typically this value is low. Exceptions include systems on high-speed networks where the sender is a more powerful Data Domain system than the replica, or where the replica has a higher workload than the sender because the replica is the destination for multiple replication contexts.

Waiting Network

The time spent sending file data and metadata and waiting for replies from the server on what data needs to be sent. This is typically the highest of the four values. This value exceeds 100 regularly if the sender is able to replicate multiple files in parallel.

Note

If the Network column has the highest time values among Reading, Meta, Waiting, and Network, and if the Network KB/sec value is lower than expected, a network problem may be present. For example, packet loss may be causing reduced throughput.

```
replication show stats [destination | tenant-unit tenant-unit | all]
```

Display statistics for all replication pairs, a tenant unit, or a specific destination pair. Output format is based on replication type. In collection replication, the difference in values between Post-comp Bytes Sent and Post-comp Bytes Received is expected behavior. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-user, tenant-admin.

Output Definitions

Network bytes sent to destination

The number of physical bytes sent to the destination over the wire.

Pre-compressed bytes written to source

The number of bytes received by the source, including logical bytes associated with the file being replicated.

Pre-compressed bytes sent to destination

The number of bytes sent to the destination, including logical bytes associated with the file being replicated.

Pre-compressed bytes remaining

For directory replication only, this is the sum of file sizes remaining to be replicated for the context. Output includes the entire logical size of the current file being replicated. If a large file is being replicated, this number may take a lot of time to change. The number changes only after the current file finishes.

Compression ratio

The ratio of the value of logical bytes ingested to the source to physical bytes actually sent to the destination over the wire.

Sync'ed-as-of time

The time when the most recently replicated data on the destination was generated on the source. A value of Unknown appears during replication initialization.

replication status

```
replication status [destination | tenant-unit tenant-unit | all] [detailed]
```

Show the current status of replication. Role required: admin, limited-admin, security, user, backup-operator, none, tenant-user, tenant-admin.

Example 165

```
# repl status mtree://dd860-79.chaos.local/data/coll/m1
CTX: 1
Mode: source
Destination: mtree://dd860-79.chaos.local/data/coll/m1
Tenant-unit: tul
Enabled: no
Low bandwidth optimization: disabled
Replication encryption: disabled
Replication propagate-retention-lock: enabled
Local filesystem status: enabled
Connection: idle since Mon Jun 16 14:39:32
State: normal
Error: no error
Sync'ed-as-of time: Mon Jun 16 14:39
Current throttle: 983040 bps
Max-repl-streams: 32 (default)
```

replication sync

```
replication sync [and-verify] [destination]
```

Synchronize replication between the source and destination and wait for replication to complete. You must first configure the source and destination and initialize the context. Role required: admin, limited-admin, security, backup-operator.

replication throttle

```
replication throttle add [destination host | default] sched-
specrate
```

Change the rate of network bandwidth used by replication. By default, network bandwidth use is unlimited, meaning it continuously runs as fast as possible. If you set a throttle, replication runs at the given rate until the next scheduled change, or until new throttle command options force a change. Throttle is usually set at the source Data Domain system, but can optionally be set at the destination. Role required: admin, limited-admin.

To limit replication to 5 megabits per second for a destination Data Domain system named ddr1-ny, starting on Tuesdays and Fridays, at 10:00 a.m., enter:

```
# replication throttle add destination ddr1-ny tue fri 2200 5Mbps
```

```
replication throttle del [destination host | default] sched-
spec
```

Remove one or more throttle schedule entries. Role required: admin, limited-admin.

To remove an entry for Mondays at 1:00 p.m., enter:

```
# replication throttle del mon 1300
```

```
replication throttle reset [destination host | default]
{current | override | schedule | all}
```

Reset a throttling schedule. Role required: admin, limited-admin.

```
replication throttle set current [destination host | default]
rate
```

Set the throttle rate until the next scheduled change or a system reboot. Setting the throttle to *current* cannot be done if `replication throttle set override` is in effect. Role required: *admin*, *limited-admin*.

```
replication throttle set override [destination host | default]
rate
```

Set the throttle rate until another *override* is issued. Throttle override cannot be set if `replication throttle set current` is in effect. Role required: *admin*, *limited-admin*.

```
replication throttle show [destination host | default | all]
```

Show throttle configuration. If no option is specified, *all* is the default option. Role required: *admin*, *limited-admin*, *security*, *user*, *backup-operator*, *none*.

```
replication throttle show performance [destination host |
default | all] [interval sec] [count count]
```

Show current throttle throughput for an optionally specified number of times and interval. If no option is specified, *all* is the default option. Role required: *admin*, *limited-admin*, *security*, *user*, *backup-operator*, *none*.

To specify that results be shown exactly 7 times, at 2 second intervals (for a total of 14 seconds), enter:

```
# replication throttle show performance all interval 2 count
710/16 10:15:18
usr1-dl.datadomain
      [8000K bps]
      -----
          (0 bps)
          (0 bps)
          (0 bps)
          (0 bps)
          (0 bps)
          (0 bps)
          (0 bps)
SE@usr1-dd1## date
Wed Oct 16 10:15:34 PDT 2013
```

Argument Definitions

all

Removes and resets current or override settings and removes all scheduled changes. This option returns the system to the default settings.

count

The number of times the results will be shown. The default is unlimited (the command will run until it is ended by the user).

current

Removes and resets the rate set by a previous `replication throttle set current`.

host

The destination hostname when you are setting up a destination throttle.

override

Removes and resets the rate set by a previous `replication throttle set override`.

rate

The rate, which can be the word `unlimited`; or a number; or `disable`, `disabled`, or `zero` (any of the last three will stop replication until the next rate change). If set to zero, new contexts are also throttled to zero. The system enforces a minimum rate of 98,304 bits per second (about 100 Kbps) and a maximum of 34,358,689,792 bits per second (about 34 Gbps). The number can include a tag for bits or bytes per second. Do not use a space between the number and the bits or bytes specification. The default rate is bits per second. In the rate variable:

- `bps` equals raw bits per second
- `Kbps` or `kbps` equals 1000 bits per second
- `Mbps` or `mbps` equals 1×10^6 bits per second
- `Gbps` or `gbps` equals 1×10^9 bits per second

`Kib` = Kibibits, the base-2 equivalent of `Kb` or Kilobits. `KiB` = Kibibytes, the base-2 equivalent of `KB` or Kilobytes.

sched-spec

One or more three-letter days of the week (such as `mon`, `tue`, or `wed`), or the word `daily` (to set the schedule every day). This argument can also specify a time of day in 24-hour format.

schedule

Removes and resets scheduled changes.

sec

The number of seconds for the interval between displaying the results. The default is five seconds.

replication watch

`replication watch destination`

Display the progress of a replication initialization, resynchronization process, or recovery operation. Role required: `admin`, `limited-admin`, `security`, `user`, `backup-operator`, `none`.

Example 166

During initialization:

```
# repl init rctx://14
(00:00) Initialize started.
Use 'replication watch rctx://14' to monitor progress.

# repl watch rctx://14
Use Control-C to stop monitoring.
(00:00) Replication initialize started...
(00:08) 100%: pre-initialize
(00:08) initializing 3/3:
(00:33) : 60% completed, pre-comp: 0 KB/s, network: 6 KB/s
```


Example 166 (continued)

When initialization completes:

```
# repl init rctx://14
(00:00) Initialize started.
Use 'replication watch rctx://14' to monitor progress.

# repl watch rctx://14
Use Control-C to stop monitoring.
(00:00) Replication initialize started...
(00:08) 100%: pre-initialize
(00:08) initializing 3/3:
(00:49) : 100% completed, pre-comp: 0 KB/s, network: 6 KB/s
(00:49) Replication initialize completed.
```

replication

CHAPTER 35

route

The `route` command is an alias for the `net route` command.

The `route` command manages routing between a Data Domain system and backup hosts. An additional routing rule in the Kernel IP routing table and in the Data Domain system Route Config list shows a list of static routes reapplied at each system boot. Each interface is assigned a route based on the address assigned to it. Also, depending on the default gateway subnet, a route is added to an interface automatically if the interface is in the subnet of the default route address.

Federal certification requirements state the DD OS must be IPv6-capable and that interoperability with IPv4 be maintained in a heterogeneous environment. As a result, several `net` command options now include arguments for both versions of Internet Protocol. Data Domain customers select which version to use, based on the type of configuration.

This chapter contains the following topics:

• route change history	340
• route guidelines and restrictions	340
• route add	340
• route del	342
• route reset	342
• route set	343
• route show	343
• route trace	343

route change history

There have been no changes to this command in this release.

route guidelines and restrictions

- Changes to Ethernet interfaces made with `net` command options flush the routing table. All routing information is lost and data movement using routing is cut off immediately. If possible, make interface changes only during scheduled downtime. You must reconfigure routing rules and gateways after making interface changes.

route add

```
route add [ipversion {ipv4 | ipv6}] <route spec>
```

The IPv4 *route spec* syntax is: *ipv4address* [*netmask*] [*gw gateway*]

The IPv6 *route spec* syntax is: *ipv6address* [*gw gateway*]

Add an IPv4 or IPv6 static route for a network or network host. Role required: admin, limited-admin.

```
route add [ipversion {ipv4 | ipv6}] [type {fixed | floating}]
<route spec>
```

Add a fixed or floating static route in a high-availability (HA) system. Role required: admin, limited-admin.

```
route add gateway ipv4address [interface name]
```

Add a gateway address to the list of gateway addresses on the Data Domain system. Optionally specify a specific interface to associate with the gateway address. If the gateway is unreachable, the system displays a warning, but still adds the gateway. Role required: admin, limited-admin.

Argument Definitions

fixed

Specifies that the static route is fixed.

floating

Specifies that the static route is floating.

gw gateway

Specifies the IP address of the gateway to use to reach the destination network or host. If no gateway is specified, the route uses the default gateway.

ipv4address

Specifies the IPv4 address for the destination network or host. If no gateway is specified, the command fails if the destination host is not found on the local network or through the default gateway.

ipv6address

Specifies the IPv6 address for the destination network or host. If no gateway is specified, the command fails if the destination host is not found on the local network or through the default gateway.

ipversion ipv4

Specifies that the new route is for IPv4 routing. If this is omitted, the route is applied to the IPv4 routing table if it is valid.

ipversion ipv6

Specifies that the route is for IPv6 routing. This argument is not required when an IPv6 address is specified.

-netmask

Specifies the network mask that applies to the destination network or network host.

type

Specifies the type of static route is either a fixed or floating IP.

Example 167

This example adds an IPv4 route to network 192.168.1.0 with netmask 255.255.255.0 using the *svr12* gateway.

```
# route add 192.168.1.0 netmask 255.255.255.0 gw svr12
```

Example 168

This example adds an IPv4 route to a host named *user24* through the *srvr12* gateway.

```
# route add user24 gw srvr12
```

route del

```
route del [ipversion {ipv4 | ipv6}] <route spec>
```

The IPv4 *route spec* syntax is: *ipv4address* [*netmask*] *gw gateway*

The IPv6 *route spec* syntax is: *ipv6address gw gateway*

Delete an IPv4 or IPv6 static route for a network or network host. Role required: admin, limited-admin.

```
route del gateway {ipv4address | routing-table-name name
```

Deletes the specified gateway, along with its associated route entries and route rules.

Role required: admin, limited-admin.

Argument Definitions

gw gateway

Specifies the IP address of the gateway used to reach the destination network or host.

ipv4address

Specifies the IPv4 address of the destination network or host.

ipv6address

Specifies the IPv6 address of the destination network or host.

ipversion ipv4

Specifies that the route is an IPv4 route. If this is omitted, the route is deleted from the IPv4 routing table.

ipversion ipv6

Specifies that the route is an IPv6 route. If this is omitted, the route is deleted from the IPv4 routing table.

-netmask

Specifies the network mask that applies to the destination network or network host.

Example 169

To delete a route, for example, 192.168.1.0 and a netmask of 255.255.255.0, enter:

```
# route del 192.168.1.0 netmask 255.255.255.0 gw 10.25.160.
```

route reset

```
route reset gateway [ipversion {ipv4 | ipv6}]
```

Delete the configured routing gateway for the specified protocol. If no protocol is specified, the IPv4 gateway configuration is deleted. Role required: admin, limited-admin.

route set

```
route set gateway {ipaddr | ipv6addr}
```

Configure the IP address for the IPv4 or IPv6 default gateway. When the default gateway is added or changed, the Data Domain operating system automatically adds a route to default gateway for each interface with the same subnet. Role required: admin, limited-admin.

Note

When configuring an IPv6 address, a command failure might not produce an error message in the CLI. If the new gateway is not visible using the `route show gateway` and `route show table` commands, check the `messages` log file for information on why the command failed.

Example 170

The following example configures the device at 192.168.10.1 as the default IPv4 gateway.

```
# route set gateway 192.168.10.1
```

route show

```
route show config [routing-table-name name]
```

Display the configured static routes. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
route show gateways [detailed] [ipversion {ipv4 | ipv6}]
[<ipv4addr>]
```

Displays the configured or DHCP-supplied IPv4 and IPv6 gateways as specified. The `detailed` option displays the network interface or type, associated routing tables, interface addresses, and owners if applicable. If no IP version is specified, both gateways are displayed. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
route show tables [table-name-list | gateway ipv4addr |
ipversion {ipv4 | ipv6}]
```

Displays the IPv4 and IPv6 routing tables as specified. If no IP version is specified, both tables are displayed. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
route show [ipversion {ipv4 | ipv6}] [type {fixed | floating}]
```

Displays the type of IP address as specified.

route trace

```
route trace ipv4addr | ipv6addr | {hostname [ipversion {ipv4 |
ipv6}]}
```

Display a route used by a Data Domain system to connect with the specified destination. Role required: admin, limited-admin, security, user, backup-operator, or none.

Example 171

To trace the route to `svr24`, enter:

```
# route trace svr24
Traceroute to svr24.yourcompany.com (192.168.1.6), 30 hops max, 38
byte packets
1 svr24 (192.168.1.6) 0.163 ms 0.178 ms 0.147 ms
```


CHAPTER 36

scsitarget

The `scsitarget` command manages the SCSI (Small Computer System Interface) target subsystem configuration on single-node DD systems and on systems using Extended Retention.

The SCSI target subsystem configuration comprises several SCSI target entities:

- services (VTL, DD Boost, and vdisk)
- transports (Fibre Channel)
- transport endpoints (Fibre Channel port)
- endpoints (such as VTL tape drives)
- logical devices
- host initiators
- access groups

In some cases, mostly group management, individual services provide interfaces more tailored to the service, for example, `vtl group`. These services may be more convenient for daily use than the generic `scsitarget` interface.

This chapter contains the following topics:

• scsitarget change history	346
• scsitarget device	346
• scsitarget disable	346
• scsitarget enable	346
• scsitarget endpoint	346
• scsitarget group	350
• scsitarget initiator	352
• scsitarget option	353
• scsitarget persistent-reservation	355
• scsitarget port	355
• scsitarget reset	359
• scsitarget service	359
• scsitarget show	359
• scsitarget status	359
• scsitarget trace	360
• scsitarget transport	361

scsitarget change history

There have been no changes to this command in this release.

scsitarget device

```
scsitarget device show detailed [device-spec] [service service-name] [group group-spec]
```

Show detailed information for SCSI target or vdisk devices. Role required: admin, limited-admin, security, user, backup-operator, none.

```
scsitarget device show list [device-spec] [service service-name] [group group-spec]
```

List summary information for SCSI target or vdisk devices. If no arguments are selected, the output includes basic information for all device criteria, including vdisk devices. Role required: admin, limited-admin, security, user, backup-operator, none.

Argument Definitions

device-spec

A list of devices that may use wildcards. This can be a vdisk *device-spec*.

group-name

The name of the SCSI target access group. These names are case-insensitive and case-preserving. They cannot include colons, question marks, commas, asterisks, forward or backward slashes, open or closed parentheses, or the words **summary**, **all**, or **VTL**.

service-name

A SCSI target service: vtl, ddboost, or vdisk.

scsitarget disable

```
scsitarget disable
```

Disable the SCSI target subsystem. Role required: admin, limited-admin.

scsitarget enable

```
scsitarget enable
```

Enable the SCSI target subsystem. Role required: admin, limited-admin.

scsitarget endpoint

```
scsitarget endpoint add endpoint-name system-address address
[primary-system-address address] [secondary-system-address
{address-list | none}] [wwpn {auto | wwpn}] [wwnn {auto |
wwnn}] [fc2-retry {disable | enable | default}]
```

Add a SCSI target endpoint. For the Fibre Channel transport, NPIV must be enabled to have more than one endpoint per system address. Endpoints are added as disabled; they must be explicitly enabled using `scsitarget endpoint enable`. This allows other properties of the endpoint to be changed before enabling the endpoint. Role required: admin, limited-admin.

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, 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.

Example 172

```
# scsitarget endpoint add endpoint-2 system-address 10a wwpn
20:10:33:44:55:66:77:88 wwnn 20:00:33:44:55:66:77:88
Endpoint "endpoint-2" created OK.
```

```
scsitarget endpoint connection-reset endpoint-spec
```

Reset one or more SCSI target endpoints. Be aware that resetting endpoint connections during a backup may disrupt the backup operation. If NPIV is enabled, this command resets the port instance currently associated with the endpoint, if any. Use `scsitarget port connection-reset` to reset all connections for a port. Role required: admin, limited-admin.

```
scsitarget endpoint del endpoint-spec
```

Delete one or more endpoints. This may be used to delete an endpoint if the underlying hardware is no longer available. 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. Role required: admin, limited-admin.

```
scsitarget endpoint disable endpoint-spec
```

Disable one or more SCSI target endpoints. 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. Use `scsitarget port disable` to explicitly disable a port and all endpoints using it. Role required: admin, limited-admin.

```
scsitarget endpoint enable endpoint-spec
```

Enable one or more SCSI target endpoints. Enabling an endpoint enables the port only if it is currently disabled, that is, you are in non-NPIV mode. Use `scsitarget port enable` to explicitly enable a port and all endpoints using it. Role required: admin, limited-admin.

```
scsitarget endpoint modify endpoint-spec [system-address
address] [primary-system-address address] [secondary-system-
address {address-list | none}] [wwpn {auto | wwpn}] [wwnn {auto
| wwnn}]
```

Modify one or more endpoints. Role required: admin, limited-admin.

Note

When using NPIV, it is recommended that you use only one protocol (that is, 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.

Example 173

The following example sets the primary system address endpoint-fc-1 to 5a and the secondary system address to 6a,6b and verifies both addresses. Note that the system-address is changed in this case because the current system address is set to the new primary system address as part of the change.

```
# scsitarget endpoint modify endpoint-fc-1 system-address "5a"
secondary-system-address "6a,6b"
Endpoint 'endpoint-fc-1' successfully modified.

# scsitarget endpoint show detailed endpoint-fc-1
Endpoint:                endpoint-fc-1
Current System Address:  5a
Primary System Address:  5a
Secondary System Address: 6a,6b
Enabled:                 Yes
...
```

Example 174

The following example changes the primary system address of a failed-over endpoint so that on failback it moves to a different port.

```
# scsitarget endpoint modify endpoint-fc-1 primary-system-address "8a"
Endpoint 'endpoint-fc-1' successfully modified.

# scsitarget endpoint show detailed endpoint-fc-1
Endpoint:                endpoint-fc-1
Current System Address:  6a
Primary System Address:  8a
Secondary System Address: 6a,6b
Enabled:                 Yes
...
```

```
scsitarget endpoint rename src-endpoint-name dst-endpoint-name
```

Rename an endpoint. Role required: admin, limited-admin.

```
scsitarget endpoint show detailed [endpoint-spec] [system-address system-address-spec]
```

Show detailed information about one or more endpoints. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 175

```
# scsitarget endpoint show detailed endpoint-fc-1
Endpoint:                endpoint-fc-1
Current System Address:  5a
Primary System Address:  5a
Secondary System Address: 6a,6b
Enabled:                 Yes
Status:                  Online
Transport:               FibreChannel
```

Example 175 (continued)

```
FC WWNN:                25:80:00:21:88:00:61:d3
FC WWPN:                25:00:00:21:88:00:61:d3
```

```
scsitarget endpoint show list [endpoint-spec] [system-address
system-address-spec]
```

Show summarized list of configured endpoints. If no argument is selected, the output will be basic information for all endpoint criteria. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 176

```
# scsitarget endpoint show list
system-address 5a,5b
Endpoint      System Address      Transport      Enabled      Status
-----
endpoint-fc-0 5a                    FibreChannel   Yes          Online
endpoint-fc-3 5a                    FibreChannel   No           Offline
x1            5b                    FibreChannel   Yes          Online
hdp1         5b                    FibreChannel   Yes          Online
-----
```

```
scsitarget endpoint show stats [endpoint-spec] [interval
interval] [count count]
```

Periodically list I/O statistics on one or more endpoints. If no endpoints are specified, the output will be a single-line total for each interval. Role required: admin, limited-admin, security, user, backup-operator, none.

```
scsitarget endpoint use endpoint-spec { primary | secondary }
```

Change the in-use system address for one or more endpoints. Role required: admin, limited-admin.

Example 177

```
# scsitarget endpoint use endpoint-fc-1 primary
```

Argument Definitions**count**

The number of objects on which to perform the action, as specified by the command option.

endpoint-name

The name of the endpoint (which, in a SCSI target architecture, corresponds to a virtual port on a DD system). The name must not conflict with any other endpoint name currently in the system.

endpoint-spec

A list of endpoints (which, in a SCSI target architecture, corresponds to a virtual port on a DD system) that may use wildcards.

fcp2-retry

A port option.

interval *interval*

The time interval in seconds. The default is 2 seconds.

primary-system-address

The primary system address for the endpoint. The primary and any secondary system address must be different. The current system address must always be in the set of primary and secondary addresses.

secondary-system-address

The secondary system address for the endpoint. The primary and any secondary system address must be different. If the endpoint is failed-over, the current system address must remain in the secondary address list. If this is set to none, the endpoint cannot failover. The current system address must always be in the set of primary and secondary addresses. If multiple secondary addresses are given, when failover occurs the system will automatically pick one of the addresses from that list to make the current system address.

system-address

A system-specific name that identifies a specific SCSI target transport interface. For the Fibre Channel transport, the system address is the name of the HBA (host bus adapter) port used, for example, 5a. This name must match a currently valid system address in the system. It must always be in the set of primary and secondary addresses.

topology

The Fibre Channel topology for the endpoint. Values include: `loop-preferred`, `loop-only`, `point-to-point`, `default`.

wwpn

The worldwide port name (WWPN) for the endpoint, which must follow existing rules for WWPN conflict. If you do not provide a `wwpn`, and the transport uses `wwpn`, it is assigned by default.

wwnn

The worldwide node name (WWNN) for the endpoint, which must follow existing rules for WWNN conflict. If you do not provide a `wwnn`, and the transport uses `wwnn`, it is assigned by default.

scsitarget group

```
scsitarget group add group-name device device-spec [lun lun]
[primary-endpoint {all | none | endpoint-list}] [secondary-
endpoint {all | none | endpoint-list}]
```

Add SCSI target or vdisk devices to a group. Role required: admin, limited-admin.

```
scsitarget group add group-name initiator initiator-spec
```

Add one or more initiators to a group. Role required: admin, limited-admin.

```
scsitarget group attach group-name device device-name lun lun
primary-endpoint {all | none | endpoint-list} secondary-
endpoint {all | none | endpoint-list}
```

Attach an additional LUN to a SCSI target or vdisk device in a group. This may be used to expose a device with different LUNs through different endpoints. Role required: admin, limited-admin.

```
scsitarget group create group-name service service-name
```

Create a new group associated with a specific service, which can be a vdisk service. Role required: admin, limited-admin.

```
scsitarget group del group-name device device-spec
```

Delete one or more SCSI target or vdisk devices from a group. Role required: admin, limited-admin.

```
scsitarget group del group-name initiator initiator-spec
```

Delete one or more initiators from a group. Role required: admin, limited-admin.

```
scsitarget group destroy group-name
```

Destroy a group. Role required: admin, limited-admin.

```
scsitarget group detach group-name device device-name lun lun
```

Detach a SCSI target or vdisk device from a LUN in a group. There must be at least one LUN for a device in a group. Role required: admin, limited-admin.

```
scsitarget group modify group-name device device-spec [lun lun]
[primary-endpoint {all | none | endpoint-list}] [secondary-
endpoint {all | none | endpoint-list}]
```

Modify SCSI target or vdisk device attributes in a group. If a device is attached to multiple LUNs, the *lun* argument, if specified, indicates which LUN to update. Role required: admin, limited-admin.

```
scsitarget group rename src-group-name dst-group-name
```

Rename a group. Role required: admin, limited-admin.

```
scsitarget group show detailed [group-spec] [device device-
spec] [initiator initiator-spec] [service service-name]
```

Show detailed information on specific groups, based on selected arguments. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 178

```
#scsitarget group show detailed vdisk_g1
```

```
Group: vdisk_g1
Service: VDISK
Active state: active
Initiators: None
Devices: None
```

```
scsitarget group show list [group-spec] [device device-spec]
[initiator initiator-spec] [service service-name]
```

Display a list of groups based on selected arguments. If no arguments are selected, output displays basic information on all group criteria, including vdisk devices. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 179

```
# scsitarget group show list
```

Group Name	Service	# Initiators	# Devices
TapeServer	VTL	0	0
disk1	VDISK	0	0
test1	VTL	0	0
vdisk_g1	VDISK	0	0
vdisk_g2	VDISK	0	0

```
scsitarget group use group-name device device-spec {primary |
secondary}
```

Switch the in-use endpoint lists for one or more SCSI target or vdisk devices in a group between primary and secondary endpoint lists. For best results, do not run this command option during heavy VTL usage. Role required: admin, limited-admin.

Argument Definitions

all

Shows all information about the object specified by the command option.

device-name

The name of the SCSI target or vdisk device. These names are case-insensitive and case-preserving. They cannot include colons, question marks, commas, asterisks, forward or backward slashes, open or closed parentheses, or the word **all**.

device-spec

A list of devices that may use wildcards. This can be a vdisk *device-spec*.

group-name

The name of the SCSI target access group. These names are case-insensitive and case-preserving. They cannot include colons, question marks, commas, asterisks, forward or backward slashes, open or closed parentheses, or the words **summary**, **all**, or **VTL**.

group-spec

A list of access groups that may use wildcards. This can be a vdisk *group-spec*.

initiator-spec

A list of initiators that may use wildcards.

lun

A device address to pass to the initiator. The maximum logical unit number (LUN) is 16383. A LUN must be unique within a group, but need not be unique across the system. LUNs for VTL devices within a group must start with zero and be contiguous numbers.

primary-endpoint

The primary endpoint on which the SCSI target devices are visible. By default, or if you specify **all**, SCSI target devices are visible on all ports. Specify **none** if the devices should not be visible on any ports.

secondary-endpoint

The secondary endpoint on which the SCSI target devices are visible. By default, the devices are visible on all ports. The secondary port list supports path redundancy.

service-name

A SCSI target service: **vtl**, **ddboost**, or **vdisk**.

scsitarget initiator

```
scsitarget initiator add initiator-name system-address system-address
```

Add an initiator with the specified system address. An initiator may be added before it is visible on a port, which allows for early provisioning. Role required: admin, limited-admin.

```
scsitarget initiator del initiator-spec
```

Delete an initiator. Note that if the initiator remains visible, it may be automatically rediscovered. Role required: admin, limited-admin.


```
scsitarget initiator modify initiator-spec [address-method
{auto | vsa | default}]
```

Modify one or more initiators. Role required: admin, limited-admin.

```
scsitarget initiator rename src-initiator-name dst-initiator-
name
```

Rename an initiator. Role required: admin, limited-admin.

```
scsitarget initiator show detailed [initiator-spec] [endpoint
endpoint-spec] [group group-spec]
```

Show detailed information for one or more initiators, based on selected arguments. Role required: admin, limited-admin, security, user, backup-operator.

```
scsitarget initiator show list [initiator-spec] [endpoint
endpoint-spec] [group group-spec]
```

Display a list of initiators based on selected arguments. If no arguments are selected, the output consists of basic information for all initiator criteria. Role required: admin, limited-admin, security, user, backup-operator.

Argument Definitions

auto

The device address method chosen based on the numeric LUN range being reported: 0 - 255, peripheral device addressing is used, 256 - 16383, flat device addressing is used (default).

endpoint-spec

A list of endpoints (which, in a SCSI target architecture, corresponds to a virtual port on a DD system) that may use wildcards.

group-spec

A list of access groups that may use wildcards. This can be a vdisk *group-spec*.

initiator-name

The name of SCSI target host initiator.

initiator-spec

A list of initiators that may use wildcards.

system-address

A system-specific name that identifies a specific SCSI target transport interface. For the Fibre Channel transport, the system address is the name of the HBA (host bus adapter) port used, for example, 5a. This name must match a currently valid system address in the system. It must always be in the set of primary and secondary addresses.

vsa

Volume set addressing (VSA). This method is used primarily for addressing virtual buses, targets, and LUNs. The HP-UX operating system selects the volume set addressing method based on inquiry data and LUN information returned by the SCSI-3 REPORT LUNS command.

scsitarget option

```
scsitarget option reset {option-name | all}
```

Reset SCSI target global options. Role required: admin, limited-admin.

```
scsitarget option set [failover-delay delay-secs] [failback-
delay delay-secs] [automatic-failback {enabled | disabled}]
Set SCSI target global options. Role required: admin, limited-admin.
```

Note

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.

Note

Here is 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.
 - VTL Fibre Channel operation is expected to be interrupted when the VTL Fibre Channel endpoints failover. You may need to perform discovery (that is, operating system discovery and configuration of 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.
-

Example 180

To set a delay of 60 seconds before starting failover and to perform automatic failback when a port has been normal for 300 seconds:

```
# scsitarget option set failover-delay 60 automatic-failback enabled
failback-delay 300
```

```
scsitarget option show {option-name | all}
```

Show SCSI target global options. Role required: admin, limited-admin, security, user, backup-operator, none.

Argument Definitions

automatic-failback

A SCSI target global option, which provides the option to automatically failback or not. Values are enabled or disabled (default).

failback-delay

A SCSI target global option, which is the time to wait before attempting automatic failback when the interface is normal. The default is 120 seconds, the minimum is 30 seconds, and the maximum is 600 seconds.

failover-delay

A SCSI target global option, which is the time to delay before performing a failover. The default is 90 seconds, the minimum is 10 seconds, and the maximum is 300 seconds.

option-name

List of SCSI target global options. One or more may be specified.

```
automatic-failback
```

```
failback-delay
```

```
failover-delay
```

scsitarget persistent-reservation

```
scsitarget persistent-reservation clear [device device-spec]
[initiator initiator-name]
```

Clear SCSI persistent reservations. Role required: admin, limited-admin.

Example 181

To clear all persistent reservations set by an initiator no longer visible to the system, enter:

```
# scsitarget persistent-reservation clear initiator ibm-initiator-17
```

```
scsitarget persistent-reservation disable [service service-name]
```

Disable SCSI persistent reservations. Role required: admin, limited-admin.

```
scsitarget persistent-reservation enable [service service-name]
```

Enable SCSI persistent reservations. Role required: admin, limited-admin.

```
scsitarget persistent-reservation show detailed [device device-spec]
[initiator initiator-name]
```

Show detailed information for SCSI persistent reservations. Be aware that if a device does not include a reservation key, or is using a shared key, a series of zeros (0X0000000000000000) will be displayed in the Reservation Key category, instead of n/a, which is the expected behavior. Role required: admin, limited-admin, security, user, backup-operator, none.

```
scsitarget persistent-reservation show list [device device-spec]
[initiator initiator-name]
```

Show summary information for SCSI persistent reservations. Role required: admin, limited-admin, security, user, backup-operator, none.

Argument Definitions

device-spec

A list of devices that may use wildcards. This can be a vdisk *device-spec*.

initiator-name

The name of SCSI target host initiator.

service-name

A SCSI target service: vtl, ddboost, or vdisk.

scsitarget port

```
scsitarget port connection-reset system-address-spec
```

Reset all connections for the given SCSI target port. Role required: admin, limited-admin.

```
scsitarget port disable system-address-spec [failover-endpoints]
```

Disable one or more SCSI target ports, along with any endpoints currently using that port. If *failover-endpoints* is used, any endpoints that use the port for their primary system address will be disabled or failed-over. Endpoints that are already disabled by administrative operation prior to a port being disabled are remembered as manually disabled. This state will be restored when that port is later enabled. Role required: admin, limited-admin.

Example 182

```
# scsitarget port disable 5a failover-endpoints
```

```
scsitarget port enable system-address-spec [failback-endpoints]
```

Enable one or more SCSI target ports, along with any endpoints currently using that port. If *failback-endpoints* is used, any endpoints that use the port for their primary system address, and are failed-over, will be failed-back. Role required: admin, limited-admin.

Example 183

```
# scsitarget port enable 5a failback-endpoints
System address '5a' successfully enabled.
```

```
scsitarget port modify system-address-spec [topology {loop-
preferred | loop-only | point-to-point | default}] [speed {auto
| speed}] [base-wwpn {auto | wwpn}] [base-wwnn {auto | wwnn}]
[npiv {auto | disabled}] [fc2-retry {disable | enable |
default}]
```

Modify options for SCSI target ports. Role required: admin, limited-admin.

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.

Example 184

```
# scsitarget port modify 6b topology point-to-point speed 16
```

```
scsitarget port reset detailed-stats
```

Reset cumulative information about SCSI target ports. Note that these statistics are gathered by autosupport. As with other similar always-incrementing statistics, autosupport itself may periodically reset the statistics. Role required: admin, limited-admin.

```
scsitarget port show detailed [system-address-spec] [transport
transport-name]
```

Show configured SCSI target ports in detailed form. Role required: admin, limited-admin.

Example 185

```
# scsitarget port show detailed 5a
System Address:      5a
```

Example 185 (continued)

```

Enabled:                Yes
Status:                 Online
Transport:              FibreChannel
FC Port:                5a
FC Operational Status: Normal
FC NPIV:                Enabled (auto)
Port ID:                0x0b0000
Model:                  QLE2562
Firmware:               5.08.00
FC WWNN:                25:80:00:21:88:00:61:d3
FC WWPN:                25:00:00:21:88:00:61:d3
Connection Type:       N-Port
Link Speed:             8 Gbps (auto)
FC Topology:            Default

```

```

Endpoints for port 5a:
Endpoint      Enabled Status  Current Instance
-----
endpoint-fc-1  Yes  Online      5a:1
endpoint-fc-2  Yes  Online      5a:3
endpoint-fc-7  Yes  Online      5a:2
endpoint-fc-10 No   Offline     n/a
-----

```

```
scsitarget port show detailed-stats
```

Show detailed cumulative information about SCSI target ports. Note that these statistics are gathered by autosupport. As with other similar always-incrementing statistics, autosupport may periodically reset the statistics. Role required: admin, limited-admin.

Example 186

```
# scsitarget port show detailed-stats
```

```
Statistics for Transport: FibreChannel
```

System Address	Control Commands	Write Commands	Read Commands	In (MiB)	Out (MiB)
5a	7	32	0	1	0
5b	22	0	0	0	0

System Address	Link Failures	LIP Count	Sync Losses	Signal Losses	Prim Seq Proto Errors	Invalid Tx Words	Invalid CRCs
5a	0	1	0	0	0	0	0
5b	0	2	0	0	0	0	0

```
scsitarget port show list [system-address-spec] [transport transport-name]
```

List configured SCSI target ports in summary form. Role required: admin, limited-admin.

Example 187

To list addresses starting with "5":

```
# scsitarget port show list 5*
```

System Address	Transport	Enabled	Online Status	Operation Status	#Endpoints
5a	FibreChannel	Yes	Online	Normal	4

Example 187 (continued)

```
5b          FibreChannel  Yes  Online  Marginal  1
-----
```

```
scsitarget port show stats [system-address-spec] [interval
interval] [count count]
```

Show I/O-oriented statistics related to one or more SCSI target port. This command is functionally and syntactically equivalent to `scsitarget endpoint show stats`, except that it shows values by port instead of by endpoint. Role required: admin, limited-admin.

Example 188

To show I/O statistics every 30 seconds, two times, for system address 5a:

```
# scsitarget port show stats 5a interval 30 count 2
08/13 15:52:12
System Address  Ctrl/s  Read/s  Read MiB/s  Write/s  Write MiB/s
-----
5a              0       10      5           0         0
-----

08/13 15:52:42
System Address  Ctrl/s  Read/s  Read MiB/s  Write/s  Write MiB/s
-----
5a              0       10      5           0         0
-----
```

Argument Definitions**base-wwnn**

The base worldwide node name (WWNN) for the port. This is used only when NPIV is enabled for the port. Values are auto (default) or a valid WWNN.

base-wwpn

The base worldwide port name (WWPN) for the port. This is used only when NPIV is enabled for the port. Values are auto (default) or a valid WWPN.

fcp2-retry

A port option.

npiv

Enables NPIV (N_Port ID Virtualization) support for the port. Values are auto (default) or disabled. If set to disabled, NPIV is always disabled for the port. NPIV is a Fibre Channel feature in which multiple Fibre Channel node port (N_Port) IDs can share a single physical N_Port.

speed

The preferred link speed for the port. Values can be auto (default) or a number in Gb/s (1, 2, 4, 8, or 16).

system-address-spec

A list of port system addresses that may use wildcards.

topology

The Fibre Channel topology for the endpoint. Values include: `loop-preferred`, `loop-only`, `point-to-point`, `default`.

scsitarget reset

```
scsitarget reset detailed-stats
```

Reset detailed statistics for a SCSI target subsystem. Role required: admin, limited-admin, security, user, backup-operator, none.

scsitarget service

```
scsitarget service refresh [service]
```

Refresh SCSI target service configuration. All services, including vdisk, within the SCSI target system configuration will be re-created. Role required: admin, limited-admin.

```
scsitarget service show list
```

Display a list of configured services, including vdisk, and current state. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 189

```
# scsitarget service show list
SCSI Target Services
Service      Status
-----
VTL          Running
DD-Boost FC  Shutdown/Inactive
VDISK       Running
-----
```

scsitarget show

```
scsitarget show config
```

Show SCSI target configuration. Role required: admin, limited-admin, security, user, backup-operator, none.

```
scsitarget show detailed-stats
```

Show detailed statistics for the SCSI target subsystem. Role required: admin, limited-admin, security, user, backup-operator, none.

scsitarget status

```
scsitarget status
```

Show SCSI target status.

- The `administrative` state shows the overall state of the SCSI target subsystem.
- The `process` state shows if the SCSI target management process is currently running.
- The `module` state shows if required system modules have been loaded prior to starting the management process.

If the status shows an administrative state of `enabled` but a process state of `stopped`, you can use `scsitarget enable` to request a start of the SCSI target

subsystem. Role required: admin, limited-admin, security, user, backup-operator, none.

scsitarget trace

```
scsitarget trace disable [component {all | user | kernel |
default | component-list}]
```

Disable SCSI target tracing. Role required: admin, limited-admin.

```
scsitarget trace enable [component {all | user | kernel |
default | component-list}] [level {all | high | medium | low}]
[timeout {never | timeout-mins}] [service service-name]
```

Enable SCSI target tracing. If no components are specified, the default components are used. If no timeout is given, a 10-minute timeout is used. Use `scsitarget trace show` to see which components are available for each type (all, default, user, kernel). Role required: admin, limited-admin.

```
scsitarget trace show [component {all | user | kernel | default
| component-list}]
```

Show SCSI target trace status, which includes vdisk service. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 190

```
# scsitarget trace show
Component Name  Level      Timeout (min)  Service
-----
service         medium      9 mins        VDISK
```

Argument Definitions

all

Shows all information about the object specified by the command option.

component

Components available for tracing: all, default, user, kernel.

component-list

List of tracing components. One or more may be specified.

service

device

group

transport

initiator

endpoint

failover

op

system

event


```

comm
monitor
persistent-reservation
session
port

```

level

The degree of debugging verbosity to enable (`all` | `high` | `medium` | `low` | `none`).

service-name

A SCSI target service: `vtl`, `ddboost`, or `vdisk`.

timeout

The length of time that debugging is enabled for the specified components.

scsitarget transport

```
scsitarget transport option reset {option-name | all}
```

Reset a SCSI target transport option (`loop-id`, `npiv`, or `wwnn-scope`). Role required: `admin`, `limited-admin`.

```
scsitarget transport option set option-name value
```

Set a SCSI target transport option (`loop-id`, `npiv`, or `wwnn-scope`). Role required: `admin`, `limited-admin`.

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 (up to a maximum of 8 endpoints per 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.
- Multiple DD systems can be consolidated into a single DD system, however, the number of HBAs remains the same on the single DD system.
- Multiple endpoints can be configured on the single DD system, providing equivalent access to the DD systems that were previously consolidated.

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 0xffffffff. Check the Online Support website for details on updating the system ID.

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.
-

Example 191 Some ports have loop-preferred topology

```
# scsitarget port show detailed
System Address:      6a
Enabled:             No
...
FC Topology:        Loop-Only

System Address:      6b
...
FC Topology:        Loop-Preferred

# scsitarget transport option set npiv enabled
Transport option npiv set for transport fc.

scsitgtd.info log
07/15 11:17:57.234 (tid 0x7f46183b59d0): NPIV is allowed but NPIV for
port 6a must be disabled or reconfigured to a different topology to
function properly.
```

Example 192 No ports have loop-preferred topology

```
scsitarget port show detailed
System Address:      6a
Enabled:             No
...
FC Topology:        Loop-Only

System Address:      6b
...
FC Topology:        Loop-Only

# scsitarget transport option set npiv enabled
Failed to set option npiv for transport fc.

**** NPIV is not allowed. Topology for some or all ports must be
reconfigured to loop preferred.
```

Example 192 No ports have loop-preferred topology (continued)

If port topology is not configured to loop preferred then npiv must be disabled on the port to function properly.

```
scsitarget transport option show {option-name | all}
```

Show SCSI target transport options and currently assigned values. Role required: admin, limited-admin, security, user, backup-operator, none.

Example 193

```
# scsitarget transport option show all
SCSI Target Transport Options
Option          Value
-----
loop-id         1
wwnn-scope     global
npiv            enabled
-----
```

```
scsitarget transport show stats
```

This command is deprecated. Use `scsitarget port show stats`, `scsitarget port show detailed-stats`, and/or `scsitarget port reset detailed-stats` instead.

Argument Definitions**option-name**

The specific SCSI target transport option, which can be `loop-id`, `npiv`, or `wwnn-scope`.

value

The value for the specific option.

scsitarget

CHAPTER 37

smt

The `smt` command manages the *Data Domain Secure Multitenancy* software option, available on DD OS versions 5.5 and later. See the *Data Domain Operating System Administration Guide* for instructions on how to create and administer multiple Tenant Units on a single DD system.

This chapter contains the following topics:

- [smt change history](#)366
- [smt disable](#) 366
- [smt enable](#) 366
- [smt status](#) 366
- [smt tenant](#) 366
- [smt tenant-unit](#) 368

smt change history

There have been no changes to this command in this release.

smt disable

```
smt disable
```

Disable the SMT (Secure Multitenancy) feature. Prior to running this command, you must unassign Tenant Unit resources and destroy Tenant Units. When an MTree is unassigned from a Tenant Unit, the MTree remains on the DD system, and functionality is unaffected. Role required: admin, limited-admin.

smt enable

```
smt enable
```

Enable the SMT (Secure Multitenancy) software option. Required role: admin, limited-admin.

smt status

```
smt status
```

View the status of the SMT (Secure Multitenancy) software option – either enabled or disabled. Required role: admin, limited-admin.

smt tenant

```
smt tenant add tenant-name tenant-units tenant-unit-list
```

Add Tenant Units to a Tenant. Role required: admin, limited-admin.

Example 194

```
# smt tenant add tenant1 tenant-units tu1,tu2
Tenant-unit "tu1" added to tenant "tenant1".
Tenant-unit "tu2" added to tenant "tenant1".
```

```
smt tenant create tenant-name [tenant-uuid uuid]
```

Create a Tenant. When setting up SMT-aware MTree replication, on the destination machine, a Tenant has to be created with the same UUID as the UUID of the Tenant at the source DD system that is associated with the MTree being replicated. MTree replication protocol does an SMT security check during replication initialization, to check that the Tenant UUID at the source and destination are not different. Role required: admin, limited-admin.

Example 195

```
# smt tenant create tenant1 tenant-uuid
659b71dada2f0025:1cf71c412c48fe77
Tenant "tenant1" created with tenant-uuid
"659b71dada2f0025:1cf71c412c48fe77".
```

```
smt tenant del tenant-name tenant-units tenant-unit-list
```

Delete Tenant Units from a Tenant. Role required: admin, limited-admin.

Example 196

```
# smt tenant del tenant1 tenant-units tu1,tu2
Tenant-unit "tu1" deleted from tenant "tenant1".
Tenant-unit "tu2" deleted from tenant "tenant1".
```

```
smt tenant destroy tenant-name
```

Destroy a Tenant. Role required: admin, limited-admin.

Example 197

```
# smt tenant destroy tenant1
Tenant "tenant1" destroyed.
```

```
smt tenant rename tenant-name new-tenant-name
```

Rename a Tenant. Role required: admin, limited-admin.

Example 198

```
# smt tenant rename tenant1 tenant2
Tenant "tenant1" renamed to "tenant2".
```

```
smt tenant show detailed [all | tenant tenant-name]
```

Show detailed information about all Tenants or a specific Tenant. Role required: admin, limited-admin, security, user, backup-operator.

Example 199

```
# smt tenant show detailed all
Tenant: t1
Tenant UUID: f15c920502a3a7f8:e3ca68e199aac091
Tenant Pre-Comp (GiB): 360
Tenant-units:
  Name      Tenant      Number of      Types      Pre-Comp
  -----      Self-service      Mtrees      -----      (GiB)
  tu1      Disabled      1      DD Boost      120.0
  tu2      Enabled      2      CIFS, DD Boost      240.0
  tu3      Enabled      2      240.0
  tu4      Enabled      2      240.0
  -----
Management-IP:
IP Address      Type      Tenant-unit
-----
10.25.246.70      local      tu1
10.25.246.80      remote      tu2
2001:0db8:85a3:0000:0000:8a2e:0370:7334      remote      tu3
-----
Management-User:
  Name      Role      Tenant-unit
  -----
  u1      tenant-admin      tu1
      tenant-admin      tu2
  u2      tenant-user      tu4
      tenant-user      tu3
  -----
Management-Group:
  Name      Role      Tenant-unit
```

Example 199 (continued)

```

-----
g1  tenant-admin      tu1
    tenant-admin      tu2
    tenant-user        tu4
g2  tenant-user        tu3
-----

Tenant:                t2
Tenant UUID:           f15c920502a3a7f8:e3ca68e199aac092
Tenant Pre-Comp (GiB): 0
Tenant-units:
  No tenant-units.
Management-IP:
  No management-ips.
Management-User:
  No management-users.
Management-Group:
  No management-groups.

```

```
smt tenant show list [all | tenant tenant-name]
```

Show a summary for all known Tenants or a specific Tenant. Role required: admin, limited-admin, security, user, backup-operator.

Example 200

```

# smt tenant show list all
Name          Number of      Pre-comp      UUID
              Tenant-units  (GiB)
-----
tenant1        2              120.0        f15c920502a3a7f8:e3ca68e199aac091
tenant2        2              120.0        f15c920502a3a7f8:e3ca68e199aac091
tenant3        2              120.0        f15c920502a3a7f8:e3ca68e199aac091
-----

Total Pre-Comp (GiB): 360.0      <-- when filesystem is up and running
Total Pre-Comp (GiB):
The filesystem is not running.   <-- when filesystem is not running

```

Argument Definitions**new-tenant-name**

This must be a unique name.

tenant-name

This must be a unique name.

tenant-unit-list

A comma-separated list of Tenant Units.

tenant-uuid

This must be a unique identifier.

smt tenant-unit

```
smt tenant-unit create tenant-unit
```

Create a Tenant Unit. Tenant Units are initially created using the SMT (Secure Multitenancy) configuration wizard. Role required: admin, limited-admin.


```
smt tenant-unit data-ip add local ipv4-ipv6-list [remote
ipaddr-list] tenant-unit tenant-unit
```

Add a local or remote data IP address to a Tenant Unit. To add a local IP address, that IP address must already be configured on an existing interface on the Data Domain system. A local IP address must be associated with the Tenant Unit before a remote association is permitted. Role required: admin, limited-admin.

```
smt tenant-unit data-ip del local ipv4-ipv6-list [remote
ipaddr-list] tenant-unit tenant-unit
```

Delete a local or remote data IP address from a Tenant Unit. Role required: admin, limited-admin.

```
smt tenant-unit data-ip show [tenant-unit | all]
```

Show the local and remote data IP addresses for a Tenant Unit. Role required: admin, limited-admin, user, security, backup-operator.

```
smt tenant-unit destroy tenant-unit
```

Destroy a Tenant Unit. Tenant Units must be destroyed before the SMT software option can be disabled. Role required: admin, limited-admin.

```
smt tenant-unit gateway add ipv4-address tenant-unit tenant-
unit
```

Add a gateway to a Tenant Unit. A Tenant Unit must have a local IP address on the same subnet as the gateway before the gateway can be added. The command will fail if the system does not detect a local IP address on the same subnet as a the gateway. Role required: admin, limited-admin.

```
smt tenant-unit gateway del ipv4-address tenant-unit tenant-
unit
```

Delete a gateway from a Tenant Unit. Role required: admin, limited-admin.

```
smt tenant-unit gateway show [tenant-unit | all]
```

Show the gateway for a Tenant Unit. Role required: admin, limited-admin, user, security, backup-operator.

```
smt tenant-unit hostname reset tenant-unit tenant-unit
```

Delete a hostname from a Tenant Unit. Role required: admin, limited-admin.

```
smt tenant-unit hostname set hostname tenant-unit tenant-unit
```

Add a hostname to a Tenant Unit. Role required: admin, limited-admin.

```
smt tenant-unit hostname show [tenant-unit | all]
```

Show the hostname for a Tenant Unit. Role required: admin, limited-admin, user, security, backup-operator.

```
smt tenant-unit management-group assign group group-type
{active-directory | nis} tenant-unit tenant-unit [role {tenant-
admin | tenant-user}]
```

Assign an active directory or NIS management group to a Tenant Unit in the role of tenant-admin or tenant-user. Role required: admin, limited-admin.

Example 201

```
# smt tenant-unit management-group assign group1 group-type nis
tenant-unit tul role tenant-admin
Management group "group1" with type "nis" is assigned to tenant-unit
"tul" as "tenant-admin".
```

```
smt tenant-unit management-group show [tenant-unit | all]
```

Show the management group assigned to one or all Tenant Units. Role required: admin, limited-admin, tenant-admin, tenant-user.

```
smt tenant-unit management-group unassign group group-type
{active-directory | nis} tenant-unit tenant-unit
```

Unassign an active directory or NIS (Network Information System) management group from a Tenant Unit. Role required: admin, limited-admin.

Example 202

```
# smt tenant-unit management-group unassign group1 group-type nis
tenant-unit tu1
Management group "group1" with type "nis" is unassigned from tenant-
unit "tu1".
```

```
smt tenant-unit management-ip add {ipaddr | ip6addr} tenant-
unit tenant-unit type {local | remote}
```

Add a management IP address to a Tenant Unit. Use management IP addresses when you want to restrict self-service access to specific IP address(es). *Local* IP is the IP at the DD system, which is associated with a Tenant Unit for providing self-service. *Remote* IP is a remote IP address (of the client), from which a self-service user can log in to the DD system. Role required: admin, limited-admin.

Example 203

```
# smt tenant-unit management-ip add 10.25.246.190 tenant-unit tu1
type remote
Remote/Local management access ip "10.110.250.31" deleted from tenant-
unit "tu1".
```

```
smt tenant-unit management-ip del {ipaddr | ip6addr} tenant-
unit tenant-unit
```

Delete a management IP address from a Tenant Unit. Role required: admin, limited-admin.

Example 204

```
# smt tenant-unit management-ip del 10.110.250.31 tenant-unit tu1
Remote management access ip "10.25.246.190" added to tenant-unit
"tu1".
```

```
smt tenant-unit management-ip show [tenant-unit | all]
```

Show management IP address information for all Tenant Units. Role required: admin, limited-admin, tenant-admin, tenant-user.

Example 205

```
# smt tenant-unit management-ip show
Tenant-unit: tu1
```

IP Address	Type
10.25.246.190	remote
10.110.250.31	local
2001:0db8:85a3:0000:0000:8a2e:0370:7334	remote

```
Tenant-unit: tu2
No management-ips.
```

```
smt tenant-unit management-user assign user tenant-unit tenant-
unit [role {tenant-admin | tenant-user}]
```

Assign a user from a management group to a Tenant Unit in the role of tenant-admin or tenant-user. Role required: admin, limited-admin.

```
smt tenant-unit management-user show [tenant-unit | all]
```

Show user access information for a specific Tenant Unit or for all Tenant Units. Role required: admin, limited-admin, security, user, backup-operator, tenant-admin, tenant-user.

```
smt tenant-unit management-user unassign user tenant-unit
tenant-unit
```

Unassign a management group user from a Tenant Unit. Role required: admin, limited-admin.

```
smt tenant-unit option reset tenant-unit {all | self-service |
security-mode}
```

Reset options for the specified Tenant Unit. Role required: admin, limited-admin.

Example 206

```
#smt tenant-unit option reset tu1 security-mode
Security-mode reset to default for tenant-unit "tu1".
```

```
smt tenant-unit option set tenant-unit {self-service {enabled |
disabled} | security-mode {default | strict}}
```

Set options for the specified Tenant Unit. Security mode sets a check for SMT-aware replication. The *default* mode makes sure that the source and destination do not belong to *different* Tenants. The *strict* mode makes sure they belong to the *same* Tenant. In the latter case, the UUID for both the source and destination Tenant must have been set and must be identical. Role required: admin, limited-admin.

Example 207

```
#smt tenant-unit option set tu1 security-mode strict
Security-mode set to "strict" for tenant-unit "tu1".
```

```
smt tenant-unit option show {tenant-unit | all}
```

Show options for a specified Tenant Unit or for all Tenant Units. Role required: admin, limited-admin.

Example 208

```
#smt tenant-unit option show all
Tenant-unit   Self-service   Security-mode
-----
tu1           Disabled      strict
tu2           Enabled       default
-----
```

```
smt tenant-unit rename tenant-unit new-name
```

Rename a Tenant Unit. Role required: admin, limited-admin.

```
smt tenant-unit setup tenant-unit
```

Lets you enter Tenant Unit values as prompted by the SMT configuration wizard. Role required: admin, limited-admin.

```
smt tenant-unit show detailed [tenant-unit | all]
```

Show detailed information for specific Tenant Units or for all Tenant Units. Role required: admin, limited-admin, tenant-admin, tenant-user.

Example 209

```
# smt tenant-unit show detailed tu1
Tenant-unit: "tu_test_1"
  Summary:
  Tenant-unit  Tenant  Pre-Comp  Number of  Type(s)  Self-service
  Security-mode Hostname      (GiB)
  Mtrees
-----
default      tu_test_1      0.0      0          Disabled
             host1.datadomain.com
-----
  Data-IP:
  Local      Remote
  Address    Address
  -----
  10.20.30.22  172.16.204.122
  -----
  Management-IP:
  No management-ips.
  Gateway:
  Tenant-unit  Gateway IP
  -----
  tu_test_1    10.20.30.1
  -----
  Management-User:
  No management-users.
  Management-Group:
  No management-groups.
  DDBoost:
  Storage-units not found
  Getting users with default-tenant-unit (null)
  DD Boost users not found
  Mtrees:
  **** There are no MTrees configured.
  Quota:
  **** There are no MTrees configured that match the pattern .
  Replication:
  **** License for "REPLICATION" does not exist.
  Alerts:
  No such active alerts.
```

```
smt tenant-unit show list [tenant-unit | all]
```

Show a list of all Tenant Units or for a specific Tenant Unit. Role required: admin, limited-admin, tenant-admin, tenant-user.

Example 210

```
# smt tenant-unit show list
Tenant-unit  Tenant  Pre-Comp  Number of  Type(s)  Self-service  Security-mode
Hostname
(GiB)
```

Example 210 (continued)

Mtrees						
-----	-----	-----	-----	-----	-----	-----
tu_test_1		0.0	0		Disabled	default
host1.datadomain.com						
-----	-----	-----	-----	-----	-----	-----

smt

CHAPTER 38

snapshot

The `snapshot` command manages MTree snapshots. MTrees add granularity to filesystem-type operations, allowing operations to be performed on a specific MTree instead of the entire filesystem. Snapshots are useful for avoiding version skew when backing up volatile data sets, such as tables in a busy database, and for restoring previous versions of a deleted directory or file.

A snapshot is a read-only copy of the Data Domain MTree from the top of each MTree: `/data/col1/mtree-name`. The MTree `/data/col1/backup` is the default directory created in the system during installation. It is also the MTree that is refreshed during an upgrade procedure. The directory `/backup` points to the default MTree. Snapshots can be accessed from the directories `/backup/.snapshot` or `/data/col1/mtree-name/.snapshot`.

This chapter contains the following topics:

- [snapshot change history](#) 376
- [snapshot create](#) 376
- [snapshot expire](#) 376
- [snapshot list](#) 377
- [snapshot rename](#) 377
- [snapshot schedule](#) 378

snapshot change history

There have been no changes to this command in this release.

snapshot create

```
snapshot create snapshot mtree mtree-path [retention {date | period}]
```

Create a snapshot. Naming conventions for creating MTrees include uppercase and lowercase letters A-Z, a-z), numbers 0-9, single, non-leading embedded space, exclamation point (!), hash (#), dollar sign (\$), ampersand (&), caret (^), tilde (~), left and right parentheses ((or)), left and right brackets ([or]), left and right curly braces ({ or }). Role required: admin, limited-admin, backup.

Argument Definitions

snapshot

A name for the snapshot.

mtree mtree-path

The pathname of the MTree for which the snapshot is being created. The base of the path must be `/data/coll/mtree_name` or `/backup`.

retention *date*

A four-digit year, two-digit month, and two-digit day separated by dots, slashes, or hyphens. For example, 2013.03.23. The snapshot is retained until midnight (00:00, the first minute of the day) of the designated date.

retention *period*

Number of days, weeks (*wks*), or months (*mos*) to retain a snapshot. Note there is no space between the number and time period; for example, 4wks. Also, one month equals 30 days. The snapshot is retained until the same time of day it was created.

Example 211

If a snapshot was created at 8:48 a.m. on March 1, 2013 with a retention period of one month, it would be retained for 30 days.

```
# snapshot create test22 mtree /backup retention 1mos
```

snapshot expire

```
snapshot expire snapshot mtree mtree-path [retention {date | period | forever}]
```

Set or reset the retention time of a snapshot. To expire a snapshot immediately, use the `snapshot expire` operation with no options. An expired snapshot remains available until the next `filesystem clean` operation. Role required: admin, limited-admin.

Argument Definitions

snapshot

The name of the snapshot.

mtree mtree-path

The pathname of the MTree for which the snapshot is being created.

retention date

A four-digit year, two-digit month, and two-digit day separated by dots (.), slashes (/), or hyphens (-). With a retention *date*, the snapshot is retained until midnight (00:00, the first minute of the day) of the designated date.

retention period

Number of days, weeks (wks), or months (mos) to retain snapshot. Note there is no space between the number and time period; for example, 4wks. Also, one month equals 30 days. The snapshot is retained until the same time of day it was created. The retention period must be set in days only.

retention forever

The snapshot does not expire.

snapshot list

```
snapshot list mtree mtree-path | tenant-unit tenant-unit
```

View a list of snapshots of a specific MTree. The display shows the snapshot name, the amount of pre-compression data, the creation date, the retention date, and the status. The status may be blank or expired. Role required: admin, limited-admin, user, backup-operator, tenant-admin, tenant-user, security, none.

Example 212 Argument Definitions

tenant-unit

A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system.

```
# snapshot list mtree /data/col1/ddmtree1
```

snapshot rename

```
snapshot rename snapshotnew-name mtree mtree-path
```

Rename a snapshot for a specific MTree. Role required: admin, limited-admin.

Example 213

To change the name from snap1 to new-snap1 for an MTree named /newMTree, enter:

```
# snapshot rename snap1 new-snap1 mtree /backup
```

snapshot schedule

```
snapshot schedule add name mtrees mtree-list
```

Add multiple MTrees to a single snapshot schedule. Separate multiple MTrees with colons. Role required: admin, limited-admin.

```
snapshot schedule create name [mtrees mtree-list] [days days]
time time [,time ...] [retention period] [snap-name-pattern
pattern]
```

```
snapshot schedule create name [mtrees mtree-list] [days days]
time time every mins [retention period] [snap-name-pattern
pattern]
```

```
snapshot schedule create name [mtrees mtree-list] [days days]
time time-time [every <hrs | mins>] [retention period] [snap-
name-pattern pattern]
```

Use these commands to create a snapshot schedule for multiple MTrees. Command arguments determine the duration of the schedule. (Note the different arguments for specifying time interval.) Role required: admin, limited-admin.

CAUTION

The retention period must be set in days only. If set in hours an error message appears stating your data is in danger and the command will fail.

Example 214

In the following example, snapshots are spaced one minute apart.

```
# snapshot schedule create sml mtrees /data/coll/m1 time
00:00-23:00 every 1mins retention 1days
```

```
snapshot schedule del name mtrees mtree-list
```

Remove a list of MTrees from a schedule. Separate multiple MTrees with colons. Role required: admin, limited-admin.

```
snapshot schedule destroy [name | all]
```

Remove the name of a schedule. Role required: admin, limited-admin.

```
snapshot schedule modify name [mtrees mtree-list] [days days]
time time [,time ...] [retention period] [snap-name-pattern
pattern]
```

```
snapshot schedule modify name [mtrees mtree-list] [days days]
time time every mins [retention period] [snap-name-pattern
pattern]
```

```
snapshot schedule modify name [mtrees mtree-list] [days days]
time time-time [every hrs | mins] [retention period] [snap-name-
pattern pattern]
```

Use these commands to modify a snapshot schedule. Command arguments determine the duration of the schedule. (Note the different arguments for specifying time interval.) Role required: admin, limited-admin.

```
snapshot schedule reset
```

Reset a snapshot schedule and delete all snapshot schedules. Role required: admin, limited-admin.

⚠ CAUTION

This command deletes the previous schedule without prompting the user.

```
snapshot schedule show [name | mtrees mtree-list | [tenant-unit  
tenant-unit]]
```

Show a specific schedule and show schedules associated with a specific MTree.

Separate multiple MTrees with colons. To show a list of schedules, enter the command with no options. Role required: admin, limited-admin, user, backup-operator, tenant-admin, tenant-user, security, none.

Argument Definitions**tenant-unit (Optional)**

The basic unit of a multi-tenancy configuration. A tenant unit is a secure, isolated partition for tenant-specific data and control flow within a Data Domain system.

snapshot

CHAPTER 39

snmp

The `snmp` command enables or disables SNMP access to a Data Domain system, adds community strings, gives contact and location information, and displays configuration settings.

SNMP management requires two primary elements: an SNMP manager and an SNMP agent. An SNMP *manager* is software running on a workstation from which an administrator monitors and controls the different hardware and software systems on a network. These devices include, but are not limited to, storage systems, routers, and switches.

An SNMP *agent* is software running on equipment that implements the SNMP protocol. SNMP defines how an SNMP manager communicates with an SNMP agent. For example, SNMP defines the format of requests that an SNMP manager sends to an agent and the format of replies the agent returns.

From an SNMP perspective a Data Domain system is a read-only device, with one exception: A remote machine can set the SNMP location, contact, and system name on a Data Domain system. The `snmp` command enables administrative users to configure community strings, hosts, and other SNMP MIB variables on the Data Domain system.

With one or more trap hosts defined, a Data Domain system takes the additional action of sending alert messages as SNMP traps, even when the SNMP agent is disabled.

This chapter contains the following topics:

• snmp change history	382
• snmp guidelines and restrictions	382
• snmp add	382
• snmp_debug	383
• snmp del	383
• snmp disable	384
• snmp enable	384
• snmp reset	384
• snmp set	385
• snmp show	385
• snmp status	386
• snmp user	386

snmp change history

There have been no changes to this command in this release.

snmp guidelines and restrictions

- Data Domain systems support MIB access from management stations using SNMPv1, v2C, and v3.
- Data Domain system can send traps using SNMP v2c or SNMP v3.
- Default port 161 is used for inbound/outbound, read/write SNMP access. Default port 162 is used for outbound traffic for SNMP traps.
- Spaces, tabs, colons, semicolons, U.S. dollar signs, and quotation marks cannot be used in community strings.
- To change multiple settings quickly and avoid restarting SNMP, run the `snmp disable` command, change the settings, and then run `snmp enable`.

snmp add

`snmp add ro-community community-string-list [hosts host-list]`
 Add one or more community strings for read-only access to the Data Domain system. A common string for read-only access is *public*. To grant access to specific hosts, replace *host-list* with one or more hostnames. Role required: admin, limited-admin.

Example 215

A valid host list can include both hostnames and IP addresses.

```
hostnameA,hostNameB 10.10.1.2,10.10.1.310.**
```

Example 216

The following command adds the *public* community string for read-only access from host *host.emc.com*.

```
# snmp add ro-community public hosts host.emc.com
```

`snmp add rw-community community-string-list [hosts host-list]`
 Add one or more community strings for read/write access to the Data Domain system. A common string for read/write access is *private*. To grant access to specific hosts, replace *host-list* with one or more hostnames. Role required: admin, limited-admin.

Example 217

The following command adds the *private* community string for read-write access from host *host.emc.com*.

```
# snmp add rw-community private hosts host.emc.com
```

```
snmp add trap-host host-name-list[:port] [version {v2c | v3}]  

[{community community | user user}]
```

Add one or more trap hosts to receive the SNMP traps generated by the Data Domain system. Note that alerts are also sent as traps, even when the local SNMP agent is disabled.

Replace *host-name-list* with one or more hostnames or IP addresses. By default, port 162 is used to send traps, but another port may be assigned. For SNMPv1 and v2c specify the version and the pre-existing community (username). For SNMPv3, specify the SNMPv3 username. Role required: admin, limited-admin.

Example 218

The following command adds trap host *admin12*.

```
# snmp add trap-host admin12 version v2c community public
```

snmp_debug

```
snmp debug-level set {none | packet | error | all}
```

Provides debug options for the SNMP component. This command allows debug levels of none, packet, error, and all. The "none" turns off additional debug information. "Packet" provides the input and output packets in a readable format. "Error" is an intermediate level and only shows output errors. All is a verbose level that displays all possible information.

You can display debug information by using the `log watch` as shown in the following example:

```
log watch debug/sm/snmpd.log
  2015-12-21 10:23:12 NET-SNMP version 5.4.2.1
  2015-12-21 10:23:28 Received SNMP packet(s) from UDP:
[127.0.0.1]->[127.0.0.1]:47393
```

Once the required information is retrieved, you should reset the debug level back to the default ("none"). In addition, you should monitor the size of the file `/ddr/var/log/debug/sm/snmpd.log` to ensure it is not too large. If the file is too large, you can disable the snmp service, delete the file, and then reenabling the service. Role required: admin, limited-admin.

```
snmp debug-level show
```

Shows the current debug status. Role required: admin, limited-admin.

snmp del

```
snmp del ro-community community-string-list [hosts host-list]
```

Delete one or more community strings or hosts from the read-only access list. To display the read-only access list, enter `snmp show ro-communities`. Role required: admin, limited-admin.

```
snmp del rw-community community-string-list [hosts host-list]
```

Delete one or more community strings or hosts from the read-write access list. To display the read-write access list, enter `snmp show rw-communities`. Role required: admin, limited-admin.

Example 219

Example 219 (continued)

The following command deletes host *myhost.emc.com* from the community string *private*.

```
# snmp del rw-community private hosts myhost.emc.com
```

Example 220

The following command deletes the community *private* and all associated hosts.

```
# snmp del rw-community private
```

```
snmp del trap-host host-name-list [version {v2c | v3}]
```

Delete one or more hosts from the list of SNMP trap hosts. To display the list of trap hosts, enter `snmp show trap-hosts`. Include the SNMP version in the command to list the trap hosts for that version. Role required: admin, limited-admin.

Example 221

The following command deletes trap host *admin12*.

```
# snmp del trap-host admin12
```

snmp disable

```
snmp disable
```

Disable SNMP and close port 161. Role required: admin, limited-admin.

snmp enable

```
snmp enable
```

Enable SNMP and open port 161. Role required: admin, limited-admin.

snmp reset

```
snmp reset
```

Reset the SNMP agent configuration to the default values. Role required: admin, limited-admin.

```
snmp reset ro-communities
```

Reset the list of read-only community strings to the default values. Role required: admin, limited-admin.

```
snmp reset rw-communities
```

Reset the list of read-write community strings to the default values. Role required: admin, limited-admin.

```
snmp reset sysContact
```

Reset the SNMP administrative contact MIB variable to the default value or to an empty string if the system value is empty. Role required: admin, limited-admin.

```
snmp reset sysLocation
```


Reset the system location MIB variable to the default value or to an empty string if the system value is empty. Role required: admin, limited-admin.

```
snmp reset trap-hosts
```

Reset the list of SNMP trap receiver hosts to default values. Role required: admin, limited-admin.

snmp set

```
snmp set engineID
```

Configure a unique SNMP engine ID for the Data Domain system. The engine ID must be between 5 and 34 hexadecimal characters. Role required: admin.

```
# snmp set engineID 00112233445566778899AABBCCDDEEFF
SNMP engineID: [Hex] 00112233445566778899AABBCCDDEEFF
```

```
snmp set sysContact sysContact
```

Set the SNMP administrative contact MIB variable using a text string such as an email address. The SNMP `sysContact` MIB variable differs from the value set with the `config set admin-email` command option. However, if the SNMP MIB variables are not set with the SNMP commands, the variables default to the system values given with the `config set` commands. Role required: admin, limited-admin.

```
snmp set sysLocation sysLocation
```

Set the SNMP physical location MIB variable using a text string. The SNMP `sysLocation` MIB variables differs from the value set with the `config set location` command option. However, if the SNMP MIB variables are not set with the SNMP commands, the variables default to the system values given with the `config set` commands. Role required: admin, limited-admin.

```
snmp set sysNotes sysNotes
```

Set the SNMP system notes MIB variable to a text string to record system-specific data not stored in other SNMP variables. Role required: admin, limited-admin.

snmp show

```
snmp show config [version {v2c | v3}]
```

Use this command to display all SNMP configuration parameters or only those for SNMP V2C or V3. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
snmp show engineID
```

Show the configured SNMP engine ID for the Data Domain system. Role required: admin.

```
# snmp show engineID
SNMP engineID: [Hex] 00112233445566778899AABBCCDDEEFF
```

```
snmp show ro-communities
```

Show the configured SNMP read-only communities and hosts. Role required: admin, limited-admin.

```
snmp show rw-communities
```

Show the configured SNMP read/write communities and hosts. Role required: admin, limited-admin.

```
snmp show stats
```

Show the SNMP operating statistics. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
snmp show sysContact
```

Show the configured SNMP administrative contact. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
snmp show sysLocation
```

Show the configured SNMP physical location MIB variable. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
snmp show sysNotes
```

Show the configured SNMP system notes MIB variable. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
snmp show trap-hosts [version {v2c | v3}]
```

Use this command to display all configured trap hosts or only those for SNMP V2C or V3. Role required: admin, limited-admin, security, user, backup-operator, or none.

snmp status

```
snmp status
```

Display whether SNMP is enabled or disabled. Role required: admin, limited-admin, security, user, backup-operator, or none.

snmp user

```
snmp user add user-name access {read-only | read-write}
[authentication-protocol {MD5 | SHA1} authentication-key auth-
key [privacy-protocol {AES | DES} privacy-key priv-key]]
```

Add an SNMPv3 user to the local system specifying the access rights, authentication protocol, and privacy protocol. The authentication key is used when calculating the digest for the authentication protocol. The privacy key is used as input for the privacy protocol. Role required: admin, limited-admin.

```
snmp user del user-name
```

Delete an SNMPv3 user. To view the configured SNMPv3 users, enter `snmp show config version v3`. Role required: admin, limited-admin.

```
snmp user modify user-name access {read-only | read-write}
[authentication-protocol {MD5 | SHA1} authentication-key auth-
key [privacy-protocol {AES | DES} privacy-key priv-key]]
```

Modify SNMPv3 user settings such as access rights, authentication protocol, and privacy key. To view the configured SNMPv3 users, enter `snmp show config version v3`. Role required: admin, limited-admin.

```
snmp user reset
```

Reset the list of SNMPv3 users to the default values. Role required: admin, limited-admin.

```
snmp user show user-name
```

Display SNMPv3 user information. To view the configured SNMPv3 users, enter `snmp show config version v3`. Role required: admin, limited-admin, security, user, backup-operator, or none.

CHAPTER 40

storage

The `storage` command adds, removes, and displays disks and LUNs belonging to active and archive storage tiers. Tiered storage enables the Data Domain system to use different types of storage devices.

System storage for a filesystem or associated RAID disk group consists of two storage tiers: one active and one archive. The active tier uses one active unit of storage, and the archive tier uses one or more retention units of storage.

This chapter contains the following topics:

- [storage change history](#) 388
- [storage guidelines and restrictions](#) 388
- [storage add](#) 388
- [storage migration](#) 389
- [storage remove](#) 393
- [storage sanitize](#) 394
- [storage show](#) 394

storage change history

There have been no changes to this command in this release.

storage guidelines and restrictions

- After adding disks or LUNs to storage tiers, the storage must be provisioned by creating or expanding the filesystem.
- Available LUNs may be removed from a tier to use as a RAID hot spare.

storage add

```
storage add [tier {active | archive | cache |cloud}]
{enclosures <enclosure-list> | disks <disk-list> | <LUN-list>
[spindle-group <1-16>]}
```

Add storage devices to a tier. Device types include all disks in an enclosure, multiple enclosures, one or more disks, one or more LUNs, or spindle group. Disks or LUNs must be in the `Unknown` state to be added to the designated tier, after which the state changes to `Available`. This command cannot be used on dataless head (DLH) units. The default spindle group is 1.

When adding storage, consider the following guidelines.

- Specify `<enclosure-list>` as: A comma or space-separated list formatted `{<enclosure-id>[:<pack-id>]}` or `<enclosure-id>-<enclosure-id>`.
- Specify `<disk-list>` as: A comma or space-separated list formatted `<enclosure-id>.<disk-id>`.
- Specify `<LUN-list>` as: A comma or space-separated list formatted `{dev<disk-id> or dev<disk-id>-<disk-id> .`
- Each system model tier supports a maximum storage quantity based on the system capacity and the installed memory. It is a good practice to add no more than the storage quantity supported by your system. Although you can use the `storage add` command to add storage beyond the supported capacity, an error is reported if you attempt to use the unsupported storage with the `filesystems create` or `filesystems expand` command.
- If adding a disk to an enclosure on the active tier and if there is already a disk group in the enclosure, the disk becomes a spare, not available. This is because if you add a disk and it becomes available, there is no way for the available disk to become spare. Spares are only created when a disk group is created within the enclosure. This rule also applies to the head unit.
- If there is not a disk group in the enclosure (other disks are available or spare), the disk becomes available.
- If the `tier` option is excluded, the storage will be added to the active tier by default.
- For DD3300 DD3300 systems, specify the value in the `Disk` column of the `disk show hardware` command output as the disk ID when adding storage.

Note

The `storage add dev` *disk-id* command option is allowed only after running the command option `storage add enclosure` *enclosure-id* to add the shelf.

Role required: admin, limited-admin.

Argument Definitions**dev** *disk-id* [*spindle-group 1-16*]

Specifies the device and spindle group to be added. To see the available disk IDs, enter `disk show hardware`.

disk *enclosure-id.disk-id*

Specifies the disk to be added with the associated enclosure. To see the available disk IDs with enclosure information, enter `disk show hardware`.

enclosure *enclosure-id:[pack-id]*

Adds all disks in the specified enclosure or pack to the specified tier. If you do not specify a pack within a DS60 enclosure, the system adds all valid packs that have not been previously added. Use this command multiple times to add some packs to the active tier and some to the archive tier.

tier {*active | archive | cache*}

Specifies whether the storage is to be added to the active, or archive, or cache tier. If no tier is specified, the storage is added to the active tier. If an Extended Retention license is present, the storage tier must be specified.

Example 222

To add disks in two different enclosures to the active tier:

```
# storage add enclosure 2
# storage add enclosure 5
```

Figure 5 Output: `storage add enclosure 2 tier cache`

```
#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
Notify filesystem...done
```

storage migration

`storage migration finalize`

Finalize the storage migration. Remove the configuration associated with the source enclosures (such as disk groups), remove the migration destination flag on the destination enclosures, and restart the filesystem using only the storage on the destination enclosures. Role required: admin, limited-admin.

Figure 6 Output: 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.
```

storage migration option reset throttle

Reset the throttle used during storage migration to the default value. Role required: admin, limited-admin.

storage migration option set throttle {low | medium | high}

Set the throttle that is used during storage migration. 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. Role required: admin, limited-admin.

Example 223

```
# storage migration option set throttle high
Throttle for storage migration set to high
```

storage migration option show throttle

Display the throttle that is used during storage migration. Role required: admin, limited-admin.

```
storage migration precheck source-enclosures enclosure-list
destination-enclosures enclosure-list
```

Perform checks to determine if data migration can proceed from one or more source enclosures to one or more destination enclosures. For example, the storage on the destination enclosure set must be equal to or larger than the storage on the source enclosure set.

The enclosure list can include one or more enclosure numbers, separated by commas or space characters. To specify a pack within a DS60 enclosure, use the format *enclosure_number.pack*. Role required: admin, limited-admin.

```
#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"
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 resume

Resume a suspended storage migration. Role required: admin, limited-admin.

storage migration show history

Display the history of completed migrations. Role required: admin, limited-admin.

Figure 7 Output: storage migration show history

```
# storage migration show history
Id  Source      Source Enclosure  Dest  Dest Enclosure  Status  Start Time  End Time
   Enclosure*  Serial No.      Enclosure*  Serial No.
--  -
2   9:0         SHU952400106A23  7:0   SHU9524084G055B Finalized Sat Aug 8 11:59:37 2015 Mon Aug 10 11:10:11 2015
1   9:0         SHU952400106A23  7:0   SHU9524084G055B Finalized Thu Aug 6 16:39:55 2015 Fri Aug 7 10:28:07 2015
   8:0         SHU9524084G04LR
--  -
(*) Enclosure ids at migration start time.
```

storage migration start source-enclosures *enclosure-list*

destination-enclosures *enclosure-list*

Initiate data migration from one or more source enclosures to one or more destination enclosures. The storage on the destination enclosure set must be equal to or larger than the storage on the source enclosure set. If necessary the command reserves space inside the file system to account for different types of drives.

Note

Storage migration does not start when disks are rebuilding in the source enclosures. If a disk in any enclosure requires rebuilding after storage migration starts, the migration is suspended to speed up the rebuilding process. When the rebuild is complete, the migration process automatically resumes.

The enclosure list can include one or more enclosure numbers, separated by commas or space characters. To specify a pack within a DS60 enclosure, use the format *enclosure_number.pack*. Role required: admin, limited-admin.

```
#storage migration start 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

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

Display the status of storage migration. Role required: admin, limited-admin.

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.

```
#storage migration status
```

Migration status needs to report 'reserving space' and 'preparing' as one of the 'States' now.

Id	Source Enclosure (s)	Destination Enclosure (s)	State	Percent Complete	Estimated Time to Complete	Current Throttle Setting
1	2	11	Reserving Space	40%	-	-
1	2	11	Preparing	10%	-	-
1	2	11	Migrating	1%	80 hrs 25 mins	medium

```
storage migration suspend
```

Pause the storage migration. You can continue a suspended storage migration using `storage migration resume`. Role required: admin, limited-admin.

storage remove

```
storage remove {enclosures <enclosure-list> | disks <disk-list> | <LUN-list> [spindle-group <1-16>]}
```

Remove storage devices from the tier, including all disks in an enclosure, multiple enclosures, one or more disks, or one or more LUNs. You can also remove a disk from a DLH unit. When a device is removed the state changes to `Unknown`.

- Specify `<enclosure-list>` as: A comma or space-separated list formatted `{<enclosure-id>[:<pack-id>]}` or `<enclosure-id>-<enclosure-id>`.
- Specify `<disk-list>` as: A comma or space-separated list formatted `<enclosure-id>.<disk-id>`.
- Specify `<LUN-list>` as: A comma or space-separated list formatted `{dev<disk-id> or dev<disk-id>-<disk-id> .`

This command cannot remove an In Use disk if doing so exceeds the minimum number allowed by the RAID scheme. This command also cannot remove a disk if the disk is a spare or an In Use LUN. Role required: admin, limited-admin.

Argument Definitions

dev *disk-id*

Removes the specified device.

disk *enclosure-id.disk-id*

Removes the specified disk.

enclosure *enclosure-id[:pack-id]*

Removes all disks in the specified enclosure or pack.

storage sanitize

`storage sanitize abort enclosure enclosure-id[:pack-id]`
 Abort the storage sanitize task for the specified storage location. Role required: admin, limited-admin.

`storage sanitize resume enclosure enclosure-id[:pack-id]`
 Resume a suspended storage sanitize task at the specified storage location. Role required: admin, limited-admin.

`storage sanitize start enclosure enclosure-id[:pack-id]`
 Initiate the storage sanitize task to remove (zero out) all data from all disks in the specified storage location. All disks in the specified location must be in the unknown state. You cannot use this command to sanitize disks that are being used by the system. Role required: admin, limited-admin.

Example 224

```
# storage sanitize start enclosure 2
** This operation will take approximately 50+ hours to complete.
   Do you want to continue? (yes|no) [no]:
```

`storage sanitize status [enclosure enclosure-id[:pack-id]]`
 Display the status of all sanitize tasks or only the task for specified storage location. The status can be any of the following: COMPLETED, STARTED, STOPPED, SUSPENDED or ABORTED. This command also displays the percent complete for unfinished sanitize tasks. Role required: admin, limited-admin.

`storage sanitize suspend enclosure enclosure-id[:pack-id]`
 Pause the storage sanitize task for the specified storage location. You can continue a suspended sanitize task using `storage sanitize resume`. Role required: admin, limited-admin.

storage show

`storage show {all | summary | tier {active | archive | cache | cloud}}`

Display information about the disk groups, disks, and storage capacity of the file system. The information that appears depends on the system configuration. All systems display the Active tier details table and summary information about the storage tiers. If Extended Retention is in use, the Archive tier details table also appears. Additional tables may appear for Storage addable disks, Storage expandable disks, and Shelf Capacity License information. Role required: admin, limited-admin, security, user, backup-operator, or none.

Argument Definitions

all

Displays storage information for the active and archive tiers.

summary

Displays the states of the disk drives and a count of the disks in each state.

tier {active | archive | cache | cloud}

Specifies the tier for which you want to display storage information.

Figure 8 Output: storage show all

```
#storage show all
Active tier details:
Disk      Disks              Count  Disk      Additional
Group     -----            -----  Size     Information
-----
dg0       1.1-1.7, 1.9-1.12  11     931.5 GiB
(spare)   1.8                1     931.5 GiB
-----

Current active tier size: 8.1 TiB
Active tier maximum capacity: 3.6 TiB

Cache tier details:
Disk      Disks              Count  Disk      Additional
Group     -----            -----  Size     Information
-----
dgX       2.1-2.12          12     931.5 GiB
-----

Current cache tier size: 8.1 TiB
Cache tier maximum capacity: 3.6 TiB

Storage addable disks:
Disk      Disks              Count  Disk      Enclosure  Shelf Capacity  Additional
Type     -----            -----  Size     Model     License Needed  Information
-----
(unknown) 2.13-2.15        3     931.5 GiB  ES30      N/A             cache only
-----

Shelf Capacity License:
License   Model   Total    Used    Remaining
-----
CAPACITY-ACTIVE DS60   43.6 TiB  0.0 TiB  43.6 TiB
<show license for cache tier>
-----
```

Figure 9 Output: disk show hardware

```
#disk show hardware
```

Disk (enc/disk)	Slot	Manufacturer/Model	Firmware	Serial No.	Capacity	Type
1.1	1	WDC_WD1002FBYS-02A6B0	03.00C06	WD-WMATV1675254	931.51 GiB	SATA
1.2	2	WDC_WD1002FBYS-02A6B0	03.00C06	WD-WMATV1677124	931.51 GiB	SATA
...						
2.1	1	Hitachi_HUA721010KLA330	GKAOA70M	GTF002PBHJ6L9F	931.51 GiB	SAS_SSD
2.2	2	Hitachi_HUA721010KLA330	GKAOA70M	GTE002PBHGHU1E	931.51 GiB	SAS_SSD
2.3	3	Hitachi_HUA721010KLA330	GKAOA70M	GTE002PBHH7AVE	931.51 GiB	SAS_SSD
2.4	4	Hitachi_HUA721010KLA330	GKAOA70M	GTF002PBHJ7Z1F	931.51 GiB	SAS_SSD
2.5	5	Hitachi_HUA721010KLA330	GKAOA70M	GTE002PBHHWUGE	931.51 GiB	SAS_SSD
2.6	6	Hitachi_HUA721010KLA330	GKAOA70M	GTE002PBHHV3RE	931.51 GiB	SAS_SSD
2.7	7	Hitachi_HUA721010KLA330	GKAOA70M	GTE002PBHH1M2E	931.51 GiB	SAS_SSD
2.8	8	Hitachi_HUA721010KLA330	GKAOA70M	GTE002PBHJ3VVE	931.51 GiB	SAS_SSD
2.9	9	Hitachi_HUA721010KLA330	GKAOA70M	GTE002PBHH4Y1E	931.51 GiB	SAS_SSD

Output Definitions

Additional Information

Displays additional information regarding the disk group.

Expandable

The `Expandable` entry indicates a system or enclosure that can provide additional storage if the proper license is added. For expandable systems, the `Storage expandable disks` table appears and the `Additional Information` column in that table displays the capacity in use.

Migration destination

The disk is in use as the destination for storage migration.

Migration source

The disk is in use as the source for storage migration.

Pack *n*

This entry identifies the DS60 pack that contains the disks.

Capacity License Needed

If a license is needed to use the full capacity of the enclosure, the license is indicated in this column in the `Storage expandable disks` table. The abbreviation `N/A` in this column indicates that the enclosure does not require a capacity license, or that part of the enclosure is within a tier and the capacity license for the entire enclosure is accounted for.

Count

Shows a count of the disks in the disk group or in the disk state identified in the `Disk Group` column.

Disk Group

Identifies the configured disk groups and the state of disk slots that are not actively participating in a disk group. The following are the disk states that can appear in the `Disk Group` column.

absent

No disk is in the disk slot.

available

Any of the following:

- A previously unknown disk or LUN added to a tier by the `storage add enclosure` command option.
- DD Extended Retention system only: a previously `In Use` disk or LUN deleted from a retention unit by the `filesys archive unit del` command option. This operation reverts the disk or LUN to available storage in the archive tier.
- A previously failed disk in an expansion shelf populated with other disks belonging to a tier that is not primarily composed of disk group disks, and whose partition was destroyed by the `disk unfail` command.

Failed

Tiered storage (`Available`, `Spare`, or `In Use`) removed from the tier automatically by the disk subsystem, or explicitly by an administrative user. `Failed` may also indicate unknown or foreign storage explicitly changed to the `Failed` state.

Foreign

A disk belonging to a third-party vendor.

In Use

Storage that is part of an active filesystem or associated RAID disk group.

Spare

A disk that can be used as a RAID hot spare through RAID reconstruction. Spare disks can be used to create or expand the filesystem.

Spare (reconstruction)

A spare disk that is pending or undergoing RAID reconstruction, which puts filesystem data into what the formerly spare disk and then makes the disk an integral part of a disk group. After RAID reconstruction of a spare disk completes, the disk is part of a RAID disk group.

unknown

A blank disk inserted into the disk slot, or a disk failed by a RAID system.

Disk States

When the summary argument is specified, this column displays the operational states of system disks.

Destination

The disks in this row are in use as a destination for storage migration.

In Use

The disks in this row are part of the system storage in use.

Migrating

The disks in this row are in use as a source for storage migration.

Spare

The disks in this row are reserved for use as a spares.

Disk Size

The size of the disks in the disk group or disk state.

Disks

Identifies the disks within a disk group using the format enclosure_number.disk_number.

Enclosure Model

This column appears in the Storage expandable disks table and indicates the enclosure model that contains the expandable disk group.

Shelf Capacity License

This table lists the storage related licenses in use by the system enclosures.

License

Displays either Capacity-Active for an active tier license or Capacity-Archive for an extended retention tier.

Model

Displays the enclosure model number to which the license applies.

Remaining

Displays the enclosure capacity available for use.

Total

Shows the total enclosure capacity supported by the license.

Used

Displays the enclosure capacity in use.

storage

CHAPTER 41

support

The `support` command manages bundles (Data Domain log files), traces (performance log files, also known as `perf.logs`), and file lists (file names under `/ddvar`) from a customer Data Domain system. This command also configures the ConnectEMC transport feature for securely transmitting information to Data Domain.

This chapter contains the following topics:

- [support change history](#) 402
- [support bundle](#) 403
- [support connectemc](#) 403
- [support coredump](#) 404
- [support notification](#) 406

support change history

New commands in DD OS 6.1.2

`support coredump split filename by n {MiB|GiB}`

Split the specified coredump file into chunks of the specified size. A single file cannot be split into more than 20 chunks. The smallest allowed size for a chunk is 1 MB. An MD5 checksum is created for the split coredump file. Role required: admin.

Modified arguments in DD OS 6.1.2

`support connectemc device register ipaddr esrs-gateway [host-list] [ha-peer ipaddr]`

The `esrs-gateway host-list` parameter supports specifying multiple Secure Remote Services gateways.

`support connectemc device unregister [host-list]`

The `esrs-gateway host-list` parameter supports specifying multiple Secure Remote Services gateways.

`support connectemc device update [new-ipaddr] [ha-peer new-ipaddr] [esrs-gateway host-list]`

The `esrs-gateway host-list` parameter supports specifying multiple Secure Remote Services gateways.

`support bundle create {files-only file-list | traces-only} [and-upload [transport {http|https}]]`

Removed support for the `[and-upload [transport {http|https}]]` option.

`support bundle create default [with-files file-list] [and-upload [transport {http|https}]]`

Removed support for the `[and-upload [transport {http|https}]]` option.

Removed commands in DD OS 6.1.2

`support bundle option reset http-proxy`

Reset the proxy server configuration. Role required: admin, limited-admin.

`support bundle option set http-proxy proxy-server name or ip address [http-proxy-port <port number>]`

Specify a proxy server to use for forwarding support bundles to Data Domain. The default value, port 80, applies if the port number is omitted. Role required: admin, limited-admin.

`support bundle option show http-proxy`

Display the proxy server configuration. Role required: admin, limited-admin.

`support bundle resume bundle-name [transport {http|https}]`

If a support bundle upload fails before completion, use this command to resume the upload from where it left off (instead of starting from the beginning). When

uploading using an HTTP proxy server, you must include the `transport http` option. The default transport is HTTPS. Role required: admin, limited-admin.

```
support bundle upload bundle-name [transport {http | https}]
```

Upload the specified bundle to the support server. To list the system bundles available, enter `support bundle list`. When uploading using an HTTP proxy server, you must include the `transport http` option. The default transport is HTTPS. Role required: admin, limited-admin.

```
support connectemc status
```

Display the ConnectEMC status. Role required: admin, limited-admin.

support bundle

```
support bundle create {files-only file-list | traces-only}
```

Compress listed files into bundle and upload if specified. File names in a list must be separated by a space or a comma. The system automatically deletes the oldest support bundle if five support bundles exist on the system. Role required: admin, limited-admin.

```
support bundle create default [with-files file-list]
```

Compress default and listed files into bundle and upload if specified. File names in a list must be separated by a space or a comma. The maximum number of support bundles allowed is five. Role required: admin, limited-admin.

```
support bundle delete {bundle-name-list | all}
```

Delete some or all of the support bundles on the system. File names in a bundle list must be separated by a space or a comma. To list the system bundles available, enter `support bundle list`. Role required: admin, limited-admin.

```
support bundle list
```

List all support bundles on system. Role required: admin, limited-admin.

support connectemc

```
support connectemc device register ipaddr esrs-gateway [host-list] [ha-peer ipaddr]
```

Register the system to the Secure Remote Services gateway. Use the *host-list* parameter to specify multiple Secure Remote Services gateways to provide redundancy.

Role required: admin, limited-admin.

```
support connectemc device unregister [host-list]
```

Unregister from the Secure Remote Services gateway. If the notification method for communicating ASUPs and alerts to EMC is ConnectEMC, this command switches the notification method to email after the system unregisters from the last Secure Remote Services gateway. Use the *host-list* parameter to unregister multiple from Secure Remote Services gateway IP addresses. Role required: admin, limited-admin.

Note

Running this command without specifying a gateway unregisters from all the Secure Remote Services gateways configured on the system.

```
support connectemc device update [new-ipaddr] [ha-peer new-ipaddr] [esrs-gateway host-list]
```

Update the system's IP address for the Secure Remote Services gateway. Use the *host-list* parameter to update the system IP address on multiple Secure Remote Services gateways. Running this command without specifying a gateway updates the system IP address on all the Secure Remote Services gateways configured on the system. Role required: admin, limited-admin.

```
support connectemc config show
```

Display the ConnectEMC configuration. Role required: admin, limited-admin.

Example 225

```
# support connectemc config show
ConnectEMC Configuration:
  ESRS gateway IP/hostname:    esrs-gateway-1
  Registered device IP(s):     10.1.1.2
  ESRS gateway IP/hostname:    esrs-gateway-2
  Registered device IP(s):     10.1.1.2
```

```
support connectemc show history [last n {hours | days | weeks}]
```

List the ConnectEMC event messages sent during the specified period. If you do not specify a period, the command displays all messages from the last 24 hours. Role required: admin, limited-admin.

```
# support connectemc show history
Message Type           Time Completed           Result
-----
Autosupport            "2018-04-03 07:01:02"   Success
Alert-summary-email    "2018-04-03 08:00:09"   Success
-----
```

```
support connectemc test
```

Send a test message to Support through the Secure Remote Services gateway to test ConnectEMC operation. Test messages are not included in the message history list. Role required: admin, limited-admin.

Example 226

```
# support connectemc test
Sending test message through ConnectEMC...
Test message successfully sent through ConnectEMC.
```

support coredump

```
support coredump delete {core-file-list | all}
```

Delete the specified coredump files or delete all coredump files. File names in a list must be separated by a space or a comma. To display the coredump files on the system, enter `support coredump list`. Role required: admin, limited-admin.

```
support coredump list
```

List the coredump files on the system. Role required: admin, limited-admin.

```
support coredump split filename by n {MiB|GiB}
```

Split the specified coredump file into chunks of the specified size. A single file cannot be split into more than 20 chunks. The smallest allowed size for a chunk is 1 MB. An MD5 checksum is created for the split coredump file. Role required: admin.

```
# 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.

support notification

```
support notification disable {autosupport | alerts | all}
```

Disable email notification to Data Domain for the specified option. Disabling autosupport disables the daily autosupport email. Disabling alerts disables all alert email, including both current alerts and summary reports. The all option specifies that reporting of both autosupport and alerts is to be disabled. Role required: admin, limited-admin.

```
support notification enable {autosupport | alerts | all}
```

Enable email notification to Data Domain for the specified option. Enabling autosupport enables the daily autosupport email. Enabling alerts enables all alert email, including both current alerts and summary reports. The all option specifies that reporting of both autosupport and alerts is to be enabled. Role required: admin, limited-admin.

```
support notification method reset
```

Use this command to reset the notification method selection from ConnectEMC to legacy email. Role required: admin, limited-admin.

```
support notification method set {email | connectemc}
```

Select `email` to use the legacy unsecure method of sending autosupport and alert messages to Data Domain. Select `connectemc` to send ConnectEMC secure messages. The ConnectEMC method requires a configured system administrator email address (`config set admin-email`). The default method is legacy email. Role required: admin, limited-admin.

```
support notification method show {email | connectemc}
```

Display which Data Domain notification method is selected. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
support notification show {autosupport | alerts | all}
```

Show the notification configuration for the autosupport option, the alerts option, or both. Role required: admin, limited-admin, security, user, backup-operator, or none.

CHAPTER 42

system

The `system` command enables administrative users to perform standard tasks on Data Domain systems, configure a system for Retention Lock Compliance, and view system-level information.

This chapter contains the following topics:

- [system change history](#) 408
- [system availability](#) 408
- [system bash](#) 408
- [system headswap](#) 408
- [system option](#) 408
- [system package](#) 409
- [system passphrase](#) 410
- [system poweroff](#) 412
- [system reboot](#) 413
- [system retention-lock](#) 413
- [system sanitize](#) 413
- [system set](#) 414
- [system show](#) 414
- [system status](#) 423
- [system upgrade](#) 424

system change history

There have been no changes to this command in this release.

system availability

```
system availability reset
```

Reset the system availability information. Role required: admin, limited-admin.

```
system availability show
```

Show the system availability information. Role required: admin, limited-admin, user, backup-operator, or none.

system bash

```
system bash enter [timeout n {hr | min}]
```

Access BASH on the Data Domain system. If a timeout value is not specified, the default timeout value is 10 minutes. A BASH token issued by Data Domain support is required to gain access to BASH. If security officer oversight is enabled on the Data Domain system, the system also prompts for the security officer credentials.

The minimum value allowed for the timeout is 5 minutes, and the maximum is 119,000 hours.

Type `exit` to exit BASH.

Role required: admin.

system headswap

```
system headswap
```

Restore the configuration to a system after replacing the head unit. For additional instructions, see the Chassis Replacement FRU document for the system mode and the *Data Domain System Controller Upgrade Guide*. Role required: admin, limited-admin.

Note

After you enter this command, the system displays a message that reminds you that you need the passphrase for the old system if encryption was enabled on that system. You must type `yes` to continue.

system option

```
system option reset {login-banner}
```

Delete the configuration for the login banner, so that no login banner is displayed. Role required: admin, limited-admin.

```
system option set console {serial | lan | monitor}
```

Set the active console option to one of the following.

- For a Serial Over LAN (SOL) connection, enter `system option set console lan`.
- For a console connection through the serial port, enter `system option set console serial`.

- For a console connection through the monitor port (which is not available on all systems), enter `system option set console monitor`.

Role required: admin, limited-admin.

```
system option set login-banner file
```

Set the login banner file. Role required: admin, limited-admin.

Example 227

To create a banner message for your system, mount the Data Domain system directory, `/ddvar`, from another system, create a text file with your login message in `/ddvar`, and then enter the command to use the system banner. The following command selects a file named `banner` in `/ddvar`:

```
# system option set login-banner /ddvar/banner
```

```
system option show
```

View the configuration for the login banner file and the active console. Role required: admin, limited-admin.

system package

```
system package del file
```

Deletes the specified package file. Role required: admin, limited-admin.

```
system package list [file]
```

If the `file` attribute is omitted, this command lists all files in the `/ddvar/releases` directory, which is where package files are stored. If the `file` attribute is specified, this command lists information about the specified package file. In either case, this command indicates whether the package is signed and if so, whether the signature is valid, and whether the package is production software or support software. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system package show installed
```

Lists the packages that are installed on the system. The listing includes three pieces of information: the package name, the package version, whether or not the package is an upgrade relative to the base DD OS version, and whether the package is production software or support software. Role required: admin, limited-admin, security, user, backup-operator, or none.

Example 228

The following example shows the output if there are add-on packages:

```
sysadmin@koala39: system package show installed
Package      Version
-----
ddr           0.6000.12.2-497289
upgrade*     6.0.0.0-000000
-----
(*) Upgraded package
```

Example 229

The following example shows the output if there are no add-on packages:

```
sysadmin@koala39: system package show installed
Package      Version
-----
```

Example 229 (continued)

```

ddr          0.6000.12.2-497289
upgrade     6.0.0.0-000000
-----

```

system passphrase

```
system passphrase change
```

Change the passphrase used to access the system. You must disable the file system before using this command, and the new passphrase must contain the minimum number of characters configured with the `system passphrase option set min-length` command. Role required: admin, limited-admin. This command requires security officer authorization.

Example 230

The following is an example of a successful passphrase change.

```

# system passphrase change
This command requires authorization by a user having a 'security'
role.
Please present credentials for such a user below.
    Username:
    Password:
Enter current passphrase:
Enter new passphrase:
Re-enter new passphrase:
Passphrases matched.
The system passphrase has changed

```

Example 231

The passphrase change fails in the following example because the new passphrase does not conform to the configured minimum length.

```

# system passphrase change
This command requires authorization by a user having a 'security'
role.
Please present credentials for such a user below.
    Username:
    Password:
Enter current passphrase:
Enter new passphrase:
Re-enter new passphrase:
Passphrases matched.

*** New passphrase does not meet the minimum length policy.

```

```
system passphrase option reset min-length
```

Reset the system passphrase min-length option to the default value of 9. Role required: admin, limited-admin. This command requires security officer authorization.

Example 232

```

# system passphrase option reset min-length
This command requires authorization by a user having a 'security'
role.

```

Example 232 (continued)

```
Please present credentials for such a user below.
Username:
Password:
Passphrase option "min-length" is reset to default(9).
** The current passphrase (length 6) must be changed to meet the new
min-length requirement.
```

```
system passphrase option set min-length length
```

Set the minimum length for the system passphrase. No minimum length is defined for new systems. The range for the minimum length is 1 to 255 characters. Role required: admin, limited-admin. This command requires security officer authorization.

Example 233

```
# system passphrase option set min-length 16
This command requires authorization by a user having a 'security'
role.
Please present credentials for such a user below.
Username:
Passphrase option "min-length" set to 16.
```

Example 234

If you set a passphrase minimum length that is longer than the current passphrase length, DD OS displays a message to remind you to change the current passphrase.

```
# system passphrase option set min-length 20
This command requires authorization by a user having a 'security'
role.
Please present credentials for such a user below.
Username:
Passphrase option "min-length" set to 20.
** The current passphrase (length X) must be changed to meet the new
min-length requirement.
```

```
system passphrase option show [min-length]
```

Show the system passphrase minimum length configuration. Role required: admin, limited-admin.

Example 235

```
# system passphrase option show
Option      Value
-----
min-length  16
-----
```

```
system passphrase set
```

For fresh installations, set the passphrase used to access the system. The passphrase length must be longer than the configured minimum and cannot exceed 255 characters. Role required: admin, limited-admin.

Example 236

```
# system passphrase set
Enter new passphrase:
Re-enter new passphrase:
```

Example 236 (continued)

```
Passphrases matched.  
The passphrase is set.
```

system poweroff

```
system poweroff
```

Shut down the Data Domain system. The command performs an orderly shutdown of file system processes. This command does not power off external storage. Role required: admin, limited-admin.

system reboot

`system reboot`

Shut down and reboot a Data Domain system. The command automatically performs an orderly shutdown of file system processes. Role required: admin, limited-admin.

system retention-lock

`system retention-lock compliance configure`

Configure Retention Lock Compliance on the Data Domain system. Role required: admin, limited-admin. This command option requires security officer authorization.

Note

If a system is currently configured for retention-lock compliance, the interface displays this fact and the message that reconfiguration of retention-lock compliance is disabled.

`system retention-lock compliance enable`

Enable Retention Lock Compliance on the Data Domain system. Role required: admin, limited-admin. This command option requires security officer authorization. See the *Data Domain Operating System Administration Guide* for instructions on configuring and enabling Retention Lock.

`system retention-lock compliance status`

Display the status of the Retention Lock Compliance policy on the system, including system clock skew. Role required: admin, limited-admin. This command option requires security officer authorization.

system sanitize

`system sanitize abort`

Stop the system sanitization process. Role required: admin, limited-admin.

`system sanitize start`

Start the system sanitization process. Note that prior to running sanitization, snapshots created during a previous replication process by another user may continue to hold deleted data. To ensure data is removed from replication snapshots during system sanitization, synchronize all replication contexts prior to beginning the procedure. Role required: admin, limited-admin.

Note

When the `system sanitize start` command is run on a Cloud Tier enabled system, an incorrect status message is displayed saying that sanitization has started. The message should indicate that the command failed.

`system sanitize status`

Check system sanitization process status. Role required: admin, limited-admin.

`system sanitize watch`

Monitor the progress of system sanitization. Role required: admin, limited-admin.

For more information on sanitization and task-based instructions, see the *Data Domain Operating System Administration Guide*.

system set

```
system set date MMDDhhmm [ [CC] YY ]
```

Set the system date and time. Do not use this command if Network Time Protocol (NTP) is enabled. This command option requires security officer authorization if the system is enabled for Retention Lock Compliance.

The data and time format uses the following elements.

- Two digits for the month, *MM* (01 through 12).
- Two digits for the day of the month, *DD* (01 through 31).
- Two digits for the hour, *hh* (00 through 23).
- Two digits for minutes, *mm* (00 through 59).
- Optional: Two digits for the century *CC* and two digits for the year *YY*.

The hour *hh* and minute *mm* variables are entered in 24-hour format with no colon between the hours and minutes. 2400 is an invalid entry. The entry 0000 equals midnight. Role required: admin, limited-admin.

Example 237

You can use either of the following commands (two- or four-digit year) to set the date and time to April 23, 2013, at 3:24 p.m.

```
# system set date 0423152413
# system set date 042315242013
```

system show

```
system show all
```

Show all system information. Note that newer systems, such as DD4500 and DD7200, display the product serial number in the Serial number row and the chassis serial number in the Chassis serial number row. On legacy systems, such as DD990 and earlier, the Serial number row displays the chassis serial number and the Service tag row displays the product serial number. The product serial number remains the same during many maintenance events, including chassis upgrades. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system show date
```

Display the system clock. Role required: admin, limited-admin, tenant-admin, security, user, tenant-user, backup-operator, or none.

```
system show detailed-version
```

Show the version number and release information. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system show eula
```

View the End User License Agreement (EULA). Note if the user is not present during system installation, the Data Domain Technical Consultant can temporarily bypass license acceptance and continue with the installation by pressing Ctrl-C. Otherwise, the user must press Enter to accept the license, which is displayed the first time he or she logs in to the system. See the *Data Domain Operating System Initial Configuration Guide* for details. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system show hardware
```

Display information about slots and vendors and other hardware in a Data Domain system. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system show managing-system
```

Identify on which Data Domain Management Console the Data Domain system was added. Also display details about the Data Domain Management Console, such as the outbound proxy host and port, the date for which the system became managed, and the date of last contact. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system show meminfo
```

Display summary of system memory usage. Output differs between newer systems, such as DD4500 and DD7200, and legacy systems, such as DD990 and earlier. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system show modelno
```

Display the hardware model number of a Data Domain system. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system show nvram
```

Display information about NVRAM cards. If output indicates one or more component errors, an alert notification is sent to the designated group and the Daily Alert Summary email includes an entry citing details of problem.

The normal charge level for batteries is 100 percent, and the normal charging status is enabled. Exceptions occur when the system is new or the card is replaced. In both cases the charge may be less than 100 percent initially; however, if it does not reach 100 percent within three days, or if a battery is not enabled, the card must be replaced. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system show oemid [name | value]
```

Show the system OEM IDs for the system controller and any shelves. The OEM ID consists of a name and a value. Omit the options to display both the name and value, or specify one option to display only that data. On systems with head units and shelves, the OEM identifier of the head unit is displayed first. The output includes IDs for connected enclosures only. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system show performance [raw | fsop | view {legacy | default}
custom-view {state | throughput | protocol | compression |
streams |utilization | mtree-active},...] [duration duration
{hr | min} [interval interval {hr | min}]]
```

Display system performance statistics for a designated interval. If you enter this command without the custom-view argument, the standard performance report appears. Role required: admin, limited-admin, security, user, backup-operator, or none.

Argument Definitions

custom-view

Specifies a custom report that includes only those performance statistics that you specify. To display multiple performance statistics, enter multiple labels in the order in which you want the statistics to appear. For example, `system show performance custom-view state streams`.

duration *duration* {hr | min}

The hours or minutes prior to the current time for which to show data.

fsop

Display the number of each filesystem operation performed per minute.

interval *interval* {hr | min}

The time between each line in the display. To specify the interval, you must also specify the duration.

raw

Show unformatted statistics.

view {legacy | default}

Selects one of two predefined views.

Example 238

To show performance figures of the prior 30-minute duration only, enter:

```
# system show performance duration 30 min
```

Example 239

To show performance figures of the prior 30-minute duration with an interval of 5 minutes between each set of figures, enter:

```
# system show performance duration 30 min interval 5 min
```

Output Definitions: Cache Miss**data**

Percent of data segment lookups that miss in the cache. A high percent indicates poor data prefetching.

meta

Percent of metadata segment lookups that miss in the cache. For each data access, first perform a metadata lookup followed by a data lookup. A high percent indicates poor metadata prefetching.

ovhd

Percent of a compression unit cache block that is unused. Compression regions are stored in fixed size (128 KB) blocks. A high ovhd relative to unus indicates space is being wasted due to cache block fragmentation. In the ideal case, ovhd exactly equals unus.

thra

Percent of compression units that were read and discarded without being used. A high percent indicates cache thrashing.

unus

Percent of compression unit data that is unused. Because a compression unit contains multiple segments, not all segments in a compression region may be used. A high percent indicates poor data locality.

Output Definitions: Compression**gcomp**

Global compression rate.

lcomp

Local compression rate.

Output Definitions: IOPS**total**

Operations per second.

read seq/rand

Sequential and random read operations per second.

write seq/rand

Sequential and random write operations per second.

Output Definitions: Protocol Latency**avg/sdev ms**

The average and standard deviation of the response time for ddfs to service all protocol requests, excluding the time to receive or send the request or reply.

read seq/rand

The response time for sequential and random reads.

write seq/rand

The response time for sequential and random writes.

Output Definitions: SS Load Balance (user/repl)

Indicates the relative load balance across segment storage (segstore) instances. Information under (user/repl) denotes all user-plus-Replicator traffic.

prefetch avg/sdev

Prefetch requests.

stream avg/sdev

The average number of open streams and the standard deviation.

rd

The number of read requests.

rd

Read processes.

tot

The total number of requests.

SS Load Balance (gc)

Denotes type and number of expunge (gc) processes.

wr

The number of write requests.

wr
Write processes.

tot
The total number of gc processes.

Output Definitions: MTree Active

rd
The number of active read streams.

wr
The number of active write streams.

Output Definitions: Protocol

data (MB/s in/out)
Protocol throughput. Amount of data the filesystem can read from and write to the kernel socket buffer.

load
Load percentage (pending ops/total RPC ops *100).

ops/s
Operations per second.

wait (ms/MB in/out)
Time taken to send and receive 1MB of data from the filesystem to kernel socket buffer.

Note

Protocol data includes NFS, CIFS, DD Boost over IP, and DD Boost-managed replication and optimized duplication. Data does *not* include Replication, VTL over Fibre Channel, or DD Boost over Fibre Channel.

Output Definitions: State

C
Cleaning

D
Disk reconstruction

F
Archive data movement

I
Container verification (scrubbing)

M
Fingerprint merge

P
Physical space measurement

- R**
Archive space reclamation
- S**
Summary vector checkpoint
- t**
Storage migration in progress.
- V**
File verification running

Output Definitions: Streams

read seq/rand

The number of sequential and random read streams.

write seq/rand

The number of sequential and random write streams.

r+

The number of reopened read file streams in the past 30 seconds.

rd

The number of active read streams.

Repl in

The number of incoming replication streams.

Repl out

The number of outgoing replication streams.

w+

The number of reopened write file streams in the past 30 seconds.

wr

The number of active write streams.

Output Definitions: Throughput (MB/s)

Read

The read throughput data from the Data Domain system.

Repl Network (in/out)

Network replication throughput into and out of the Data Domain system.

Repl Pre-comp (in/out)

Replication pre-compressed (logical) throughput into and out of the Data Domain system. The value is always zero for collection replication.

Write

The write throughput data to the Data Domain system.

Note

Throughput Read and Write data includes NFS, CIFS, DD Boost over IP and Fibre Channel, VTL, Replication, DD Boost-managed replication and optimized duplication.

Output Definitions: Time Stamp**Date**

The date system performance is being viewed.

Time

The time system performance is being viewed.

Output Definitions: Utilization**CPU avg/max %**

The average and maximum percentage of CPU utilization.

Disk max %

The maximum percentage of disk utilization.

```
system show ports
```

Display information about ports. VTL-related ports do not display unless VTL is enabled. Role required: admin, limited-admin, security, user, backup-operator, or none.

Output Definitions**Connection Type**

The type of connection, such as Ethernet, SAS, VTL, etc.

Firmware

The Data Domain system HBA firmware version.

Hardware Address

A MAC address or WWN. An address followed by an Ethernet port number is a MAC address. WWN is the world-wide name of the Data Domain system SAS HBA on a system with expansion shelves.

Link Speed

The speed in Gbps (Gigabits per second).

Port

The port number. See the model-specific installation and setup guide to match a slot to a port number.

```
system show serialno [detailed]
```

Display the system serial number and also shows whether encryption is enabled. On newer systems, such as DD4500 and DD7200, the system serial number is the product serial number, which remains the same during many maintenance events, including chassis upgrades. On legacy systems, such as DD990 and earlier, the system serial number is the chassis serial number. When the `detailed` argument is specified, the output displays both the system serial number and the chassis serial number, which is different on newer systems and identical on legacy systems. Role required: admin, limited-admin, security, user, backup-operator, or none.

Example 240

The following example is from a newer system that displays the product serial number as the system serial number as well as support for data encryption:

```
# system show serialno detailed
Serial number: APM00141269680
Chassis Serial number: FCNME140100235
```

Example 240 (continued)

```
Data Encryption Supported: Yes
Data Deduplication Supported: No
```

```
system show stats [view {nfs | cifs | repl | net | iostat |
sysstat | ddbost},...] [custom-view view-spec,...] [interval
nsecs] [count count]
```

Display the system statistics collected since the last reboot. If you enter this command without the view or custom-view arguments, the standard statistics report appears.

If the system is too busy to determine a value, the column shows a dash instead of a number. Role required: admin, limited-admin, security, user, backup-operator, or none.

Argument Definitions**column**

Displays output of for each node in column format. Column headings indicate type of stat value.

count *count*

Specifies how many times to display the results. The default count is one. If interval is specified and count is omitted, the count is set to infinite, or until the user presses Ctrl-C.

custom-view *view-spec*,...

Specifies a custom report that includes only those statistics that you specify. Valid entries include any column section label in the standard reports: cpu, state, nfs, cifs, net (for network), disk, nvram, and repl (for replication). To display multiple column sections, enter the column labels in the order in which you want the sections to appear.

interval *nsecs*

When specifying intervals for collecting statistics, the first report is for current activity. Subsequent reports show activity performed during [*interval nsecs*]. The default interval is five seconds.

row

Output for each node is in row format, displayed as a single line for each interval.

view {nfs | cifs | repl | net | iostat | sysstat | ddbost},...

Specifies a variation of the standard statistics report that provides additional statistics for the feature specified. Valid entries are nfs for NFS, cifs for CIFS, repl for replication, net for network, iostat for I/O statistics, sysstat for system statistics, and ddbost for DD Boost statistics.

Example 241

```
# system show stats view nfs interval 2
```

Output Definitions: CPU**aggr busy %**

Average of busy percentage of all CPUs.

aggr max %

Amount of data sent through all interfaces.

State

Indicates system state. See the description of the show system performance command for details on what each letter represents.

Protocol Aggr

total

Operations per second.

load %

Load percentage (pending ops/total RPC ops *100).x.)

data in % MB/s

Protocol throughput. Amount of data the filesystem can read from and write to the kernel socket buffer.

data out % MB/s

Protocol throughput. Amount of data the filesystem can write to the kernel socket buffer.

CIFS

ops/s

I/O and metadata operations per second.

ops/s

I/O and metadata operations per second.

ops/s

I/O and metadata operations per second.

in MB/s

Write throughput.

out MB/s

Read throughput.

NFS

ops/s

I/O operations.

load %

Load percentage (pending ops/total RPC ops *100).x.)

data in % MB/s

Protocol throughput. Amount of data the filesystem can read from and write to the kernel socket buffer.

data out % MB/s

Protocol throughput. Amount of data the filesystem can write to the kernel socket buffer.

wait in ms/MB

Average amount of time spent in ms to receive the amount of data.

wait out ms/MB

Average amount of time spent in ms to send the amount of data.

Net**aggr in MB/s**

Amount of data received through all interfaces.

aggr out MB/s

Amount of data sent through all interfaces.

aggr drop (in)

Number of incoming connections dropped by all interfaces.

aggr drop (out)

Number of outgoing connections dropped on all interfaces.

<Interface> drop (in)

Number of incoming connections dropped by a specific interface.

<Interface> drop (out)

Number of outgoing connections dropped on a specific interface.

```
system show uptime
```

Display the filesystem uptime, the time since the last reboot, the number of users, and the average load. Role required: admin, limited-admin, security, user, backup-operator, or none.

```
system show version
```

Display the Data Domain OS version and build identification number. Role required: admin, limited-admin, tenant-admin, security, user, tenant-user, backup-operator, or none.

Example 242

```
sysadmin@apollo65# system show version
Data Domain OS 0.6000.0.0-528922
sysadmin@apollo65#
sysadmin@apollo65# system show meminfo
Memory Usage Summary
Total memory:      773941 MiB
Free memory:      476625 MiB
Inactive memory:  3353 MiB
Total swap:       5119 MiB
Free swap:        5119 MiB
```

system status

```
system status
```

Display status of fans, internal temperatures, and power supplies. Information is grouped separately for the Data Domain system and each expansion shelf connected to the system. See the *Data Domain Operating System Administration Guide* for details. Role required: admin, limited-admin, security, user, backup-operator, or none.

system upgrade

```
system upgrade [precheck] file
```

Note

This command is deprecated and will be removed from a future release. Use `system upgrade start` or `system upgrade precheck`.

```
system upgrade history
```

Display the history of system upgrades. Role required: admin, limited-admin.

Example 243

```
# system upgrade history
Version      Partition  State      Time
-----
5.5.0.7-448183 1          INSTALL   09/24/14 17:16:17
5.5.1.1-450844 2          UPGRADE   09/25/14 12:50:23
5.5.2.0-461195 1          UPGRADE   12/01/14 08:16:37
5.5.2.0-461359 2          UPGRADE   12/03/14 08:21:00
5.6.0.0-466502 1          UPGRADE   01/13/15 14:12:11
5.7.0.0-466745 2          UPGRADE   01/15/15 07:51:05
5.7.0.0-473398 1          UPGRADE   02/11/15 09:42:44
5.7.0.0-475811 2          UPGRADE   03/04/15 14:32:10
5.7.0.0-475968 1          UPGRADE   03/09/15 11:33:05
-----
```

```
system upgrade option show [support-software]
```

Display whether the Data Domain system allows or prevents the installation of support software. Support software may cause the system to function in unexpected ways. Contact Support for additional information about support software. Role required: admin, limited-admin.

```
system upgrade precheck file
```

Evaluate whether the operating system can be upgraded to the version in the specified file. The system searches for the file in the `/ddvar/releases` directory and on a USB key. The precheck does not upgrade or modify the system, and it does not impact system performance. The precheck evaluates a set of parameters and indicates whether or not there are any issues that will prevent an upgrade.

A precheck cannot start until any active upgrade or precheck completes. Role required: admin, limited-admin.

Note

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

```
system upgrade start file
```

Upgrade the Data Domain system software to the version in the specified file. The system searches for the file in the `/ddvar/releases` directory and on a USB key. An upgrade cannot start until any active upgrade or precheck completes. The upgrade starts with a precheck of the system parameter values required to support an upgrade. If the precheck fails, the command terminates without modifying the system.

Note

System upgrades are not supported while a `storage migration finalize` process is active. If an upgrade fails while a `storage migration finalize` process is active, restart the upgrade after the process completes.

If the precheck is successful, the upgrade begins and the console displays the upgrade status. Other users can monitor the upgrade by entering `system upgrade watch`. To stop the upgrade status display, press `Ctrl-C`. Terminating the upgrade status display does not stop the upgrade. Once started, the upgrade continues to completion.

When the upgrade is nearly complete, the system shuts down the filesystem and reboots. The upgrade may require over an hour, depending on the amount of data on the system. During a system upgrade, a banner appears to warn all logged-in users and any new logged-in users that an upgrade is in progress and that not all DD OS operations are available. See the *Data Domain Operating System Release Notes* for instructions on upgrading Data Domain systems. Role required: admin, limited-admin.

Note

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

```
system upgrade start rpm [force ] [local]
```

Upgrade the Data Domain software RPM with the local parameter specified to also allow the standby node in a high-availability (HA) system to reboot. Role required: admin, limited-admin.

The standby node rejects the command unless the local parameter is specified.

Note

The RPM can be a system RPM, a bundle RPM, a component RPM, or an add-on RPM.

If the upgrade request succeeds, the command polls the status of the upgrade until the current node reboots; the operation finishes with a success status or until an error is reported. The CLI is rejected on the passive node in an HA system unless the local parameter is specified. Yes/no confirmation is required in interactive mode.

```
system upgrade status
```

When an upgrade is in progress, this command displays the current upgrade status and upgrade procedure phase, and then the command terminates. This command does not continually display the upgrade progress.

If no upgrade is in progress, the system displays the completion time and status of the last upgrade. The status shown is not affected or updated in response to any corrective measures or configuration changes made after the completion time. Role required: admin, limited-admin.

Example 244

```
# system upgrade status
Current Upgrade Status: DD OS upgrade succeeded
End time: 2015.04.27:09:08
```

```
system upgrade uninstall package [local]
```

Request that the named package be uninstalled. The package must have been installed as an add-on.

```
system upgrade watch
```

This command displays the current status throughout the precheck or upgrade process, and this command is not limited to the user who initiated the upgrade. You can enter this command after reboot to continue monitoring an upgrade. To terminate the display before the process completes, press Ctrl-C. Role required: admin.

Example 245

```
# system upgrade watch  
There is no upgrade or precheck in progress.
```

CHAPTER 43

user

The `user` command adds and deletes users, manages password aging and strength policies, and displays user roles. A role determines the type of operations a user can perform on the Data Domain system. See the *Data Domain Operating System Administration Guide* for details.

The default administrative account is `sysadmin`. You can change the `sysadmin` password but cannot delete the account.

This chapter contains the following topics:

- [user change history](#) 428
- [user add](#) 428
- [user change](#) 429
- [user del](#) 429
- [user disable](#) 430
- [user enable](#) 430
- [user password](#) 430
- [user reset](#) 435
- [user show](#) 435

user change history

There have been no changes to this command in this release.

user add

```
user add user [role {admin | limited-admin | security | user |
backup-operator | none}] [min-days-between-change days] [max-
days-between-change days] [warn-days-before-expire days]
[disable-days-after-expire days] [disable-date date] [force-
password-change {yes | no}]
```

Add a new locally defined user. A user name must start with a number or a letter. Special characters cannot be used. The user names `root` and `admin` are default names on each Data Domain system and are not available for general use.

The following list describes the roles that can add new users, and the level of users that they can add:

- `sysadmin`: Can add `admin`, `limited-admin`, `backup-operator`, `user`, `none`.
- `admin`: Can add `limited-admin`, `backup-operator`, `user`, `none`.
- `limited-admin`: Can add `backup-operator`, `user`, `none`.

Admin users can create the first security officer role. After the first security-role user is created, only security-role users can add or delete other security-role users. After creating a security role, you must enable security authorization using the `authorization policy` command. See the *Data Domain Operating System Administration Guide* for details on user roles.

Argument Definitions

disable-date

Account is disabled on this date. If not specified, account never expires.

disable-days-after-expire

Account is disabled if inactive for the specified number of days past expiration.

force-password-change

Require that the user change the password during the first login when connecting using SSH or Telnet or through DD System Manager. The default value is `no`, do not force a password change.

max-days-between-change

Maximum number of days before password expires.

min-days-between-change

Minimum number of days allowed before the password can be changed again.

role

The type of user permissions allowed. The default role is `none`. For SMT configurations, the only user role that can be assigned to SU under tenant-units is `none`. See the *Data Domain Operating System Administration Guide* for details.

warn-days-before-expire

Number of days of warning before a password expires.

user change

```
user change password [user]
```

Change the password of a locally defined user. Admin-role users can change the password for any user, and security-role users can change the passwords for other security role users. Users in all other management roles can change only their own passwords. Passwords must comply with the password strength policy, which you can check with the command option `user password strength show`. To display a list of all locally defined users, enter `user show list`.

The following list describes the roles that can change passwords and the level of users that they can change passwords for:

- **sysadmin:** Can change password for itself, admin, limited-admin, backup-operator, user, none.
- **admin:** Can change password for itself, limited-admin, backup-operator, user, none.
- **limited-admin:** Can change password for itself, backup-operator, user, none.
- **security officer:** Can change its own password only.

Role required: admin, tenant-admin, security, user, tenant-user, backup-operator, or none.

```
user change role user {admin | limited-admin | user | backup-operator | none}
```

Change the role of a user.

The following list describes the roles that can change the roles of other users, and the level of users that they can act upon:

- **sysadmin:** Can change admin, limited-admin, backup-operator, user, none.
- **admin:** Can change limited-admin, backup-operator, user, none.
- **limited-admin:** Can change backup-operator, user, none.

Users cannot promote other users to greater or equal roles. For example, an admin user cannot promote a limited-admin user to admin or sysadmin.

No management role is permitted to change the role of a security-role user. 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.

To display a list of all locally defined users, enter `user show list`. See the *Data Domain Operating System Administration Guide* for more information on user roles.

user del

```
user del user
```

Remove any locally defined user except sysadmin and DD Boost users. The sysadmin user cannot be deleted. To delete a user name in use by DD Boost, delete the DD

Boost user first, then use this command to delete the user name. To display a list of all locally defined users, enter `user show list`.

The following list describes the roles that can delete users, and the level of users that they can delete:

- `sysadmin`: Can delete `admin`, `limited-admin`, `backup-operator`, `user`, `none`.
- `admin`: Can delete `limited-admin`, `backup-operator`, `user`, `none`.
- `limited-admin`: Can delete `backup-operator`, `user`, `none`.

Security-role users can delete only security-role users.

Example 246

```
# user del ddbboost1
o ddbboost1 cannot be deleted if referenced by ddbboost
```

user disable

```
user disable user
```

Disable the specified locally defined user account so that the user cannot log on to the Data Domain system. To display a list of all locally defined users, enter `user show list`.

The following list describes the roles that can disable users, and the level of users that they can disable:

- `sysadmin`: Can disable `admin`, `limited-admin`, `backup-operator`, `user`, `none`.
- `admin`: Can disable `limited-admin`, `backup-operator`, `user`, `none`.
- `limited-admin`: Can disable `backup-operator`, `user`, `none`.

Security-role users can only disable security-role users.

user enable

```
user enable user [disable-date date]
```

Enable the specified locally defined user account so that the user can log on to the Data Domain system. To display a list of all locally defined users, enter `user show list`. Admin-role users can enable users in all management roles except the security role.

The following list describes the roles that can enable users, and the level of users that they can enable:

- `sysadmin`: Can enable `admin`, `limited-admin`, `backup-operator`, `user`, `none`.
- `admin`: Can enable `limited-admin`, `backup-operator`, `user`, `none`.
- `limited-admin`: Can enable `backup-operator`, `user`, `none`.

Security-role users can only enable security-role users.

user password

```
user password aging option reset {all | [min-days-between-
change] [max-days-between-change] [warn-days-before-expire]
[disable-days-after-expire]}
```

Reset one or more rules in the default password aging policy to the current default values. New accounts inherit the policy in effect at the time they are created, unless you set different aging options with the `user add` command. The argument definitions are the same as for `user password aging option set`. Role required: admin.

Example 247

```
# user password aging option reset all
Password aging options have been reset.
```

```
user password aging option set {[min-days-between-change days]
[max-days-between-change days] [warn-days-before-expire days]
[disable-days-after-expire days]}
```

Set the default values for the password aging policy. Role required: admin.

Argument Definitions

min-days-between-change *days*

The minimum number of days between password changes that you allow a user. This value must be less than the `max-days-between-change` value minus the `warn-days-before-expire` value. The default setting is 0.

max-days-between-change *days*

The maximum number of days between password changes that you allow a user. The minimum value is 1. The default setting is 90.

warn-days-before-expire *days*

The number of days to warn the users before their password expires. This value must be less than the `max-days-between-change` value minus the `min-days-between-change` value. The default setting is 7.

disable-days-after-expire *days*

The system disables a user account after password expiration according to the number of days specified with this argument. Enter `never` or a number equal to or greater than zero. The default setting is `never`.

Example 248

```
# user password aging option set warn-days-before-expire 14
Password aging options have been set.
```

```
user password aging option show
```

Display the default password aging policy. Role required: admin.

Example 249

```
# user password aging option show
Minimum Days Between Password Change: 0
Maximum Days Between Password Change: 90
Warning Days Between Password Change: 14
Disable Days After Expire: never
```

```
user password aging reset user {all | [min-days-between-change]
[max-days-between-change] [warn-days-before-expire] [disable-
days-after-expire]}
```

Reset one or more rules in the password aging policy for the specified locally defined user to the current default values set by the `user password aging option set` command. The argument definitions are the same as for `user password aging set`. Role required: `admin` for all except security-role users, `security` for security-role users.

```
user password aging set user [min-days-between-change days]
[max-days-between-change days] [warn-days-before-expire days]
[disable-days-after-expire days]
```

Set the password aging policy for the specified locally defined user. To display the locally defined users, enter `user show list`. Role required: `admin` for all except security-role users, `security` for security-role users.

Argument Definitions

min-days-between-change *days*

The minimum number of days between password changes that you allow a user. This value must be less than the `max-days-between-change` value minus the `warn-days-before-expire` value. The default setting is 0 and may be changed using the `user password aging option set` command.

max-days-between-change *days*

The maximum number of days between password changes that you allow a user. The minimum value is 1. The default setting is 90 and may be changed using the `user password aging option set` command.

warn-days-before-expire *days*

The number of days to warn the users before their password expires. This value must be less than the `max-days-between-change` value minus the `min-days-between-change` value. The default setting is 7 and may be changed using the `user password aging option set` command.

disable-days-after-expire *days*

The system disables a user account after password expiration according to the number of days specified with this argument. Enter `never` or a number equal to or greater than zero. The default setting is `never` and may be changed using the `user password aging option set` command.

Example 250

```
# user password aging set user disable-days-after-expire 99
User "user's" password aging information has been updated.
```

```
user password aging show [user]
```

Show the password aging policy for all locally defined users, or for a specified user. Only `admin-role` and `security-role` users can display the policy for other users. `User-role`, `backup-operator-role`, and `none-role` users can check the policy for their own account.

Example 251

```
# user password aging option show
Minimum Days Between Password Change: 0
Maximum Days Between Password Change: 99999
Warning Days Between Password Change: 0
Disable Days After Expire: never
```



```
user password strength reset {all | min-length | min-char-classes | min-one-lowercase | min-one-uppercase | min-one-digit | min-one-special | max-three-repeat | passwords-remembered}
```

Reset one or all of the password strength arguments to the default values. Role required: admin.

Argument Definitions

all

Reset the minimum length to 6 and minimum number of character classes to 1.

min-length

Reset the minimum number of characters in the password to 6.

min-char-classes

Reset the minimum number of character classes to 1.

min-one-lowercase

Reset the requirement for at least one lowercase character to disabled.

min-one-uppercase

Reset the requirement for at least one uppercase character to disabled.

min-one-digit

Reset the requirement for at least one numerical character to disabled.

min-one-special

Reset the requirement for at least one special character to disabled.

max-three-repeat

Reset the requirement for a maximum of three repeated characters to disabled.

passwords-remembered

Reset the number of remembered passwords to 1.

Example 252

```
# user password strength reset min-length
Password strength "min-length" reset to default (6).
```

```
user password strength set {[min-length length] [min-char-classes num_classes] [min-one-lowercase {enabled | disabled}] [min-one-uppercase {enabled | disabled}] [min-one-digit {enabled | disabled}] [min-one-special {enabled | disabled}] [max-three-repeat {enabled | disabled}] [passwords-remembered <0 - 24>]}
```

Set the password strength policy. Specify either min-length or min-char-classes, or both. Role required: admin.

Argument Definitions

min-length

The minimum number of characters in the password. The range is 1 to 100; the default setting is 6.

min-char-classes

The minimum number of character classes. Specify 1, 2, 3, or 4. Valid passwords must contain at least one character from the specified number of classes. The four character classes are lowercase letters, uppercase letters, digits, and special characters.

When DD OS counts the number of character classes, an uppercase letter at the beginning of the password does not count as an uppercase letter. Similarly, a digit at the end of the password does not count as a digit.

min-one-lowercase

Enable the requirement for at least one lowercase character. The default setting is disabled.

min-one-uppercase

Enable the requirement for at least one uppercase character. The default setting is disabled.

min-one-digit

Enable the requirement for at least one numerical character. The default setting is disabled.

min-one-special

Enable the requirement for at least one special character. The default setting is disabled.

max-three-repeat

Enable the requirement for a maximum of three repeated characters. The default setting is disabled.

passwords-remembered

Specify the number of remembered passwords. The range is 0 to 24. The default settings is 1.

Note

If the passwords-remembered value is reduced, the remembered password list remains unchanged until the next time the password is changed. For example, if the passwords-remembered value is changed from 4 to 3, the last four passwords are remembered until the next time the password is changed.

Example 253

```
# user password strength set min-length 10
Specified password strength requirements have been enforced.
```

```
user password strength show
```

Show the current password strength policy. Role required: admin, security, user, backup-operator, or none.

Example 254

```
# user password strength show
Option                               Value
-----                               -
Minimum password length              1
Minimum character classes             1
At least one lowercase character      disabled
At least one uppercase character      disabled
At least one digit                    disabled
At least one special character        disabled
At most three consecutive repeated characters disabled
Passwords remembered                  0
-----                               -
```

user reset

```
user reset
```

This command deletes all locally defined user accounts except *sysadmin* and user accounts for security, DD Boost (role = none), and VDISK. This command also resets the password strength and password aging options to the factory default values. Role required: admin.

Note

This command option is not allowed on Retention Lock Compliance systems.

user show

```
user show active
```

Display a list of users currently logged in. The `tty` column displays the access method for the user.

- Console access appears as `tty#`, where `#` represents a session number for that user.
- SSH, Telnet, and FTP access appears as `pts#`, where `#` represents a session number for that user.
- DD SM access appears as GUI.
- REST service access appears as WEB_SVC.
- Vdisk access appears as API.

Role required: admin, tenant-admin, security, user, tenant-user, backup-operator, or none.

user

Example 255

```
# user show active
User list from node "localhost".
Name      Idle   Login Time      Login From      tty
-----
sysadmin  20h   Thu Oct  9 17:40      ddang.datadomain.com  tty1
sysadmin  19h   Thu Oct  9 17:41      ddang.datadomain.com  pts/0
sysadmin  3m    Thu Oct  9 17:52      ddang.datadomain.com  pts/1
sysadmin  0s    Fri Oct 10 13:42      usendangdlm1.datadomain.com  pts/2
sysadmin  3m    Fri Oct 10 13:37      ::ffff:137.69.74.231  GUI
sysadmin  3m    Fri Oct 10 13:38      ::ffff:137.69.74.230  WEB_SVC
sysadmin  3m    Fri Oct 10 13:37      ddang.datadomain.com  API
-----
5 users found.
```

user show detailed [user]

Show detailed information for a specified user or for all users. Role required: admin or security.

Example 256

```
# user show detailed Tu1
User:                               Tu1
Uid:                                 501
Role:                                 user
Last Login From:                     <unknown>
Last Login Time:                     Mon Jan 14 11:55:49 2013
Status:                               enabled
Password Last Changed:               Mar 16, 2006
Disable Date:                         never
Minimum Days Between Password Change: 0
Maximum Days Between Password Change: 99999
Warning Days Between Password Change: 7
Disable Days After Expire:           never
Force Password Change at Next Login: no
Tenant-unit Roles:
  Tenant-unit      Role
  -----
  Tenant-unit1     tenant-admin
  Tenant-unit2     tenant-user
```

user show list

Display list of system users. Role required: admin or security.

Figure 10 Output: user show list

```
# user show list
User list from node "localhost".
Name      Uid   Role      Last Login From      Last Login Time      Status  Disable Date
-----
sysadmin  100  admin     carvej-dl.datadomain.com  Mon Aug 17 13:02:32 2015  enabled  never
security  500  security  <unknown>             never                 enabled  never
boost3    501  none     <unknown>             never                 enabled  never
boost4    502  none(*)  <unknown>             never                 enabled  never
user      503  user     10.30.187.142         Wed Aug 12 22:46:09 2015  enabled  never
tuesday   504  user     <unknown>             never                 enabled  never
anonymous 505  admin     <unknown>             never                 enabled  never
backup    506  backup-operator <unknown>             never                 enabled  never
new_user  507  none     <unknown>             never                 enabled  never
new1      508  none     <unknown>             never                 enabled  never
first     509  user     10.30.187.142         Tue Aug 11 23:49:20 2015  enabled  never
-----
(*) This user has additional roles for tenant-unit(s).
11 users found.
```


user

CHAPTER 44

vdisk

The `vdisk` command creates and manages virtual disk devices that can be exported as a block-level disk device to an initiator over a Fibre Channel link. These block-level devices can be accessed by an application host, backup host, or a disk sub-system for block-level backup and recovery.

This chapter contains the following topics:

• vdisk change history	440
• vdisk guidelines and restrictions	440
• vdisk config	441
• vdisk device	442
• vdisk device-group	443
• vdisk disable	444
• vdisk enable	444
• vdisk group	445
• vdisk pool	446
• vdisk property	447
• vdisk reset	448
• vdisk show	448
• vdisk static-image	450
• vdisk status	452
• vdisk trace	452
• vdisk user	452

vdisk change history

There have been no changes to this command in this release.

vdisk guidelines and restrictions

- In the current release, virtual disks are only supported for certain solutions that incorporate Data Domain and Symmetrix VMAX systems. See the *Backup Compatibility Guide, Data Domain Operating System* for information about supported configurations.
- If you use virtual disks, using either the VTL feature or DD Boost over Fibre Channel on the same Data Domain system is not supported.
- The output of `filesys show compression` may be misleading for virtual disks. Large virtual disks that contain very little data may show a total compression factor of 100%. This happens because, when you create a virtual disk, the system creates a file of the specified size (such as 2 PB) but does not actually write data to the disk, so no physical space is used. The compression ratio for that file will be extremely high until a substantial amount of real data gets written to the virtual disk.
- The output of `filesys show space` may also be misleading for virtual disks. When you create a virtual disk, the entire amount of space allocated for that virtual disk gets added to the `pre-comp` statistic for the file system. Likewise, when you delete a virtual disk, only the amount of space actually used for data gets added to the `Cleanable` statistic. This behavior is expected.
- Deleting the vdisk device that is assigned to LUN 0 from a vdisk device group is not supported, and will cause the other devices in the group to not be visible on the initiator.
- LUN 0 must be visible to all endpoints.
- Vdisk Fibre Channel operation is expected to continue without user intervention when the Fibre Channel endpoints failover.
- If a ProtectPoint backup description is longer than 1024 characters, the `vdisk static-image show detailed` command truncates the description to 1024 characters.
- If a vdisk static-image property is longer than the UI column width, the `vdisk static-image show detailed` command adds two blank spaces between each property name.
- The maximum values displayed by the `vdisk show config` are system-wide limits inherited from the MTree limit. The maximum number of pools supported ranges from 100 to 256, depending on the specific model of Data Domain system.
- The `vdisk device create` command creates a device of at least, but not necessarily the exact size specified. The only time a device of the exact size specified is created is specified is when the capacity is specified in sectors with the `capacity n sectors` option.
- Vdisk devices use device-locking to prevent data corruption when multiple hosts are running control path operations on a single vdisk device.
- DD Retention Lock is supported on vdisk devices.

vdisk config

Virtual disk clone configuration.

Virtual disk cloning creates copies of virtual disk objects.

Clone operations are supported on the following virtual disk objects:

- Devices
- Device-groups
- Pools

```
vdisk config clone device source-device device-name source-pool
pool-name source-device-group device-group-name destination-
pool pool-name destination-device-group device-group-name
```

Clone a virtual disk device configuration to a specified destination pool and device group. This command creates a new device in the destination pool and device-group with the same configuration parameters as the source device. Role required: admin, limited-admin.

```
vdisk config clone device-group source-device-group device-
group-name source-pool pool-name destination-pool pool-name
[destination-device-group device-group-name]
```

Clone a virtual disk device-group to a specified destination. This command creates a new device-group and devices in the destination pool with the same configuration parameters as the source device-group and devices. The command fails if the destination pool or device-group already exists. Role required: admin, limited-admin.

```
vdisk config clone pool source-pool pool-name destination-pool
pool-name [vdisk-user vdisk-user]
```

Clone a virtual disk pool to a specified destination pool. This command creates a new pool, device-group, and devices with the same configuration parameters as the source pool, device-group, and devices. The command fails if the destination pool. Role required: admin, limited-admin.

```
vdisk config export [{pool <pool-name>| pool <pool-name>
device-group <device-group-name>}] output-file <file-name>
```

Export a vdisk configuration to a file pathname. Optionally specify a specific pool or device-group configuration for export. Role required: admin, limited-admin.

```
vdisk config import [{pool <pool-name>| pool <pool-name>
device-group <device-group-name>}] [check-only] [skip-
initiators] [retain-serial-numbers] [on-error {continue |
stop}] input-file <filename>
```

Import a vdisk configuration that was exported to a file for migration to another location. Role required: admin, limited-admin.

Argument definitions

device-group-name

A virtual disk device group name of up to 32 characters.

device-name

The name of the virtual disk device.

pool-name

A virtual disk pool name of up to 32 characters.

vdisk-user

An authorized virtual disk user who is associated (registered) with a virtual disk pool. This user may manage all virtual disk objects that are associated with the pool.

vdisk device

Manage individual virtual disk devices. A virtual disk device is a virtualized hard disk drive that has the characteristics of a physical hard disk drive: heads, cylinders, and sectors-per-track.

```
vdisk device create [count count] capacity n {MiB|GiB|TiB|PiB|sectors} pool pool-name device-group device-group-name
```

Creates a device of at least, but not necessarily the exact size specified. The only time a device of the exact size specified is created is specified is when the capacity is specified in sectors with the `capacity n sectors` option.

```
vdisk device create [count count] heads head-count cylinders cylinder-count sectors-per-track sector-count pool pool-name device-group device-group-name
```

Add one or more new virtual disk devices. You can specify the disk size either by entering a value *n* and a unit of size, or by specifying the physical characteristics of the virtual disk. Not every possible geometry is compatible with VMAX, and not every device size can be represented as a valid geometry. Role required: admin, limited-admin.

```
vdisk device destroy device-name [destroy-static-images {yes|no}]
```

Delete a device and optionally delete its static images.

If you do not delete the static images, they become detached from (no longer associated with) the device. Use the [vdisk static-image](#) commands to attach and detach static images. Role required: admin.

Example 257

```
vdisk device create count 2 capacity 10 GiB pool p1 device-group dg1
Created VDISK device "vdisk-dev1", WWN: 60:02:18:80:00:00:00:63:05:1C:1B:02:D0:00:00
Created VDISK device "vdisk-dev2", WWN: 60:02:18:80:00:00:00:63:05:1C:1B:02:D0:00:01
2 VDISK devices created.
```

```
vdisk device modify device-name state {read-write|read-only|not-ready}
```

Modify the state of the specified vdisk device. Role required: admin, limited-admin.

```
vdisk device overwrite device-name [source-device device-name | source-pool pool-name source-device-group device-group-name source-static-image static-image-name
```

Overwrite a virtual disk device from a static-image. This command option destroys the existing data on the *device-name*. Role required: admin, limited-admin.

```
vdisk device show detailed [wwn <wwn-name>] | [<device-spec>]
[{{pool <pool-name> | pool <pool-name> device-group <device-group-name>}}
```

Show detailed information about all or selected virtual disk devices. All users may run this command option.

```
vdisk device show list [wwn <wwn-name>] | [<device-spec>]
[pool <pool-name> | pool <pool-name> device-group <device-
group-name>]
```

List all or selected virtual disk devices. All users may run this command option.

Argument definitions

capacity *n* [{MiB|GiB|TiB|PiB}]

The capacity limit of the vdisk device. The default units are GiB. Enter a value and optionally specify the units. The capacity must be between 1 GiB and 4 TiB.

count *count*

The number of vdisk devices to create. The maximum is 2048. The default is 1.

cylinders *cylinder-count*

The number of cylinders that define the disk geometry, which is used to calculate the disk capacity.

destroy-static-images {yes|no}

Whether to delete the static images of the specified device.

device-group-name

A virtual disk device group name of up to 32 characters.

device-name

The name of the virtual disk device.

heads *head-count*

The number of heads that define the disk geometry, which is used to calculate the disk capacity.

pool-name

A virtual disk pool name of up to 32 characters.

sectors-per-track *sector-count*

The number of sectors per track that defines the virtual disk device geometry, which is used to calculate the disk capacity.

static-image-name

The name of a static image.

wwn-name

A case-insensitive 30-byte string containing the hexadecimal values of the WWN. Delimiting the string with ":" characters is optional.

vdisk device-group

Manage virtual disk device groups.

A virtual disk device-group is a second-level container in a pool. It contains one or more virtual disk devices. It is represented as a subdirectory in the MTree of a virtual disk pool. The name space for a device-group name is limited to the MTree of a single vdisk pool.

Note

A device-group is not an access group.

A device group may contain:

- Devices
- Static images

`vdisk device-group create device-group-name pool pool-name`
 Create a device-group in a local pool. Role required: admin, limited-admin.

`vdisk device-group destroy device-group-name pool pool-name`
 Destroy a device-group, including all of its devices and data. Role required: admin.

`vdisk device-group rename src-device-group-name destination-group-name pool pool-name`
 Rename the device-group *src-device-group-name* to *destination-group-name*. Role required: admin, limited-admin.

`vdisk device-group show detailed [device-group-spec] [pool pool-spec]`
 Show detailed information about all or selected device-groups. All users may run this command option.

`vdisk device-group show list [device-group-spec]`
 List all or selected device groups. The output shows the name of the pool that contains these device-groups and the number of devices in each device-group. All users may run this command option.

Argument definitions

device-group-name

A virtual disk device group name of up to 32 characters.

device-group-spec

A list of virtual disk device groups that uses wildcards, such as “dg*”.

pool-name

A virtual disk pool name of up to 32 characters.

pool-spec

A list of virtual disk pools that uses a wildcard, such as “vpool*”.

vdisk disable

`vdisk disable`
 Disable the vdisk service. Role required: admin, limited-admin.

vdisk enable

`vdisk enable`
 Enable the vdisk service. Role required: admin, limited-admin.

vdisk group

Manage access between virtual devices and initiators.

Note

Use the `scsitarget group` commands to create, rename, or destroy virtual disk groups.

```
vdisk group add group-name initiator initiator-spec
```

Add an initiator to a virtual disk access group. Role required: admin, limited-admin.

```
vdisk group add group-name {device device-spec | pool pool-name
device-group device-group-name [device device-spec]} [lun lun]
[primary-endpoint {all | none | endpoint-list}] [secondary-
endpoint {all | none | endpoint-list}]
```

Add virtual disk devices to a virtual disk access group. Role required: admin, limited-admin.

```
vdisk group del group-name {device device-spec | initiator
initiator-spec | pool pool-name device-group device-group-name
[device device-spec]}
```

Remove virtual disk devices from a virtual disk access group. Role required: admin, limited-admin.

```
vdisk group modify group-name {device device-spec | pool pool-
name device-group device-group-name [device device-spec]} [lun
lun] [primary-endpoint {all | none | endpoint-list}]
[secondary-endpoint {all | none | endpoint-list}]
```

Modify the virtual disk device attributes in a virtual disk access group. Role required: admin, limited-admin.

```
vdisk group use group-name {device device-spec | pool pool-name
device-group device-group [device device-spec]} {primary |
secondary}
```

Switch the in-use endpoints list for one or more devices in a virtual disk access group between the primary port and the secondary endpoints. If you do not want to operate on all of the devices in the *device-group*, filter the devices by providing a *device-spec*. Role required: admin, limited-admin.

Argument definitions

device-group-name

A virtual disk device group name of up to 32 characters.

device *device-spec*

A list of devices that uses wildcards, such as “vdisk-dev*”.

group-spec

A list of virtual disk access groups that uses a wildcard, such as “group*”.

endpoint-list

A list of endpoints (logical names for target ports on the Data domain system).

group-name

Name of an access group.

group-spec

A list of virtual disk access groups that uses a wildcard, such as “group*”.

initiator initiator-spec

A list of initiators attached to the Data Domain system for the virtual disk service that uses a wildcard, such as “init1*”.

lun lun

A logical unit identified by a number. These are virtual disk devices exported from the Data Domain system.

pool-name

A virtual disk pool name of up to 32 characters.

vdisk pool

Manage the virtual disk pool.

A virtual disk pool is the highest-level container for virtual disk objects. It corresponds to a managed tree (MTree) on a Data Domain system.

Pools contain these lower-level objects:

- Device groups
- Devices
- Static images

```
vdisk pool create pool-name user user-name
```

Create a pool and its MTree, and assign an existing virtual disk user to the new pool. Role required: admin, limited-admin.

```
vdisk pool destroy pool-name
```

Destroy a pool, including all of its devices and data. This command cannot destroy pools that are replication destinations. Role required: admin.

```
vdisk pool modify pool-name user user-name
```

Assign an existing virtual disk user to an existing pool. Role required: admin, limited-admin.

```
vdisk pool register pool-name user vdisk-user
```

Register an existing pool or replica virtual disk MTree on a Data Domain system configured as a replication destination to vdisk, and assign an existing virtual disk user to the pool. Role required: admin, limited-admin.

```
vdisk pool rename src-pool-name dst-pool-name
```

Rename a pool from *src-pool-name* to *dst-pool-name*. Role required: admin, limited-admin.

```
vdisk pool show list [pool-spec]
```

List all pools or selected pools. All users may run this command option.

```
vdisk pool show detailed [pool-spec] [user vdisk-user]
```

Show detailed information about some or all pools. All users may run this command option.

```
vdisk pool unregister pool-name user vdisk-user
```

Unregister an existing pool or MTree from vdisk, making all the data associated with the pool inaccessible. This command requires the user to enter a password. Role required: admin, limited-admin.

Argument definitions***pool-name***

A virtual disk pool name of up to 32 characters.

pool-spec

A list of virtual disk pools that uses a wildcard, such as “vpool*”.

user-name

A Data Domain system user name with a specified role, such as backup operator.

vdisk-user

An authorized virtual disk user who is associated (registered) with a virtual disk pool. This user may manage all virtual disk objects that are associated with the pool.

vdisk property

Set or remove properties for pools, device groups, devices, and static images.

```
vdisk property reset object-name name object-type pool {all |
property-name name}
vdisk property reset object-name name object-type device-group
pool pool-name {all | property-name name}
vdisk property reset object-name name object-type device {all |
property-name name}
vdisk property reset object-name name object-type static-image
{device device-name | device-group device-group-name pool pool-
name}{all | property-name name}
```

Reset properties for a virtual disk object. Role required: admin, limited-admin.

```
vdisk property set object-name name object-type pool property-
name name property-value value
vdisk property set object-name name object-type device-group
pool pool-name property-name name property-value value
vdisk property set object-name name object-type device
property-name name property-value value
vdisk property set object-name name object-type static-image
{device device-name | device-group device-group-name pool pool-
name} property-name name property-value value
```

Set key-value pair properties for a virtual disk object. Role required: admin, limited-admin.

Example 258

```
# vdisk property set object-name pool_3 object-type pool property-name Department
property-value HR
VDISK property set for pool "pool_3"
```

Argument definitions***device-group-name***

A virtual disk device group name of up to 32 characters.

device-name

The name of the virtual disk device.

object-name *name*

A virtual disk object name. Virtual disk objects include pools, device groups, devices, and static images.

pool-name

A virtual disk pool name of up to 32 characters.

property-name *name*

A property for a virtual disk object, which can be used to identify the object. For example, a virtual disk pool named `pool-1` might have a property `department` with the value `HR`.

property-value *value*

A value for a virtual disk object property.

vdisk reset

Reset detailed virtual disk statistics. Role required: admin, limited-admin.

```
vdisk reset detailed-stats
```

vdisk show

Display information about virtual disk configuration limits and I/O statistics.

```
vdisk show config
```

Show the vdisk configuration limits. The limits displayed scale based on the Data Domain system model. Role required: admin, limited-admin.

The maximum values are system-wide limits inherited from the MTree limit. The command adjusts the output based on MTrees used by other protocols. The maximum number of pools supported ranges from 100 to 256, depending on the specific model of Data Domain system.

Example 259 DD4500 vdisk limits

```
vdisk show config
```

Name	Current	Maximum
Pools	0	127
Device-groups per pool		1024
Device-groups	0	130048
Devices	0	2048
Static images		Unlimited

Example 260 DD9500 vdisk limits

```
vdisk show config
```

Name	Current	Maximum
Pools	0	255
Device-groups per pool		1024
Device-groups	0	262143
Devices	0	2048

Example 260 DD9500 vdisk limits (continued)

```
Static images                               Unlimited
-----
```

```
vdisk show detailed-stats
```

Display detailed statistics. Role required: admin, limited-admin.

```
vdisk show stats [{pool pool-name | pool pool-name device-group
devgrp-spec}] [device device-spec] [interval interval] [count
count]
```

Periodically list I/O statistics for one or more virtual disk devices. If no pools or device-groups are specified, statistics for all virtual disk devices on the system are displayed. Specifying *count* sets the number of iterations to display. Specifying *interval* sets the amount of time between iterations. The possible combinations of *count* and *interval* create the following results:

- Neither *count* nor *interval* is specified: The system displays a single iteration of statistics.
- Both *count* and *interval* are specified: The system displays *count* number of iterations at the specified *interval*.
- If *count* is specified, but *interval* is not: The system displays *count* number of iterations at a default interval of two seconds.
- If *interval* is specified, but *count* is not: The system displays vdisk stats at the specified *interval* until the command is terminated manually.

Role required: admin, limited-admin.

Example 261

```
vdisk show stats interval 2 count 3
Start Time: 09/03 14:09:53
Interval: 2
Device      Ops/s    Read    Read    Write    Write
           -----  -----  -----  -----  -----
           Ops/s    KiB/s   Ops/s   KiB/s   Ops/s   KiB/s
vdisk-dev1      0         0         0         0         0
vdisk-dev2      0         0         0         0         0
vdisk-dev3      0         0         0         0         0
vdisk-dev4      0         0         0         0         0
vdisk-dev5      0         0         0         0         0
vdisk-dev6      0         0         0         0         0
-----
Start Time: 09/03 14:09:55
Interval: 2
Device      Ops/s    Read    Read    Write    Write
           -----  -----  -----  -----  -----
           Ops/s    KiB/s   Ops/s   KiB/s   Ops/s   KiB/s
vdisk-dev1      0         0         0         0         0
vdisk-dev2      0         0         0         0         0
vdisk-dev3      0         0         0         0         0
vdisk-dev4      0         0         0         0         0
vdisk-dev5      0         0         0         0         0
vdisk-dev6      0         0         0         0         0
-----
Start Time: 09/03 14:09:57
Interval: 2
Device      Ops/s    Read    Read    Write    Write
           -----  -----  -----  -----  -----
           Ops/s    KiB/s   Ops/s   KiB/s   Ops/s   KiB/s
vdisk-dev1      0         0         0         0         0
vdisk-dev2      0         0         0         0         0
vdisk-dev3      0         0         0         0         0
```

Example 261 (continued)

vdisk-dev4	0	0	0	0	0
vdisk-dev5	0	0	0	0	0
vdisk-dev6	0	0	0	0	0
-----	-----	-----	-----	-----	-----

Argument definitions**count** *count*

The number of iterations of statistics to display.

device *device-spec*

A list of devices that uses wildcards, such as “vdisk-dev*”.

device-group *devgrp-spec*

A collection of virtual disk devices.

endpoint-spec

A list of endpoints that uses a wildcard, such as “endpoint*”.

interval *interval*

A time window (waiting time) within which to show virtual disk I/O statistics for virtual disk devices

pool-name

A virtual disk pool name of up to 32 characters.

vdisk static-image

Manage static images for devices.

A static image is a point-in-time copy of data for a vdisk device. Static images are created within a device group. You can copy (but not move) static images to other device groups. When you create a static image, it has the same pre-compression size as the original vdisk device. As you create more static images, the pre-compression sizes of the static image files adds to the pre-compression size of the containing Mtree. Likewise, the compression ratio is affected by creating the static image files. This behavior means that you can set Mtree quotas based on the sizes of the devices and their associated static images.

A static image contains:

- A point-in-time copy of application data for a vdisk device.
- Additional metadata inserted by the vdisk feature.

```
vdisk static-image attach source-static-image-name source-pool
pool-name source-device-group device-group-name destination-
device device-name
```

Attach an existing static image to a specified destination device. This command fails if a static image is already attached to the specified destination device. Role required: admin, limited-admin.

The attach and detach command options let you organize static images by associating them with a specific device, device group, or pool. When a static image becomes detached, it is still available for use, but it is not visible if you run `vdisk static-image show` commands and you filter the output by device.

Attaching a static image to a device does not imply that the device is using the image, and does not alter the current set of active data. The `attach` operation only associates the static image with the device, for purposes of organizing the static images.

```
vdisk static-image copy src-static-image-name {source-device
device-name | source-pool pool-name source-device-group device-
group-name} {destination-device device-name | destination-pool
pool-name destination-device-group device-group-name}
```

Copy an existing static image to a specified destination. Role required: admin, limited-admin.

```
vdisk static-image create device src-device-name[destination-
device device-name | destination-pool pool-name destination-
device-group device-group-name]
```

Create a new static image of a device and attach the static image to the same device, to a different device, or to a specified device group in a specified pool. Role required: admin, limited-admin.

```
vdisk static-image destroy static-image-name [device device-
name | pool pool-name device-group device-group-name]
```

Delete a static image. Specify the device, pool, or device-group where the static image resides to delete a single copy if there are multiple copies of the static image. Role required: admin.

```
vdisk static-image detach static-image-name device device-name
```

Detach a static image from a device. Role required: admin, limited-admin.

```
vdisk static-image show detailed [static-image-spec] [device
device-name | pool pool-name [device-group device-group]
```

Show detailed information about all or specified static images. All users may run this command option.

```
vdisk static-image show list [static-image-spec] [device
device-name | pool pool-name [device-group device-group]]
```

List all or specified static images. All users may run this command option.

Argument definitions

device-group

A collection of virtual disk devices that you can use to manage the devices as a group. Device groups exist in virtual disk pools. Device group namespaces are limited to the virtual disk pool that contains the device group. Device group names may be up to 32 characters in length. The maximum number of device groups per pool is 1024. The maximum number of device groups per system is 5120.

device-group-name

A virtual disk device group name of up to 32 characters.

device-name

The name of the virtual disk device.

pool-name

A virtual disk pool name of up to 32 characters.

src-static-image-name

The name of a static image that is the source for a copy operation. The system generates these names automatically; use `vdisk static-image show list` to see the names.

static-image-name

The name of a static image.

static-image-spec

A list of static image names that uses a wildcard (*).

vdisk status

```
vdisk status
```

Show the status of the vdisk service. The output shows whether the vdisk service is enabled or disabled; whether the vdisk process is running; and whether the system has a license for the vdisk feature. Role required: admin, limited-admin.

vdisk trace

Manage tracing for virtual disk groups, initiators, and other components.

```
vdisk trace disable [component component-list]
```

Disable tracing for all or specified components. Role required: admin, limited-admin.

```
vdisk trace enable [component {all | component-list}] [level {all | high | medium | low}] [timeout {never | timeout-value-in-minutes}]
```

Enable tracing for all or specified components. By default, tracing applies to all components. If you specify a component, tracing is limited to that component, where applicable. Role required: admin, limited-admin.

```
vdisk trace show [component {all | component-list}]
```

Show tracing for all or specified components. All users may run this command option.

Argument definitions

component-list

Specify **all**, **default**, or **specific vdisk components from this list**: abnormal, device-io, fs-op, hba, obj_mgmt, procmon, scsi-other, scsi-req, scsitgt, sms-op, sys-mgmt, threads, and work-item. The default list is used by default.

level {all | high | medium | low}

Tracing level. The default is `medium`.

timeout-value-in-minutes

Timeout for tracing. The default timeout is 10 minutes.

vdisk user

Manages user privilege to access and perform tasks on virtual disks.

```
vdisk user assign vdisk-user
```

Let the specified users work with virtual disks. Separate each user name in *user-list* with a comma. Role required: admin, limited-admin.

```
vdisk user unassign vdisk-user
```

Revoke the permission for the specified users to work with virtual disks. Role required: admin, limited-admin.

```
vdisk user show
```

List the virtual disk users and the pools to which each user is assigned. All users may run this command option.

Argument definitions

vdisk-user

An authorized virtual disk user who is associated (registered) with a virtual disk pool. This user may manage all virtual disk objects that are associated with the pool.

vdisk

CHAPTER 45

vtl

Data Domain Virtual Tape Library (VTL) is a licensed software option that enables backup applications to connect to and manage a DD system running Extended Retention as a virtual tape library.

VTL pools are MTree-based (as of DD OS 5.2). Multiple MTrees let you more closely configure DD OS for data management. MTree-based pools allow MTree replication to be used instead of directory replication. Existing pools are backward compatible. You may create additional backward-compatible pools as needed. VTL pool-based replication is performed using MTree replication for MTree pools, and directory replication for backward-compatible pools. MTree-specific attributes can be applied to each VTL pool individually and include snapshots and snapshot schedules, compression information, and migration policies for Extended Retention.

The recommended number of concurrent virtual tape drive instances is platform-dependent, as is the recommended number of streams between a DD system and a backup server. This number is system-wide and includes streams from all sources, such as VTL, NFS, and CIFS. For details on the recommended number of tape drives and data streams, see the *Data Domain Operating System Administration Guide*.

Note

VTL does not protect virtual tapes from a `filesys destroy`, which will delete all virtual tapes.

This chapter contains the following topics:

• vtl change history	457
• vtl add	457
• vtl cap	457
• vtl config	458
• vtl debug	460
• vtl del	462
• vtl disable	462
• vtl drive	462
• vtl enable	464
• vtl export	464
• vtl group	464
• vtl import	466
• vtl option	468
• vtl pool	469
• vtl readahead	471
• vtl rename	471
• vtl reset	471
• vtl show	472
• vtl slot	473
• vtl status	473
• vtl tape	473

vti

vtl change history

Removed commands in DD OS 6.1.2

- `vtl initiator`
- `vtl port`

vtl add

```
vtl add vtl [model model] [slots num-slots] [caps num-caps]
```

Add a tape library. VTL supports a maximum of 64 libraries per DD system (that is, 64 VTL instances on each DD system). Role required: admin, limited-admin.

Argument Definitions

caps num-caps

The number of cartridge-access ports. The default is zero (0), and the maximum is 100 per library or 1000 per system.

model model

The name of the tape library model. See the Data Domain technical note for the model name that corresponds with your backup software.

slots num-slots

The number of slots in the library. You cannot add more drives than the number of configured slots. The maximum number of slots for all VTLs on a DD system is 32,000. The default is 20 slots.

vtl

The name of the particular virtual tape library.

vtl cap

```
vtl cap add vtl [count num-caps]
```

Add cartridge access ports (CAPs) to a virtual tape library (VTL). The total number of CAPs cannot exceed 100 per library or 1000 per system. Role required: admin, limited-admin.

```
vtl cap del vtl [count num-to-del]
```

Delete *num-to-del* CAPs from a VTL. The CAPs are deleted from the end. Role required: admin, limited-admin.

Example 262

To delete CAPs 8-10 on a VTL with 10 CAPs:

```
# vtl cap del vtl1 count 3
```

Argument Definitions

count num-caps

The number of cartridge-access ports to add. The default is 1.

count *num-to-del*

The number of objects to delete. The default is 1.

vtl

The name of the particular virtual tape library.

vtl config

```
vtl config export [vtl vtl] output-file filename
```

Export a VTL configuration to a file pathname. Role required: admin, limited-admin.

```
vtl config import [vtl vtl] [check-only] [skip-initiators]
[retain-serial-numbers] [on-error {continue | stop}] input-file
filename
```

Import a VTL configuration from a file pathname. Role required: admin, limited-admin.

Argument Definitions

check-only

This option:

- Uses the schema to validate the XML Configuration File.
- Validates the names of the following:
 - groups
 - initiators
 - endpoints
 - devices (changers, drives)
- Checks the format of the initiator system name.
- Checks whether the initiator_address_method element value belongs to one of the following:
 - SCSITGTD_INITIATOR_ADDRESS_METHOD_UNKNOWN
 - SCSITGTD_INITIATOR_ADDRESS_METHOD_AUTO
 - SCSITGTD_INITIATOR_ADDRESS_METHOD_VSA
- Checks whether the initiator_transport and endpoint_transport elements values belong to one of the following:
 - SCSITGTD_TRANSPORT_UNKNOWN
 - SCSITGTD_TRANSPORT_FC
 - SCSITGTD_TRANSPORT_FCOE
 - SCSITGTD_TRANSPORT_ISCSI
 - SCSITGTD_TRANSPORT_DUMMY
 - SCSITGTD_TRANSPORT_ALL
- Checks that the values of the following elements are BOOLEAN values (0,1 which mean FALSE and TRUE, respectively.)
 - endpoint_enabled_status
 - endpoint_online_status

- `auto_offline_option` (global)
- `auto_eject_option` (global)
- `vtl_auto_eject_option`
- `vtl_auto_offline_option`
- Validates the drive numbers.
 - Checks for repeated occurrences of the drive number.
 - Checks to make sure that the drive number does not exceed the maximum drive number allowed on the Data Domain system.
- Makes sure that the value of the VTL barcode length is appropriate, based on the model of the library.
- Does not commit the transactions or does not import any of the VTL configuration.
- When the `retain-serial-numbers` option is used, checks for the following:
 - whether the Data Domain system on which the `vtl config import retain-serial-numbers` is being used already has some VTL devices. If yes, it gives you an error.
 - validates the devices serial numbers.

input-file *filename*

The input file. Note that:

- The filename will be automatically appended with an `.xml` extension and stored in the `/ddvar/etc/vtl_configuration_files` directory.
- An `.xml` extension can also be provided explicitly. Any other extension will cause an error.

on-error {continue | stop}

Indicates what to do when an error occurs. For `stop`, the command stops, and all of the VTL configurations imported prior to the error remains, but no additional configurations are imported.

For `continue`, the action depends on the item being modified:

- Groups
 - If an error occurs while creating a group, and the group already exists, the command continues to create the next group.
 - For any other errors, the process stops.
- Endpoints
 - If an error occurs while renaming an endpoint, the command continues to configure the next endpoint.
- Initiators
 - If an error occurs while renaming an initiator, or setting an initiator alias, the command continues to configure the next initiator.
 - If an error occurs while adding an initiator to a group, the command continues to configure the next initiator.
- VTL-specific library options

- If an error occurs while configuring any of the options, the command continues to configure the next option.
- Devices
 - Changers
 - If an error occurs while creating a changer, the command continues to configure the next VTL.
 - If an error occurs while adding a changer to a group, the command continues to add the changer to other groups.
 - Drives
 - If an error occurs while adding a drive, the command continues to add the next drive.
 - If an error occurs while adding a drive to a group, the command continues to add the drive to other groups.
- Options
 - If an error occurs while enabling/disabling a VTL option, the command continues to configure the next option.

output-file *filename*

The output file. Note that:

- The filename will be automatically appended with an .xml extension and stored in the /ddvar/etc/vtl_configuration_files directory.
- An .xml extension can also be provided explicitly. Any other extension will cause an error.

retain-serial-numbers

Preserves serial numbers while creating devices on a DD system, but only when there are no pre-existing devices on that DD system. If the serial number of a device is changed in the XML configuration file, then the vdev_id of that device should also be changed to an appropriate value, because the serial number of a device is dependent on the vdev_id.

skip-initiators

Indicates:

- Skip renaming initiators, if initiators with the same system names already exist.
- Skip setting initiator aliases.
- Skip adding initiators to groups.

vtl

The name of the particular virtual tape library.

vtl debug

```
vtl debug disable [component {all | user | default | component-list}]
```

Disable debug functionality of the specified components. Role required: admin, limited-admin.

```
vtl debug enable [component {all | user | default | component-
list}] [level {high | medium | low}] [timeout {never | timeout-
value-in-minutes}]
```

Enable debug functionality for the specified components in persistent mode or for a specified `timeout` period (in minutes) at a specified debug level. Role required: `admin`, `limited-admin`.

```
vtl debug show [component {all | user | default | component-
list}]
```

Show specified components, or all components, running debug functionality. Role required: `admin`, `limited-admin`, `security`, `user`, `backup-operator`, `none`.

Argument Definitions

component {all | user | default | *component-list*}

The VTL debugging components. If you want to list them, you can include one or more of the following:

```
vhba
scst
fc
ddcl
vtc
vmc
vtlprocess
group
vscsi
vtlsm
vtc_readahead
info_cache
persistent_reservations
master_client
master_server
worker_client
worker_server
vdev_thread
registry
misc
```

Note

Components `master_client`, `master_server`, `worker_client`, and `worker_server` are used only for GDA, which is no longer supported as of 5.4.

level {high | medium | low}

The degree of VTL debugging verbosity.

timeout {never | *timeout-value-in-minutes*}

Determines the length of time (in minutes, if specified) that debugging should remain enabled for the specified components.

vtl del

```
vtl del vtl
```

Remove an existing VTL. Any tapes loaded into the library when the library is deleted are not destroyed. Instead, tapes are placed back into the virtual tape vault. Role required: admin.

Argument Definitions**vtl**

The name of the particular virtual tape library.

vtl disable

```
vtl disable
```

Close all libraries and shut down the VTL process. Role required: admin, limited-admin.

vtl drive

```
vtl drive add vtl [count num-drives] [model model]
```

Add drives to a VTL. Drives are added by starting with drive number 1 and scanning for logical unit address gaps left by `vtl drive del`. When the gaps are filled, the drives are appended to the end of the library. The number of slots within a library cannot be fewer than the number of drives in the library. If an attempt is made to add more drives than the current number of slots, the system automatically adds the additional slots required. Be aware that you cannot mix drive models within the same library. Role required: admin, limited-admin.

```
vtl drive del vtl drive drive-number [count num-to-del]
```

Delete virtual drives from a VTL. Any drive can be deleted, which means there can be gaps in the drive list. This may cause issues with some applications. Role required: admin, limited-admin.

```
vtl drive show {serial-number serial-number | vtl vtl [drive  
{drive-list}]}
```

View details of VTL drives. Role required: admin, limited-admin, security, user, backup-operator, none.

Output Definitions**Location**

Standard format location of library or drive.

Serial #

Drive serial number.

Vendor

Drive vendor identification.

Product

Drive product identification.

Product revision

Drive product revision.

Status

Drive status.

Barcode

Barcode of loaded tape.

Pool

Pool of loaded tape.

Previous Slot

Previous slot of loaded tape.

Device

SCSI device ID.

Persistent Reservation

Persistent reservation information.

Access Groups

Fibre Channel access groups for device.

Argument Definitions**count *num-drives***

The number of drives to add. The default is 1.

count *num-to-del*

The number of objects to delete. The default is 1.

drive *drive-list*

The list of drives.

drive *drive-number*

The number of the VTL drive.

model *model*

The name of the tape library model. See the Data Domain technical note for the model name that corresponds with your backup software.

serial-number *serial-number*

The serial number.

vtl

The name of the particular virtual tape library.

vtl enable

```
vtl enable
```

Enable the VTL subsystem. Before VTL can be enabled:

- You must have at least one Fibre Channel (FC) interface card installed on your DD system. VTL communicates between a backup server and a DD system through an FC interface.
- You must have previously enabled the file system and scsitarget features.
- You must have set the record (block) size for the backup software on the application host; the minimum is 64 KiB or larger. Changing the block size after the initial configuration may render unreadable any data written in the original size.

Note

VTL Fibre Channel operation is expected to be interrupted when VTL Fibre Channel endpoints failover. You may need to perform discovery (that is, operating system discovery and configuration of VTL devices) on the initiators using the affected Fibre Channel endpoint. You should expect to re-start active backup and restore operations.

Role required: admin, limited-admin.

vtl export

```
vtl export vtl {slot | drive | cap} address [count count]
```

Remove tapes from a slot, drive, or cartridge-access port (CAP) and send them to the vault. Role required: admin, limited-admin, backup-operator.

Argument Definitions

address

The address.

count *count*

The number of tapes.

vtl

The name of the particular virtual tape library.

vtl group

```
vtl group add group-name initiator initiator-alias-or -WWPN
```

Add an initiator alias or world-wide port name to the specified VTL access group. Role required: admin, limited-admin.

```
vtl group add group-name vtl vtl-name {all | changer | drive
drive-list} [lun lun] [primary-port {all | none | port-list}]
[secondary-port {all | none | port-list}]
```

Add a changer or drives to the specified VTL access group. You can add a changer or drive, optionally starting at a given logical unit number (LUN). You can optionally specify primary and secondary DD system VTL port lists. By default, the port lists contain all DD system VTL ports. Role required: admin, limited-admin.

```
vtl group create group-name
```


Create a VTL access group with the specified name. After the group is created, VTL devices (changer or drive) and initiators may then be added to the group. Role required: admin, limited-admin.

```
vtl group del group-name initiator initiator-alias-or-WWPN
```

Remove an initiator alias or world-wide port name from the specified VTL access group. Role required: admin, limited-admin.

```
vtl group del group-name vtl vtl-name {all | changer | drive drive-list}
```

Remove one or more devices from an access group. This immediately removes access from the specified initiator to the VTL devices within the group. Role required: admin, limited-admin.

```
vtl group destroy group-name
```

Remove the specified empty VTL access group. Before you can destroy a group, run `vtl group del` to remove the initiators and devices from the group. Role required: admin, limited-admin.

```
vtl group modify group-name vtl vtl-name {all | changer [lun lun] | drive drive [lun lun]} [primary-port {all | none | port-list}] [secondary-port {all | none | port-list}]
```

Modify an access group without removing and replacing devices or initiators in the group. You can use this command to change LUN assignments and primary and secondary port assignments. (Clients can access only selected LUNs from a DD system.) The main purpose of this command is to change group port assignments. VTL group changes may require the media server to rescan the SCSI bus, or you can reset the link with `scsitarget endpoint connection-reset`. Role required: admin, limited-admin.

```
vtl group rename src-group-name dst-group-name
```

Rename a VTL access group. The *dst-group-name* must not already exist. Be aware that this does not interrupt active sessions. Role required: admin, limited-admin.

```
vtl group show [ all | vtl vtl | group-name ]
```

Show information about VTL access groups. Role required: admin, limited-admin, security, user, backup-operator, none.

```
vtl group use group-name [vtl vtl-name {all | changer | drive drive-list}] {primary | secondary}
```

Switch ports in use for the specified changer in a group or library to the primary or secondary port list for the specified changer or drives. This immediately changes the access path to the primary or secondary port for the selected VTL components in an access group. When the path is restored, this will return the group to its primary port list. After you apply a group to new VTL ports, you may need to rescan the media server's SCSI bus. Also, a backup application may need to rescan available SCSI devices. This interrupts any current access to the specified group and is intended to be used during path failures. To return a group to the primary port list after the path is repaired, run `vtl group use primary`. Role required: admin, limited-admin.

Argument Definitions

drive *drive-list*

The list of drives.

dst-group-name

The name of the destination group.

group-name

The VTL group name, which must follow these rules:

- It must be unique
- It can contain only the characters 0-9, a-z, A-Z, underscore, and hyphen.
- It cannot exceed 256 characters.
- It cannot be a reserved name: TapeServer, default, all, and summary.
- A maximum of 2,048 groups is allowed.

initiator initiator-alias-or-WWPN

The initiator alias or world-wide port name.

lun *lun*

The device address to pass to the initiator. The maximum logical unit number (LUN) is 16383. A LUN must be unique within a group, but does not have to be unique across the system. LUNs for VTL devices within a group must start with zero (0) and be contiguous numbers.

port *port-list*

Includes a comma-separated list of DD system VTL ports. You can specify port names as a range separated by a hyphen (-). The ports must already exist. For multiple ports, separate each name with a comma, and enclose the list with double quotes.

primary-port

The primary VTL ports on which the devices are visible. By default, or if you specify `all`, the VTL devices are visible on all ports. Specify `none` if the devices should not be visible on any ports.

secondary-port

The secondary VTL ports on which devices are visible to `vtl group use secondary`. By default, the devices are visible on all ports. The secondary port list supports path redundancy.

src-group-name

The name of the source group.

vtl

The name of the particular virtual tape library.

vtl import

```
vtl import vtl barcode barcode [count count] [pool pool]
[element {drive | cap | slot}] [address addr]
```

Move tapes from the vault into a slot, drive, or CAP (cartridge access port). Use `vtl tape show` to display the total number of slots for a VTL and to view which slots are currently used. Use commands from the backup server to move VTL tapes to and from drives. Although `vtl import` can move tapes into tape drives, backup software commands from the backup server are more frequently used to move VTL tapes to and from drives. The default address is 1, the default element is slot, and the default pool is Default. If no address is specified, the first free slot available is used. For example if slots 1 through 4 are occupied or reserved, the address used will be 5. If the address you specify is already in use, the first free slot that is larger than the address specified is used.

The number of tapes that can be imported at one time is limited by:

- The number of empty slots. You cannot import more tapes than the number of currently empty slots.
- The number of slots that are empty and not reserved for a tape currently in a drive.
- If a tape is in a drive and the tape origin is known to be a slot, the slot is reserved.
- If a tape is in a drive and the tape origin is unknown (slot or CAP), a slot is reserved.
- A tape that is known to have come from a CAP and that is in a drive does not get a reserved slot. (The tape returns to the CAP when removed from the drive.)

The number of tapes that can be imported equals:

- The number of empty slots.
- The number of tapes that came from slots.
- The number of tapes of unknown origin.

Role required: admin, limited-admin, backup-operator.

The following two commands are equivalent:

```
# vtl import VTL1 barcode TST010L1 count 5
```

```
# vtl import VTL1 barcode TST010L1 count 5 element slot
address 1
```

Argument Definitions

address

The address.

barcode *barcode*

An eight-character virtual tape identifier. The first six characters are numbers or uppercase letters (0-9, A-Z). The last two characters are the tape code for the supported tape type: L1 (LTO-1, 100 GiB, the default capacity), LA (LTO-1, 50 GiB), LB (LTO-1, 30 GiB), LC (LTO-1, 10 GiB), L2 (LTO-2, default capacity of 200 GiB), L3 (LTO-3, default capacity of 400 GiB), L4 (LTO-4, default capacity of 800 GiB), L5 (LTO-5, default capacity of 1.5 TiB).

The default capacities are used if you do not specify the *capacity* argument when creating the tape cartridge. If you do specify a capacity, it will override the two-character tag.

When using `count` and `barcode` together, use a wild card character in the barcode to make the count valid. An asterisk matches any character in that position and all other positions. A question mark matches any character in that position.

Note

L1, LA, LB and LC tapes cannot be written on LTO-3 tape drives. L2 and L3 tapes cannot be read on LTO-1 tape drives. Also, LTO-4 will not read L2 tapes (in addition to the LA-L1 tapes).

count *count*

The number of tapes.

element

The destination element.

pool *pool*

The name of the pool. This argument is required if tapes are in a pool.

vtl

The name of the particular virtual tape library.

vtl option

```
vtl option disable option name [vtl vtl ]
```

Disable a VTL option. Optionally, you can do this only for the specified VTL. Role required: admin, limited-admin.

```
vtl option enable option name [vtl vtl ]
```

Enable a VTL option. Optionally, you can do this only for the specified VTL. Role required: admin, limited-admin.

```
vtl option reset option name [vtl vtl ]
```

Reset a VTL option to its default value. Optionally, you can do this only for the specified VTL. Role required: admin, limited-admin.

```
vtl option set option name value [vtl vtl ]
```

Set an option and value. Optionally, you can do this only for the specified VTL. Role required: admin, limited-admin.

```
vtl option show {option name | all} [vtl vtl]
```

Show settings for a specific option, all VTL options, or only for the specified VTL. Also, lists any serial-number-prefixes that are different from the default values. Role required: admin, limited-admin, security, user, backup-operator, none.

```
# vtl option show all
```

```
Global Options:
```

Name	Value
auto-eject	disabled
auto-offline	disabled
barcode-length	8
serial-number-prefix	798612

```
Options for library: lib1
```

Name	Value
serial-number-prefix	6A5690

Argument Definitions

Values for *option name* are:

auto-eject

If enabled, tapes placed into CAPs are automatically ejected to the vault.

auto-offline

If enabled, tapes being moved from a drive causes the drive to be automatically taken offline and unloaded unless the `prevent` bit is set for the drive.

barcode-length

Allows you to explicitly set the length – to either 6 or 8 – of the tape barcode that the library will report to the initiator/client.

By default (when `barcode-length` is not set), the library reports the length of the barcode depending on the type of library. For example, if tape AAA001L3 is put in an L180, DDVTL, or RESTORER-L180 library, the library will report this tape to the initiator/client as AAA001. If the same tape is placed in a TS3500, I2000, or I6000 library, the library will report this tape to the initiator/client as AAA001L3.

However, if you do specify the `barcode-length`, for example, `vtl option set barcode-length 6 vtl my_ts3500_library` on a TS3500 library, then the barcode will be reported to the initiator/client as AAA001, which is the same length as for an L180 library.

loop-id - deprecated

The Fibre Channel loop ID: 1-26. This value has been deprecated and will be removed in future releases. To set the loop ID on a Data Domain system, enter a number between 1 and 26 in the *value* field of `scsitarget transport option set`.

serial-number-prefix

The prefix of the serial number, which can be modified globally or per library.

vtl pool

```
vtl pool add pool [backwards-compatibility-mode]
```

Create a VTL pool. If `backwards-compatibility-mode` is used, a pool with backwards compatibility is created in the default directory (`/backup`). It is recommended that you create backwards-compatibility pools only if you have specific requirements, for example, replication with a pre-5.2 DD OS system. Replication of backwards-compatibility-mode pools is done using directory-based replication, as in previous releases. Role required: admin, limited-admin.

```
vtl pool del pool
```

Delete a VTL pool. You must run `vtl tape del` to remove all tapes from a pool, or use `vtl tape move` to move all tapes to another pool. Role required: admin.

```
vtl pool modify <pool-name> data-movement-policy {user-managed
| age-threshold <days> | none} to-tier {cloud} cloud-unit
<unit-name>
```

Configure the data movement policy and cloud unit information for the specified VTL pool. Role required: admin, limited-admin.

```
vtl pool rename src-pool dst-pool
```

Rename a VTL pool. A pool can be renamed only if none of its tapes is in a library. Role required: admin, limited-admin.

```
vtl pool show {all | pool}
```

List all tape pools or the contents of a specific *pool*. If `all` is used, a summary of all tape pools is provided, including the storage tier unit-name, the cloud data movement policy, the state of each pool, the number of tapes, the total usage and compression for each pool, whether a pool is a replication destination, the Retention Lock status of the pool, read/write properties, and the number of tapes in the pool. Role required: admin, limited-admin, security, user, backup-operator, none.

Output Definitions

RW

Pool has normal read/write properties.

RD

Pool is a replication destination.

RO

Pool is read-only.

RLCE

Pool is Retention Lock Compliance Enabled.

RLGE

Pool is Retention Lock Governance Enabled.

RLGD

Pool is Retention Lock Governance Disabled.

BCM

Pool is in backwards-compatibility mode.

```
vtl pool upgrade-to-mtree {pool-list | all} [check-only]
```

Upgrade a VTL pool(s) to an MTree pool(s). If *pool-list* (a colon-separated list) is specified, all pools in the list are candidates for upgrade. If *all* is specified, all backwards-compatibility mode pools are upgraded. If *check-only* is specified, the precheck is run, but no upgrade is performed, so that you can plan for these changes prior to the upgrade. If no arguments are provided, a check is made to see if an upgrade is necessary or possible. If so, the upgrade is performed, which converts the specified backwards-compatibility mode pools to MTree pools. An upgrade may be run only when VTL is disabled.

A directory pool will be converted to an MTree pool only if the following prerequisites are met:

- 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 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 Data Domain system.
 - Replicate the directory pool from the second Data Domain system to a third Data Domain system.
 - Remove the second and third Data Domain systems from the managing Data Domain system's Data Domain network.
 - On any of the systems running DD OS 5.5, run the `vtl pool upgrade-to-mtree` command.

See the *Data Domain Operating System Administration Guide* for more information about upgrading directory pools to MTree pools. Role required: admin, limited-admin.

Example 263

```
# vtl pool upgrade-to-mtree all
```

Example 264

```
# vtl pool upgrade-to-mtree old-pool check-only
```

Argument Definitions**dst-pool**

The name of the new VTL pool.

source *src-pool*

The name of the current VTL pool.

vtl readahead

```
vtl readahead reset {stats | summary}
```

Reset VTL readahead information. When VTL reads a tape file, it improves performance by reading ahead information from tape files and caching the information until needed. Role required: admin, limited-admin.

```
vtl readahead show {stats | detailed-stats | summary}
```

Display readahead information about each open tape file that has been read. Role required: admin, limited-admin, security, user, backup-operator, none.

Argument Definitions**detailed-stats**

Provides detailed statistics.

stats

Displays statistics.

summary

Shows a summary of all tapes and tape usage.

vtl rename

```
vtl rename src-vtl dst-vtl
```

Rename a virtual tape library. The source name and the destination name must differ. Role required: admin, limited-admin.

vtl reset

```
vtl reset hba - deprecated
```

This command is deprecated. Use `scsitaraget endpoint connection-reset all` instead. Role required: admin.

```
vtl reset detailed-stats
```

Reset the VTL detailed statistics. Role required: admin, limited-admin.

vtl show

```
vtl show config [vtl]
```

Show the library name and model and tape drive model for a single VTL or all VTLs. Role required: admin, limited-admin, security, user, backup-operator, none.

```
vtl show detailed-stats
```

Show a large quantity of detailed VTL statistics and information. Role required: admin, limited-admin, security, user, backup-operator, none.

```
vtl show element-address [vtl]
```

Show the following information for all VTLs, or a single VTL:

- Starting element address
- Slot count and starting address
- CAP count and starting address
- Drive count and starting address
- Changer count and starting address

Role required: admin, security, user, backup-operator, none.

```
vtl show stats [port {port-list | all}] [interval secs] [count count]
```

Show VTL I/O stats. Role required: admin, limited-admin, security, user, backup-operator, none.

```
vtl show stats vtl [drive {drive-list | changer | all}] [port {port-list | all}] [interval secs] [count count]
```

Periodically list I/O statistics for one or more VTLs. If a VTL is not specified, statistics for all VTLs on the system are displayed. Specifying *count* sets the number of iterations to display. Specifying *interval* sets the amount of time between iterations. The possible combinations of *count* and *interval* create the following results:

- Neither *count* nor *interval* is specified: The system displays a single iteration of statistics.
- Both *count* and *interval* are specified: The system displays *count* number of iterations at the specified *interval*.
- If *count* is specified, but *interval* is not: The system displays *count* number of iterations at a default interval of two seconds.
- If *interval* is specified, but *count* is not: The system displays VTL stats at the specified *interval* until the command is terminated manually.

Role required: admin, limited-admin, security, user, backup-operator, none.

Argument Definitions

count *count*

The number of tapes.

drive {*drive-list* | *changer* | *all*}

Lets you include all drives, changer, or a list of drives.

detailed-stats

Provides detailed statistics.

interval *secs*

The time interval in seconds.

port {*port-list* | all}

Lets you include all ports, or a comma-separated list of Data Domain system VTL ports. You can specify port names as a range separated by a hyphen (-). The ports must already exist. For multiple ports, separate each name with a comma, and enclose the list with double quotes.

vtl

The name of the particular virtual tape library.

vtl slot

```
vtl slot add vtl [count num-slots]
```

Add slots to a VTL. Additional slots are added to the end of the list of slots in the specified VTL. The maximum is 32,000 slots per library and 64,000 slots per system. Role required: admin, limited-admin.

```
vtl slot del vtl [count num-to-del]
```

Delete one or more slots from a VTL. Role required: admin, limited-admin.

Argument Definitions

count *num-slots*

The number of slots to add to the library. You cannot add more drives than the number of configured slots. The default is 20 slots.

count *num-to-del*

The number of slots to delete from the library.

vtl

The name of the particular virtual tape library.

vtl status

```
vtl status
```

Show the state of the VTL process. Role required: admin, limited-admin, security, user, backup-operator, none.

vtl tape

```
vtl tape add barcode [capacity capacity] [count count] [pool <pool>]
```

Add one or more virtual tapes and insert them into the vault. Optionally, add the tapes to the specified pool. Role required: admin, limited-admin.

```
vtl tape copy barcode barcode [count count] source src-pool [snapshot src-snapshot] destination dst-pool
```

Copy tapes between VTL pools. An opened writable tape in a tape drive may not be copied. Additionally, source and destination pools cannot be the same unless copying from a snapshot. If the snapshot argument is specified, tapes are copied from the snapshot of the source pool. In this case, the destination pool can be the same as the source pool. A tape in the vault or library slot/cap, or opened read-only in a tape drive, can be copied. A tape that is opened writable in a tape drive may not be copied. Role required: admin, limited-admin, backup-operator.

```
# vtl tape copy barcode AA0000LC count 100 source replica-dest
destination daily-restores
```

```
vtl tape del barcode [count count] [pool pool]
```

Delete the specified tape or one or more tapes. You cannot delete tapes that are in a VTL. Role required: admin.

NOTICE

This command deletes all data on the tapes.

```
vtl tape deselect-for-move barcode <barcode> [count <count>]
pool <pool> to-tier {cloud}
```

Deselect a specified tape for migration to the cloud. Role required: admin, limited-admin.

```
vtl tape history delete
```

Delete all VTL tape history. Role required: admin, limited-admin.

```
vtl tape history disable
```

Disable all VTL tape history. Role required: admin, limited-admin.

```
vtl tape history enable
```

Enable all VTL tape history. Role required: admin, limited-admin.

```
vtl tape history show barcode [pool pool] [start-time
MMDDhhmm[[CC]YY]] [end-time MMDDhhmm[[CC]YY]]
```

Show history of move-related events for a given tape. Role required: admin, limited-admin, security, user, backup-operator, none.

```
vtl tape history status
```

Show current state of the VTL tape history feature. Role required: admin, limited-admin, security, user, backup-operator, none.

```
vtl tape modify barcode [count count] [pool pool] retention-
lock {date | period}
```

Modify the state of retention lock of a specified tape or tapes. Change the amount of time to maintain the retention lock on the specified tape or tapes. If the volume is not mounted, the change is made immediately. Otherwise, data is synchronized first. This will fail if the file system is read-only. Role required: admin, limited-admin.

```
vtl tape modify barcode [count count] [pool pool] writeprotect
{on | off}
```

Set the write protect state of a specified tape. If the volume is not mounted, the tape file permission is changed immediately. Otherwise, outstanding writes are synchronized first. Role required: admin, limited-admin.

```
vtl tape move vtl source {slot | drive | cap} {src-address-list
| all} destination {slot | drive | cap} {dst-address-list |
auto}
```

Move one or more tapes between elements in a VTL. Values for *src-address-list* include: all, 1, 2-14, 3-5, 7-10. Values for *dst-address-list* include: 1, 2-14, 3-5, 7-10, and auto. You may specify the `auto` keyword only if moving from tapes from drives to slots. If `auto` is selected, VTL finds the previous slot the tape was in and moves it to that slot. If the slot is not empty, it moves the next available slot. Role required: admin, limited-admin, backup-operator.

```
vtl tape move barcode <barcode> [count count] source <src-pool>
destination <dst-pool>
```

Move a tape between VTL pools if it is in the vault, or in a library slot or CAP. It cannot be moved between VTL pools if the tape is open in a drive, or if it is one of the following kinds of tapes:

- Tapes open in a drive
- Tapes on a replica
- Tapes configured with Retention Lock

Role required: admin, limited-admin.

```
vtl tape recall start barcode <barcode> [count <count>] pool
<pool>
```

Recall a tape from the cloud. Role required: admin, limited-admin.

```
vtl tape select-for-move barcode <barcode> [count <count>] pool
<pool> to-tier {cloud}
```

Select a specified tape for migration to the cloud. Role required: admin, limited-admin.

```
vtl tape show {all | pool pool | vault | vtl} [cloud-unit all |
<unit-name>] [summary] [count count] [barcode barcode] [time-
display {modification | creation | retention | recalled}]
[sort-by {barcode | pool | location | state | capacity | usage
| percentfull | compression | time | modtime} [{ascending |
descending}]]
```

Display information about tapes, including modification, creation, retention, or recalled times. If `time-display` is omitted, the default is modification time for backward-compatibility-mode VTL pools. Modification times used by the system for age-based policies may differ from the last modified time displayed in the tape information sections of the GUI and CLI. This is expected behavior. If you are using Extended Retention, see the *Data Domain Operating System Administration Guide* for details on modification time. **Role required:** admin, limited-admin, security, user, backup-operator, none.

Argument Definitions

address

The address.

barcode *barcode*

An eight-character virtual tape identifier. The first six characters are numbers or uppercase letters (0-9, A-Z). The last two characters are the tape code for the supported tape type: L1 (LTO-1, 100 GiB, the default capacity), LA (LTO-1, 50 GiB), LB (LTO-1, 30 GiB), LC (LTO-1, 10 GiB), L2 (LTO-2, default capacity of 200 GiB), L3 (LTO-3, default capacity of 400 GiB), L4 (LTO-4, default capacity of 800 GiB), L5 (LTO-5, default capacity of 1.5 TiB).

The default capacities are used if you do not specify the *capacity* argument when creating the tape cartridge. If you do specify a capacity, it will override the two-character tag.

When using `count` and `barcode` together, use a wild card character in the barcode to make the count valid. An asterisk matches any character in that position and all other positions. A question mark matches any character in that position.

Note

L1, LA, LB and LC tapes cannot be written on LTO-3 tape drives. L2 and L3 tapes cannot be read on LTO-1 tape drives. Also, LTO-4 will not read L2 tapes (in addition to the LA-L1 tapes).

capacity *capacity*

The number of gibibytes (GiB) for each tape created. This value overrides default barcode capacities. The upper limit is 4,000 GiB. For best results, when data becomes obsolete (and the DD system cleaning process marks data for removal), set capacity to 100 or less for efficient reuse of DD system disk space.

GiBs equal the base-2 value of Gigabytes (GB).

count *count*

The number of tapes.

pool *pool*

The name of the pool. This argument is required if tapes are in a pool.

snapshot *src-snapshot*

A specific snapshot within a source pool.

source *src-pool*

The name of the current VTL pool.

write-protect {on | off}

Enables or disables write-protection for a tape.

cloud-unit *unit-name*

The name of the cloud unit where the VTL vault resides.

APPENDIX A

Time Zones

This appendix covers the following topics:

• Time zones overview	478
• Africa	478
• America	479
• Antarctica	480
• Asia	480
• Atlantic	481
• Australia	481
• Brazil	481
• Canada	482
• Chile	482
• Etc	482
• Europe	482
• GMT	483
• Indian (Indian Ocean)	483
• Mexico	483
• Miscellaneous	483
• Pacific	484
• US (United States)	484
• Aliases	484

Time zones overview

Time zones are used to establish your location when you initially configure your system.

Locate your time zone using the following tables.

A time zone can consist of two entries separated by a slash (/). The first entry can be a continent, nation, or region, such as Africa, the Pacific, or the United States. The second entry is the city closest to you within that area.

A time zone, and some miscellaneous entries such as GMT, Cuba, and Japan, can also be a single entry.

Examples of time zones include:

- Indiana/Indianapolis
- GMT+5
- Stockholm
- Pacific
- EasterIsland
- Japan

Africa

Table 4 African time zones

Abidjan	Accra	Addis_Ababa	Algiers	Asmara
Asmera	Bamako	Bangui	Banjul	Bissau
Blantyre	Brazzaville	Bujumbura	Cairo	Casablanca
Ceuta	Conakry	Dakar	Dar_es_Salaam	Djibouti
Douala	El_Aaiun	Freetown	Gaborone	Harare
Johannesburg	Juba	Kampala	Khartoum	Kigali
Kinshasa	Lagos	Libreville	Lome	Luanda
Lubumbashi	Lusaka	Malabo	Maputo	Maseru
Mbabane	Mogadishu	Monrovia	Nairobi	Ndjamena
Niamey	Nouakchott	Ouagadougou	Porto-Novo	Sao_Tome
Timbuktu	Tripoli	Tunis	Windhoek	

America

Table 5 American time zones

Adak	Anchorage	Anguilla	Antigua	Araguaina
Argentina/ Buenos_Aires	Argentina/ Catamarca	Argentina/ ComoRivadavia	Argentina/ Cordoba	Argentina/ Jujuy
Argentina/ La_Rioja	Argentina/ Mendoza	Argentina/ Rio_Gallegos	Argentina/ Salta	Argentina/ San_Juan
Argentina/ San_Luis	Argentina/ Tucuman	Argentina/Ushuaia	Aruba	Asuncion
Atikokan	Atka	Bahia	Bahia_Banderas	Barbados
Belem	Belize	Blanc-Sablon	Boa_Vista	Bogota
Boise	Buenos_Aires	Cambridge_Bay	Campo_Grande	Cancun
Caracas	Catamarca	Cayenne	Cayman	Chicago
Chihuahua	Coral_Harbour	Cordoba	Costa_Rica	Creston
Cuiaba	Curacao	Danmarkshavn	Dawson	Dawson_Creek
Denver	Detroit	Dominica	Edmonton	Eirunepe
El_Salvador	Ensenada	Fort_Wayne	Fortaleza	Glace_Bay
Godthab	Goose_Bay	Grand_Turk	Grenada	Guadeloupe
Guatemala	Guayaquil	Guyana	Halifax	Havana
Hermosillo	Indiana/ Indianapolis	Indiana/Knox	Indiana/ Marengo	Indiana/ Petersburg
Indiana/ Tell_City	Indiana/Vevay	Indiana/Vincennes	Indiana/ Winamac	Indianapolis
Inuvik	Iqaluit	Jamaica	Jujuy	Juneau
Kentucky/ Louisville	Kentucky/ Monticello	Knox_IN	Kralendijk	La_Paz
Lima	Los_Angeles	Louisville	Lower_Princes	Maceio
Managua	Manaus	Marigot	Martinique	Matamoros
Mazatlan	Mendoza	Menominee	Merida	Metlakatla
Mexico_City	Miquelon	Moncton	Monterrey	Montevideo
Montreal	Montserrat	Nassau	New_York	Nipigon
Nome	Noronha	North_Dakota/ Beulah	North_Dakota/ Center	North_Dakota/ New_Salem
Ojinaga	Panama	Pangnirtung	Paramaribo	Phoenix
Port-au-Prince	Port_of_Spain	Porto_Acre	Porto_Velho	Puerto_Rico

Table 5 American time zones (continued)

Rainy_River	Rankin_Inlet	Recife	Regina	Resolute
Rio_Branco	Rosario	Santa_Isabel	Santarem	Santiago
Santo_Domingo	Sao_Paulo	Scoresbysund	Shiprock	Sitka
St_Barthelemy	St_Johns	St_Kitts	St_Lucia	St_Thomas
St_Vincent	Swift_Current	Tegucigalpa	Thule	Thunder_Bay
Tijuana	Toronto	Tortola	Vancouver	Virgin
Whitehorse	Winnipeg	Yakutat	Yellowknife	

Antarctica

Table 6 Antarctic time zones

Casey	Davis	DumontDURville	Macquarie	Mawson
McMurdo	Palmer	Rothera	South_Pole	Syowa
Troll	Vostok			

Asia

Table 7 Asian time zones

Aden	Almaty	Amman	Anadyr	Aqtau
Aqtobe	Ashgabat	Ashkhabad	Baghdad	Bahrain
Baku	Bangkok	Beijing	Beirut	Bishkek
Brunei	Calcutta	Chita	Choibalsan	Chongqing
Chungking	Colombo	Dacca	Damascus	Dhaka
Dili	Dubai	Dushanbe	Gaza	Harbin
Hebron	Ho_Chi_Minh	Hong_Kong	Hovd	Irkutsk
Istanbul	Jakarta	Jayapura	Jerusalem	Kabul
Kamchatka	Karachi	Kashgar	Kathmandu	Katmandu
Khandyga	Kolkata	Krasnoyarsk	Kuala_Lumpur	Kuching
Kuwait	Macao	Macau	Magadan	Makassar
Manila	Muscat	Nicosia	Novokuznetsk	Novosibirsk
Omsk	Oral	Phnom_Penh	Pontianak	Pyongyang
Qatar	Qyzylorda	Rangoon	Riyadh	Saigon
Sakhalin	Samarkand	Seoul	Shanghai	Singapore

Table 7 Asian time zones (continued)

Srednekolymsk	Taipei	Tashkent	Tbilisi	Tehran
Tel_Aviv	Thimbu	Thimphu	Tokyo	Ujung_Pandang
Ulaanbaatar	Ulan_Bator	Urumqi	Ust-Nera	Vientiane
Vladivostok	Yakutsk	Yekaterinburg	Yerevan	

Atlantic

Table 8 Atlantic time zones

Azores	Bermuda	Canary	Cape_Verde	Faeroe
Faroe	Jan_Mayen	Madeira	Reykjavik	South_Georgia
St_Helena	Stanley			

Australia

Table 9 Australian time zones

ACT	Adelaide	Brisbane	Broken_Hill	Canberra
Currie	Darwin	Eucla	Hobart	LHI
Lindeman	Lord Howe	Melbourne	NSW	North
Perth	Queensland	South	Sydney	Tasmania
Victoria	West	Yancowinna		

Brazil

Table 10 Brazilian time zones

Acre	DeNoronha	East	West
------	-----------	------	------

Canada

Table 11 Canadian time zones

Atlantic	Central	East-Saskatchewan	Eastern
Mountain	Newfoundland	Pacific	Saskatchewan
Yukon			

Chile

Table 12 Chilean time zone

Continental	EasterIsland
-------------	--------------

Etc

Table 13 Etc time zones

GMT	GMT+0	GMT+1	GMT+2	GMT+3
GMT+4	GMT+5	GMT+6	GMT+7	GMT+8
GMT+9	GMT+10	GMT+11	GMT+12	GMT0
GMT-0	GMT-1	GMT-2	GMT-3	GMT-4
GMT-5	GMT-6	GMT-7	GMT-8	GMT-9
GMT-10	GMT-11	GMT-12	GMT-13	GMT-14
Greenwich	UCT	Universal	UTC	Zulu

Europe

Table 14 European time zones

Amsterdam	Andorra	Athens	Belfast	Belgrade
Berlin	Bratislava	Brussels	Bucharest	Budapest
Busingen	Chisinau	Copenhagen	Dublin	Gibraltar
Guernsey	Helsinki	Isle_of_Man	Istanbul	Jersey
Kaliningrad	Kiev	Lisbon	Ljubljana	London
Luxembourg	Madrid	Malta	Mariehamn	Minsk
Monaco	Moscow	Nicosia	Oslo	Paris
Podgorica	Prague	Riga	Rome	Samara

Table 14 European time zones (continued)

San_Marino	Sarajevo	Simferopol	Skopje	Sofia
Stockholm	Tallinn	Tirane	Tiraspol	Uzhgorod
Vaduz	Vatican	Vienna	Vilnius	Volgograd
Warsaw	Zagreb	Zaporozhye	Zurich	

GMT

Table 15 GMT time zones

GMT	GMT+1	GMT+2	GMT+3	GMT+4
GMT+5	GMT+6	GMT+7	GMT+8	GMT+9
GMT+10	GMT+11	GMT+12	GMT+13	GMT-1
GMT-2	GMT-3	GMT-4	GMT-5	GMT-6
GMT-7	GMT-8	GMT-9	GMT-10	GMT-11
GMT-12				

Indian (Indian Ocean)

Table 16 Indian (Indian Ocean) time zones

Antananarivo	Chagos	Christmas	Cocos	Comoro
Kerguelen	Mahe	Maldives	Mauritius	Mayotte
Reunion				

Mexico

Table 17 Mexican time zones

BajaNorte	BajaSur	General
-----------	---------	---------

Miscellaneous

Table 18 Miscellaneous time zones

Arctic/ Longyearbyen	CET	CST6CDT	Cuba	EET
Egypt	Eire	EST	EST5EDT	Factory
GB	GB-Eire	Greenwich	Hongkong	HST

Table 18 Miscellaneous time zones (continued)

Iceland	Iran	Israel	Jamaica	Japan
Kwajalein	Libya	MET	MST	MST7MDT
Navajo	NZ	NZ-CHAT	Poland	Portugal
PRC	PST8PDT	ROC	ROK	Singapore
Turkey	UCT	Universal	UTC	WET
W-SU	Zulu			

Pacific

Table 19 Pacific time zones

Apia	Auckland	Chatham	Chuuk	Easter
Efate	Enderbury	Fakaofu	Fiji	Funafuti
Galapagos	Gambier	Guadalcanal	Guam	Honolulu
Johnston	Kiritimati	Kosrae	Kwajalein	Majuro
Marquesas	Midway	Nauru	Niue	Norfolk
Noumea	Pago_Pago	Palau	Pitcairn	Pohnpei
Ponape	Port_Moresby	Rarotonga	Saipan	Samoa
Tahiti	Tarawa	Tongatapu	Truk	Wake
Wallis	Yap			

US (United States)

Table 20 US (United States) time zones

Alaska	Aleutian	Arizona	Central	East-Indiana
Eastern	Hawaii	Indiana-Starke	Michigan	Mountain
Pacific	Pacific-New	Samoa		

Aliases

GMT=Greenwich, UCT, UTC, Universal, Zulu CET=MET (Middle European Time)
 Eastern=Jamaica Mountain=Navajo